

1-2013

Managing Wind-based Electricity Generation in the Presence of Storage and Transmission Capacity

Yangfang Zhou
Carnegie Mellon University

Alan Scheller-Wolf
Carnegie Mellon University, awolf@andrew.cmu.edu

Nicola Secomandi
Carnegie Mellon University, ns7@andrew.cmu.edu

Stephen Smith
Carnegie Mellon University, ssmith@andrew.cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/tepper>

 Part of the [Business Administration, Management, and Operations Commons](#)

Managing Wind-based Electricity Generation in the Presence of Storage and Transmission Capacity

Yangfang Zhou, Alan Scheller-Wolf, Nicola Secomandi

Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh 15213
{yangfang, awolf, ns7}@andrew.cmu.edu

Stephen Smith

Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh 15213, sfs@cs.cmu.edu

Tepper Working Paper 2011-E36

October 2011; Revised: July 2012; January, 2013

Managing power generation from wind is conceptually straightforward: generate and sell as much electricity as possible when prices are positive, and do nothing otherwise. However, this strategy leads to curtailment when wind energy exceeds the transmission capacity or when prices are negative, and possible revenue dilution when current prices are low but are expected to increase in the future. Electricity storage provides a means to alleviate these problems, and also enables the purchase of electricity from the market for later resale. But the presence of storage and transmission capacity complicates the management of electricity generation from wind. Much is unknown about the management of such a generation plus storage and transmission system, for instance, how complex the optimal management of such a system is, how valuable storage is, and how relevant it is to include different factors—such as negative prices, the buying option, and future information. We answer these questions by developing and analyzing a Markov decision process model of this system.

Key words: inventory policy; Markov decision process; wind-based electricity generation; electricity storage

1. Introduction

The last decade has seen a boom in global wind-based electricity generation: since 2000, U.S. wind generation capacity has grown more than tenfold (Wiser and Bolinger 2010); similar growth has been seen in other countries as well. This global trend will probably continue for at least another decade or two, as many countries have enacted policies to promote wind energy (REN21 2010).

Managing wind energy generation is conceptually straightforward: generate and sell as much as possible when the price is positive (electricity prices can be negative), and do nothing otherwise. However, this policy may dilute a wind farm's revenue because of the following:

(i) This strategy requires curtailing wind energy when it exceeds the capacity of the transmission lines connecting wind farms (mostly remote) to electricity markets. Curtailment is a significant issue: in 2009, 17% of the wind power in the Electric Reliability Council of Texas (ERCOT)

was curtailed due to lack of transmission capacity (Wiser and Bolinger 2010). This bottleneck in transmission capacity is unlikely to disappear in the near future, as it takes much longer (about five years) to build transmission lines than to build wind farms (about one year; SECO 2011).

(ii) This strategy also leads to curtailment when electricity prices are negative. Negative prices have been observed in the New York Independent System Operator (NYISO 2011), ERCOT (ERCOT 2008), the Nordic Power Exchange (Sewalt and de Jong 2007), and the European Energy Exchange (Genoese et al. 2010, Fanone and Prokopczuk 2010). Negative prices may be caused by the high cost of ramping conventional power plants up and down: when electricity demand is low, these plant operators may try to avoid a costly shut-down by paying others to consume their excess power (Knittel and Roberts 2005, Sewalt and de Jong 2007, Genoese et al. 2010).

(iii) This strategy may sell wind energy when prices are positive, but unfavorable. For instance, wind in the U.S. tends to blow most strongly at night, when electricity prices are typically low.

These problems can be alleviated by co-locating wind farms with electricity storage facilities, for instance, grid-level electricity batteries. Storage also enables buying electricity for future resale. However, the presence of storage complicates the management of a wind farm: the operator needs to decide how much to generate, and how much to buy from the market, or to sell to the market from generation and/or inventory. It is not clear how difficult this problem is, what the dominant features of the problem are (the value of considering negative prices, versus the buying option, versus future information, versus uncertainty), or how large the value of storage might be.

We provide answers to these questions by investigating this problem of operating a wind-based electricity generation, storage, and transmission (WST) system. We model our system as a Markov decision process (MDP), analyze different control policies, and carry out numerical experiments based on price and wind models calibrated to data. We find that using a naïve policy can be significantly suboptimal, losing around 15% of the value compared to the optimal policy. This implies that this problem is non-trivial, and the value of optimization is substantial. We also find that negative prices do not matter much: the policy that ignores negative prices achieves at least 99% of the optimal value for a comprehensive range of system configurations in our experimental analysis. Along the way, we prove that this policy has a triple-threshold structure, generalizing existing ones in the literature (Nascimento and Powell 2010, Secomandi 2010b). To take advantage of negative prices, we provide a simple enhancement to this triple-threshold policy and demonstrate that it achieves 99.95% of the optimal value: the optimal policy thus appears to be consistent with the structure of this enhanced policy.

When it comes to ignoring the buying option, we find it can cause a loss of about 4% of the optimal value. We prove that the optimal policy for the model that ignores this buying option has a dual-threshold structure, the same as in Nascimento and Powell (2010) and Secomandi (2010b). In addition, we find that modeling future information and uncertainty in the optimization can be fairly important: if ignored, these cost around 8% and 11% of the optimal value, respectively.

Our experiments also quantify the monetary and energy values of storage when using our triple-threshold policy (the energy value of storage captures the change in the energy flow rate). Storage can substantially increase the monetary value of a WST system: For a typical scenario with tight transmission capacity, storage increases the monetary value of the system by 25%, of which 10% is due to reducing curtailment, 11% to time-shifting generation, and 4% to arbitrage. With less constrained transmission capacity, storage increases the monetary value of the system by 21.5%, of which 0.5% is due to reducing curtailment, 17% to time-shifting, and 4% to arbitrage.

Storage can thus greatly decrease the average wind energy curtailed (by 70% compared to the no-storage case) and increase the average energy that the system sells to the market (by 18% compared to the no-storage case) when transmission capacity is tight. However, with ample transmission capacity, adding storage may actually decrease the average *wind* energy sold: the benefit of reducing curtailment decreases, and is surpassed by the conversion loss at the storage facility.

Our results provide guidelines for practitioners to operate WST systems near optimally, help them determine which features of the problem are germane for their decision making, and also can assist them to evaluate the value of storage for wind farms. In addition, our work provides a basis for the sizing of storage facilities and transmission lines. Our structural analysis has potential implications for the development of approximate dynamic programming methods for alternative versions of our model (Nascimento and Powell 2008, 2009).

The rest of this paper is organized as follows: we review the literature in §2 and present our MDP model in §3. We discuss the framework for analyzing the relevance of many features of this problem in §4, and the framework for the valuation of storage in §5. We specify the price model and the wind energy model used in our numerical study in §6 and §7, respectively. We carry our numerical analysis in §8. We conclude and discuss future work in §9. The online appendix includes all the proofs of our structural results.

2. Literature Review

In the literature on wind-storage systems, Denholm and Sioshansi (2009) and Fertig and Apt (2011) consider the interplay of storage and transmission capacity: the former studies how to

best locate storage when transmission capacity is binding; the latter studies the optimal sizing of storage and transmission capacity. In contrast, we develop policies to manage an existing WST system. Moreover, these papers evaluate the value of storage assuming price and wind energy are deterministic processes, while our paper captures the uncertainty of both.

Other related studies examine different uses of storage for wind farms. Brown et al. (2008) focus on how to serve the load of an isolated system using wind generators and pump-hydro storage to minimize daily operating cost. Castronuovo and Lopes (2004) maximize the daily profit of a wind-hydro system in a market. Korpaas et al. (2003) consider a wind-storage system that serves load as well as trades in an electricity market. In contrast to this literature, we include transmission capacity and explore how it affects the value of storage.


Wu and Kapuscinski (2012) investigate how to curtail wind energy to minimize the total balancing cost of an electricity market (possibly in the presence of storage) from the point of view of an electricity market operator. Xi et al. (2011) optimize the use of an electricity storage facility in both the electricity energy and ancillary markets. We consider the use of co-located storage for a wind farm constrained by transmission capacity, investigating the relevance of different features of the problem from the perspective of a wind-farm manager.

Another stream of work centers on how a wind-farm manager can use storage to make better bidding decisions in a forward market. Some papers assume price is deterministic, such as Bathurst and Strbac (2003) and Costa et al. (2008); others take price uncertainty into account, such as Gonzalez et al. (2008), Löhndorf and Minner (2010), and Kim and Powell (2011). We do not consider bidding and assume any electricity offered to the market is accepted. This is realistic: many electricity markets in the U.S. treat wind generators as “must-run” in normal conditions (Wiser and Bolinger 2010); in addition, 38% of all the wind capacity developed in the U.S. in 2009 was sold through merchant/quasi-merchant projects “whose electricity sales revenue is tied to short-term contracted and/or wholesale spot electricity market prices” (Wiser and Bolinger 2010, p. 43). Furthermore, this literature ignores transmission capacity.

As electricity can be viewed as a special type of inventory, our work is also related to the literature in inventory theory (e.g. Zipkin 2000, Porteus 2002). In this literature demand is the typical source of uncertainty, but supply can also exhibit randomness. In contrast, price and wind (supply) are the sources of uncertainty in our model.

Our paper is particularly related to the literature on commodity and energy storage. Cahn (1948) introduces the classic warehouse problem, for which Charnes et al. (1966) show the optimality of a simple basestock policy. Rempala (1994) and Secomandi (2010b) extend this work to incorporate

Figure 1 The System Overview



a limit on the rate at which the commodity can be injected or withdrawn. Other related work includes Mokrian and Stephen (2006), Chen and Forsyth (2007), Boogert and de Jong (2008), Thompson et al. (2009), Lai et al. (2010a), Wu et al. (2012), and Devalkar et al. (2011) (these authors also model the operating decision of a commodity processor; see Boyabatli et al., 2011 and Boyabatli, 2011 for related work). Different from these systems, our model has a random inflow, the wind energy.

Related settings with random inflow include hydropower generation (Nasakkala and Keppo 2008) and liquified natural gas (LNG) regasification (Lai et al. 2010b). These systems differ from ours in that they store their input (water or LNG), while our system stores the output (electricity). Thus, the operating policies in these papers feature only one state-dependent sell-down threshold, while our triple-threshold policy has additional state-dependent store-up-to thresholds.

3. Model

We consider the problem of operating a WST system: a remote wind farm is co-located with a storage facility, both of which are connected to a wholesale market via a transmission line (Figure 1). The operator of this system can buy and sell electricity in this market (the transmission line is bi-directional). We assume that the WST system is small compared to the market, so the operator's decisions do not affect market prices. The operator makes trading decisions periodically over a finite horizon, specifically in time periods $t \in \mathcal{T} := \{0, 1, \dots, T\}$. In terminal period T , any electricity left in the storage facility is worthless.

Parameters. We assume the storage facility is finite in *energy* capacity (space) and *power* capacity; without loss of generality, we normalize the energy capacity to one (energy unit). If we think of the storage facility as a warehouse for electricity, the *energy* capacity is analogous to the space of the warehouse; the *power* capacity is analogous to the maximum rate of adjusting the warehouse inventory. For the rest of this paper, any capacity should be interpreted as *power* capacity, unless specified otherwise. We use the following parameters:

- K_1, K_2 : charging and discharging capacity (in energy units/period), respectively; $K_1 < 0 < K_2$.
- G, C : generation and transmission capacity (in energy units/period), respectively. We assume that $G + K_2 \geq C$. This is realistic: it is reasonable to size the transmission capacity to be less than the sum of the wind farm's generation capacity and the storage discharging capacity (if $G + K_2 < C$, the transmission capacity is never constraining).

- α, β, η : efficiency of the storage facility in charging, discharging, and storing over one period, respectively; all three parameters are in $(0, 1]$. The round-trip efficiency is defined as $\alpha \cdot \beta \cdot \eta$.

- τ : transmission efficiency, the ratio of the quantity of electricity flowing out of the transmission line to that flowing into this line, i.e., $1 - \tau$ is the ratio of the loss in the line. We apply τ at the end of the transmission in either direction.

- δ : one-period risk-free discount rate (we use risk-neutral valuation; Seppi, 2002); $0 < \delta < 1$.

State variables. A state variable with subscript t is known at time t , but unknown at earlier times, that is $0, 1, \dots, t - 1$ (Powell 2007, §5.2). At time t , the state S_t includes:

- x_t : inventory of electricity (in energy units) in the storage facility at time t . The domain of this variable is $\mathcal{X} := [0, 1 \cdot \eta]$. The maximum inventory level is η because we assume that the storage loss happens at the end of each period.

- w_t : the maximum amount of available energy (in energy units) that the wind farm can produce in period t ; $w_t = \min\{G, \text{available wind energy in period } t\}$, where G is implicitly multiplied by one period.

- \vec{p}_t : price vector of electricity in period t (\$/energy unit); each component of \vec{p}_t is a factor of the price model in §6. We denote the sum of all the components in \vec{p}_t by p_t , a scalar.

We thus define S_t as (x_t, w_t, \vec{p}_t) . In particular, S_0 is the *given* initial state (x_0, w_0, \vec{p}_0) , where w_0 and \vec{p}_0 are the given wind and price in the initial period, respectively.

Random variables. At time t , the random variables are

- w_{t+1} : the electricity that the wind farm can produce in period $t + 1$; $w_{t+1} \in \mathcal{W} \subseteq \mathbb{R}^+$. The evolution of w_{t+1} is modeled as a Markov process.

- \vec{p}_{t+1} : the price vector of electricity in period $t + 1$; $\vec{p}_{t+1} \in \mathcal{P} \subseteq \mathbb{R}^n$ (note \mathbb{R}^n instead of $(\mathbb{R}^+)^n$), where n is the number of components in \vec{p}_{t+1} . The evolution of each component (factor) in \vec{p}_{t+1} is modeled as a Markov process.

These random variables are exogenous and may be correlated.

Decision variables.

- g_t : the quantity of electricity to generate in period t ; $g_t \in \mathbb{R}^+$.
- a_t : the inventory change in period t ; $a_t \in \mathbb{R}$. If $a_t < 0$, a_t is the decrease in inventory due to selling; if $a_t \geq 0$, a_t is the increase in inventory due to generation and/or buying. It is easily shown that the option of simultaneously increasing and decreasing inventory is suboptimal.

Sequence of events. For each period $t \in \mathcal{T}$, we assume the following sequence of events:

- 1) The operator observes w_t and \vec{p}_t , and then determines the joint inventory and generation action pair (a_t, g_t) .
- 2) Electricity flows through the system, incurring the loss in discharging (selling) or charging (generation and/or buying), as well as the loss in transmission.
- 3) The associated financial settlement is completed. At the end of period t , a fraction $1 - \eta$ of any resulting inventory is lost.

Transition functions. The inventory level in period $t \leq T - 1$ is $x_{t+1} = \eta(x_t + a_t)$. The quantities w_t and \vec{p}_t evolve to w_{t+1} and \vec{p}_{t+1} according to two exogenous stochastic processes (we specify them in §6 and §7). As mentioned earlier, we assume that the operator's decisions do not affect market price.

Immediate payoff function and constraints. Let $R(a_t, g_t, p_t)$ denote the immediate payoff function (in period t) of the triple (a_t, g_t, p_t) .

- If $a_t < 0$, then the inventory decreases, which means that the total quantity sold is $(g_t - a_t\beta) \cdot \tau$. This quantity, before the transmission loss, cannot exceed the transmission capacity, so we have constraint C1: $g_t - a_t\beta \leq C$, where C is implicitly multiplied by one period.
- If $0 \leq a_t/\alpha \leq g_t$, the wind farm generates more electricity than it stores in the facility, which means that the operator sells the excess electricity to the market: $(g_t - a_t/\alpha) \cdot \tau$, which before the transmission loss cannot surpass the transmission capacity. This yields constraint C2: $g_t - a_t/\alpha \leq C$.
- If $a_t/\alpha > g_t$, more electricity is stored than is generated, which means that the operator buys the quantity $(a_t/\alpha - g_t)/\tau$ from the market. This quantity likewise cannot exceed the transmission capacity, yielding constraint C3: $(a_t/\alpha - g_t)/\tau \leq C$.

Thus, $R(a_t, g_t, p_t)$ is defined as follows:

$$R(a_t, g_t, p_t) := \begin{cases} p_t \cdot (g_t - a_t\beta) \cdot \tau, & \text{if } a_t < 0, \\ p_t \cdot (g_t - a_t/\alpha) \cdot \tau, & \text{if } 0 \leq a_t \leq g_t \cdot \alpha, \\ -p_t \cdot (a_t/\alpha - g_t)/\tau, & \text{if } a_t > g_t \cdot \alpha, \end{cases} \quad (1)$$

where the first two cases represent the selling revenue and the third case the purchasing cost. The feasible set of action pairs (a_t, g_t) , denoted by $\Psi(x_t, w_t)$, is defined by the following constraints:

$$\begin{aligned} \text{C1: } & g_t - a_t\beta \leq C, \text{ if } a_t < 0, \\ \text{C2: } & g_t - a_t/\alpha \leq C, \text{ if } 0 \leq a_t \leq g_t \cdot \alpha, \\ \text{C3: } & (a_t/\alpha - g_t)/\tau \leq C, \text{ if } a_t > g_t \cdot \alpha, \\ \text{C4: } & g_t \leq w_t, \\ \text{C5: } & -x_t \leq a_t \leq 1 - x_t, \\ \text{C6: } & K_1 \leq a_t \leq K_2; \end{aligned}$$

C4 constrains the generation by the wind energy availability (this constraint implies that $g_t \leq G$, as $w_t \leq G$); C5 limits the inventory change by the energy available in the storage facility (left) and remaining storage *energy* capacity (right); and C6 incorporates the charging and discharging capacities (K_1 and K_2 are implicitly multiplied by one period).

Objective function. We formulate our problem as an MDP. Each stage of our MDP corresponds to one time period. A policy π is a mapping from any state S_t in any stage t to a feasible action pair (a_t, g_t) . Let $A_t^\pi(S_t)$ denote the decision rule of policy π in period t , and let Π denote the set of all feasible policies. Our objective is to maximize the total discounted expected cash flows over all feasible policies:

$$\max_{\pi \in \Pi} \sum_{t=0}^T \delta^t \mathbb{E} [R(A_t^\pi(S_t), p_t) | S_0], \quad (2)$$

where the expectation \mathbb{E} is taken with respect to risk-adjusted distributions of wind and price, as well as inventory under the policy π ; recall that we use risk-neutral valuation (Seppi 2002).

For each period $t \in \mathcal{T}$, we define $V_t^*(S_t)$ as the optimal value function from period t onward. In addition, we set $V_T^*(S_T) := 0, \forall S_T$. For every period $t < T$ and state S_t , $V_t^*(S_t)$ satisfies

$$V_t^*(S_t) = \max_{(a_t, g_t) \in \Psi(x_t, w_t)} R(a_t, g_t, p_t) + \mathbb{E} [\delta V_{t+1}^*(S_{t+1}) | S_t]. \quad (3)$$

4. Analysis

In this section we build the analytical platform to investigate the difficulty of solving the MDP (3), the relevance of considering negative prices and the buying option, and the value of different types of information (future information and uncertainty). We numerically examine these issues in §8.2.

4.1. Value of negative prices

To study the importance of considering negative prices, we modify the MDP (3) to ignore them; the immediate payoff function in the MDP (3) replaces p_t with $p_t^+ := \max\{p_t, 0\}$ as follows:

$$V_t^{H1}(S_t) = \max_{(a_t, g_t) \in \Psi(x_t, w_t)} R(a_t, g_t, p_t^+) + \delta \mathbb{E} [V_{t+1}^{H1}(S_{t+1}) | S_t], \quad (4)$$

where the superscript $H1$ abbreviates ‘‘Heuristic 1’’. Note that prices in (4) can be negative, but any such negative price is taken as zero in the immediate payoff function; thus, (4) includes the special case when prices are always nonnegative.

To avoid trivial cases, we make the following benign assumption:

ASSUMPTION 1. For every period $t \in \mathcal{T}$, $\mathbb{E}[|p_k| | w_t, \vec{p}_t] < \infty$ for all $k \geq t$, $k \in \mathcal{T}$.

The absence of negative prices in the immediate reward function, combined with the lack of inventory holding cost and salvage penalty in model (4), imply that the operator is always better off having more inventory. This gives rise to Lemma 1 (see its proof in Online Appendix A):

LEMMA 1. For every period $t \in \mathcal{T}$, $V_t^{H1}(x_t, w_t, \vec{p}_t)$ is non-decreasing in x_t given any $w_t \in \mathcal{W}$ and $\vec{p}_t \in \mathcal{P}$.

Lemma 1, together with how the immediate payoff function varies with the generation decision g_t , implies that in model (4) it is optimal to generate as much energy as possible. This is Lemma 2, which is formally proved in Online Appendix B.

LEMMA 2. For each period $t \in \mathcal{T}$, the optimal generation action $g_t^{H1}(S_t)$ for model (4) is $\min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\}$.

Before we prove the optimal inventory policy structure for (4) in Proposition 2, we prove the concavity of $V_t^{H1}(x_t, w_t, \vec{p}_t)$ in Proposition 1 (see Online Appendix C for its proof).

PROPOSITION 1. For every period $t \in \mathcal{T}$, it holds that $|V_t^{H1}(x_t, w_t, \vec{p}_t)| < \infty$ and $V_t^{H1}(x_t, w_t, \vec{p}_t)$ is concave in $x_t \in \mathcal{X}$ given any $w_t \in \mathcal{W}$ and $\vec{p}_t \in \mathcal{P}$.

We now can establish the optimal inventory policy for (4), in Proposition 2 (the proof is in Online Appendix D). We define $x^- := \min\{x, 0\}$.

PROPOSITION 2. For every period $t \in \mathcal{T}$, given any $w_t \in \mathcal{W}$ and $\vec{p}_t \in \mathcal{P}$, there exist three inventory functions $\underline{X}_t^1(w_t, \vec{p}_t) \leq \underline{X}_t^2(w_t, \vec{p}_t) \leq \overline{X}_t(w_t, \vec{p}_t)$ (simplified to \underline{X}_t^1 , \underline{X}_t^2 , and \overline{X}_t , respectively) such that $a_t^{H1}(S_t)$ is computed as follows (recall that $g_t^{H1}(S_t) = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\}$):

A) If either (1) $w_t \geq C + \min\{1 - x_t, K_2\}/\alpha$ or (2) $C \leq w_t < C + \min\{1 - x_t, K_2\}/\alpha$ and $x_t > \underline{X}_t^2 - \alpha(w_t - C)^+$: $a_t^{H1}(S_t) = \alpha(g_t^{H1}(S_t) - C)^+$ (fill up);

B) If (1) $x_t + \alpha(w_t - C)^+ \leq \underline{X}_t^2$ and $w_t < C + \min\{1 - x_t, K_2\}/\alpha$, but not (2) $w_t \leq \min\{1 - x_t, K_2\}/\alpha$ and $x_t + w_t \cdot \alpha \leq \underline{X}_t^1$: $a_t^{H1}(S_t) = \min\{\underline{X}_t^2 - x_t, \alpha \cdot w_t, K_2\}$ (store up to \underline{X}_t^2);

C) If $w_t \leq \min\{1 - x_t, K_2\}/\alpha$ and $x_t + w_t \cdot \alpha \leq \underline{X}_t^1$: $a_t^{H1}(S_t) = \min\{\underline{X}_t^1 - x_t, \alpha(C\tau + w_t), K_2\}$ (store up to \underline{X}_t^1);

D) If $w_t < C$ and $\underline{X}_t^2 \leq x_t \leq \bar{X}_t$: $a_t^{H1}(S_t) = 0$ (keep inventory unchanged);

E) If $w_t < C$ and $x_t > \bar{X}_t$: $a_t^{H1}(S_t) = \max\{\bar{X}_t - x_t, (w_t - C)^-/\beta, K_1\}$ (sell down to \bar{X}_t).

Throughout this paper, \underline{X}_t^1 , \underline{X}_t^2 , and \bar{X}_t are time-dependent and shorthand notations for the functions $\underline{X}_t^1(w_t, \vec{p}_t)$, $\underline{X}_t^2(w_t, \vec{p}_t)$, and $\bar{X}_t(w_t, \vec{p}_t)$, respectively.

Proposition 2 can be interpreted as follows. Given any \vec{p}_t , and any point on the (x_t, w_t) plane, the optimal inventory action of policy *H1* depends on the region into which this point falls: in region A) it is optimal to generate, sell, and store as much as possible; in region B) it is optimal to generate and store as much as possible to reach \underline{X}_t^2 , and sell the rest; in region C) it is optimal to first generate and store, and then to buy as much as possible to reach \underline{X}_t^1 ; in region D) it is optimal to keep the inventory level unchanged; i.e., to sell whatever is generated; and in region E) it is optimal to sell down as much as possible to reach inventory level \bar{X}_t .


For *H1*, given any \vec{p}_t , the most general structure of its optimal inventory action occurs when $w_t < \min\{C, K_2/\alpha\}$. For such a case, this action is as follows (see Figure 2):

$$a_t^{H1}(S_t) = \begin{cases} \min\{\underline{X}_t^1 - x_t, \alpha(C \cdot \tau + w_t), K_2\}, & \text{if } x_t \in [0, (\underline{X}_t^1 - \alpha w_t)^+), \\ \min\{\underline{X}_t^2 - x_t, w_t \cdot \alpha, K_2\}, & \text{if } x_t \in [(\underline{X}_t^1 - \alpha w_t)^+, \underline{X}_t^2), \\ 0, & \text{if } x_t \in [\underline{X}_t^2, \bar{X}_t], \\ \max\{\bar{X}_t - x_t, (w_t - C)^-/\beta, K_1\}, & \text{if } x_t \in (\bar{X}_t, \eta]. \end{cases}$$

This means that for this case the optimal action of *H1* can be of four distinctive types: if $x_t < (\underline{X}_t^1 - \alpha w_t)^+$, generate and buy as much as possible to reach \underline{X}_t^1 ; if $(\underline{X}_t^1 - \alpha w_t)^+ \leq x_t \leq \underline{X}_t^2$, generate and store as much as possible to reach \underline{X}_t^2 ; if $\underline{X}_t^2 \leq x_t \leq \bar{X}_t$, maintain the same inventory level; if $x_t > \bar{X}_t$, sell to bring the inventory as close to \bar{X}_t as possible. Because of these three inventory threshold functions, we call the optimal policy the *triple-threshold* policy.

The existence of these three thresholds is due to the concavity of $V_t^{H1}(x_t, w_t, \vec{p}_t)$ in x_t , in the presence of two efficiency losses: transmission loss ($\tau < 1$), and the loss in conversion ($\alpha \cdot \beta < 1$). As a result, the marginal values of the following types of action differ: storing one unit bought from the market; storing one unit from generation; and selling one unit from inventory. If $\tau = 1$, then $\underline{X}_t^1 = \underline{X}_t^2$, because the first two types of action have the same marginal value; if $\alpha \cdot \beta = 1$, then $\underline{X}_t^2 = \bar{X}_t$, because the last two types of action have the same marginal value. In the case $\tau = 1$, our optimal policy reduces to that in Secomandi (2010b).

Figure 2 The Optimal Policy Structure of the Triple-threshold Policy for the Case $w_t < \min\{C, K_2/\alpha\}$ given \vec{p}_t



This triple-threshold policy deviates from the optimal policy to MDP (3) when the optimal immediate payoff function in MDP (3) is non-concave. A necessary but not sufficient condition of such deviation is the presence of efficiency loss *and* negative prices. The non-concavity of the optimal value functions in the MDP (3) originates from the *convexity* of the immediate payoff function in the decision variable a_t when prices are negative: in (1), for any $p_t < 0$, we have $-p_t/(\alpha \cdot \tau) \geq -p_t \cdot \tau/\alpha \geq -p_t \cdot \tau \cdot \beta$. Intuitively, when prices are negative, at any low inventory level, *if* the optimal action is to sell (which may happen if the expected price for the next period is more negative), then the marginal value of inventory is the (negative) price *divided* by the discharging efficiency. At any high inventory level, if the optimal action is to store, then the marginal value of inventory is the (negative) price *times* the charging efficiency. Thus, given such optimal actions, the marginal value of inventory at a low inventory level is lower than that at a high inventory level, resulting in the non-concavity of the optimal value function in inventory.

As this triple-threshold policy ignores negative prices, the difference between the value of the optimal policy to (3) and that of this policy gives the value of considering negative prices.

4.2. Value of the buying option

We now examine the value of buying by considering a variant of MDP (3) which does not allow buying, even though prices can be negative. In this model for each period $t \in \mathcal{T}$, the only difference from the original MDP (3) is in the feasible set: the new feasible set imposes the condition $a_t \leq g_t \alpha$, which is equivalent to removing constraint C3 from $\Psi(x_t, w_t)$. For convenience, we use superscript *H2* for the optimal actions and the optimal value functions of the resulting model.

Under this modification, the monotonicity of the value functions in inventory still holds, from Lemma 1. We then have:

LEMMA 3. For every period $t \in \mathcal{T}$, the generation action $g_t^{H2}(S_t)$ is determined as follows: If $p_t < 0$, $g_t^{H2}(S_t) = \min\{w_t, (1 - x_t)/\alpha, K_2/\alpha\}$; if $p_t \geq 0$, $g_t^{H2}(S_t) = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\}$.

Lemma 3 agrees with Lemma 2 when the price is nonnegative: generate as much as possible. But when the price is negative, the optimal action of $H2$ is to sell nothing and generate as much as possible. Next we show that despite the possibility of negative prices, the concavity of the value functions in inventory continues to hold for the policy $H2$, due to the absence of the buying option, as stated in Proposition 3 (see Online Appendix F for its proof):

PROPOSITION 3. For every period $t \in \mathcal{T}$, it holds that $|V_t^{H2}(S_t)| < \infty$ and $V_t^{H2}(x_t, w_t, \vec{p}_t)$ is concave in $x_t \in \mathcal{X}$ given any $w_t \in \mathcal{W}$ and $\vec{p}_t \in \mathcal{P}$.

With Proposition 3, we can prove the optimal inventory action of $H2$ in Proposition 4.

PROPOSITION 4. For every period $t \in \mathcal{T}$, the optimal inventory action $a_t^{H2}(S_t)$ is as follows: 1) if $p_t < 0$, $a_t^{H2}(S_t) = \alpha \min\{w_t, (1 - x_t)/\alpha, K_2/\alpha\}$; 2) if $p_t \geq 0$, the optimal inventory action is the same as that in Proposition 2 with the exclusion of region C .

Proposition 4 can be interpreted as follows: 1) if $p_t < 0$, sell nothing and store as much as possible from generation; 2) if $p_t \geq 0$, generate as much as possible, and the inventory policy structure is a special case of the triple-threshold policy, a *dual-threshold* policy with two threshold functions: generate-and-store-up-to, and sell-down-to. Similar dual-threshold policies have also been found in the extant literature (Secomandi 2010b, Nascimento and Powell 2010).

As this dual-threshold policy does not buy, the difference between the value of the optimal policy to (3) and that of this policy gives the value of the buying option.

Comparing the triple-threshold policy and the dual-threshold policy, the former dominates the latter when prices are non-negative. When prices are negative (in evaluating the triple-threshold policy in MDP (3)), the triple-threshold policy regards them as zero when evaluating the current cash flow, and thus would likely buy should generation falls short of filling up the storage facility and take advantage of this negative price, while the dual-threshold policy cannot. Therefore, the triple-threshold policy is likely to outperform the dual-threshold policy, which suggests that the value of negative prices is no greater than the value of the buying option. Much of this is due to the fact that without the buying option, the value of modeling negative prices is quite limited.

4.3. Value of information and optimization

In this subsection, we examine the value of including certain information or the value of optimization by studying three heuristics.

Two-period policy. This heuristic makes decisions as if the problem were a two-period problem: for each period, the value of any inventory at the end of the *second* period is zero. Specifically, for every stage t and every state S_t , we find the action of this policy by solving a modified version of the MDP (3): we set $V_k(S_k) = 0$ for all $k \in \mathcal{T}$, $k \geq t + 2$ in (3). The resulting policy ignores future information beyond two periods. Thus, the difference between the value of the optimal policy and that of this policy is the value of future information.

Rolling horizon policy. This policy makes decisions by ignoring the uncertainty in the price and wind energy evolution: for every stage and every state, we first construct a forward curve of the price and of the wind energy through all the remaining periods (each forward curve consists of expected values conditional on the current state). We then find an optimal action for every stage and state by solving a linear program, which is the deterministic version of the MDP (3).

This rolling horizon policy ignores the uncertainty in the optimization, which is different from assuming price and wind energy are deterministic: we still model the uncertainty in the evolution of these variables, but we make decisions in each stage and state as if the future were certain. Thus, the difference between the value of the optimal policy and that of this policy gives the value of *including uncertainty when making decisions*. The uncertainty is often ignored in the engineering literature, for instance, by Denholm and Sioshansi (2009) and Fertig and Apt (2011). This rolling horizon policy allows us to quantify the performance of policies that ignore uncertainty in decision-making.

Naïve policy. This is one of the simplest yet sensible policies. If the price for the current stage and state is positive, generate and sell as much as possible and store as much leftover electricity as possible; if this price is negative, generate and store as much as possible. We take the difference between the value of the optimal policy and that of this policy as *the value of optimization*.

5. Storage valuation

The value that storage provides for a WST system can be three-fold:

- Storage can reduce wind energy curtailed due to the transmission capacity constraint;
- Storage enables the wind farm to time-shift generation to periods of more favorable prices;
- Storage enables the wind farm to buy electricity from the market for future resale.

For these reasons, the value of storage can be measured both financially, as the increase in the total value of the system; or with respect to energy, as the decrease in the wind energy curtailed, or the increase in energy or wind energy sold. We compute these values in §8.3.

5.1. Monetary value of storage

We quantify the monetary value of storage as the percentage increase in the value of the system with storage under the triple-threshold policy compared to that of the system without storage. (We use the triple-threshold policy as it is more computationally tractable than the optimal policy and is yet near-optimal, as shown in §8.2.) Specifically, the value of storage is

$$\frac{V_0^{H1} - V_0^{NS}}{V_0^{NS}} \times 100 = \left(\frac{V_0^{H1} - V_0^{H2}}{V_0^{NS}} + \frac{V_0^{H2} - V_0^{H3}}{V_0^{NS}} + \frac{V_0^{H3} - V_0^{NS}}{V_0^{NS}} \right) \times 100, \quad (5)$$

where V_0^{H1} , V_0^{H2} , and V_0^{H3} are the values of the system with storage in the initial stage and state using the triple-threshold policy, the dual-threshold policy, and the naïve policy (denote by $H3$), respectively; and V_0^{NS} is the optimal value of the system without storage (the arguments of all the V_0 's are removed for simplicity). Using (5), we can interpret the value of storage as the sum of the following three components:

- $(V_0^{H1} - V_0^{H2})/V_0^{NS}$: the value of storage due to arbitrage, as $V_0^{H1} - V_0^{H2}$ is the difference between the value of the system with storage under the triple-threshold policy and the value of the same system under the dual-threshold policy (which does not buy).
- $(V_0^{H2} - V_0^{H3})/V_0^{NS}$: the value of storage due to time-shifting generation, as $V_0^{H2} - V_0^{H3}$ is the difference between the values of the system with storage under the dual-threshold policy and the naïve policy (which sells as much as possible when prices are positive).
- $(V_0^{H3} - V_0^{NS})/V_0^{NS}$: the value of storage due to reduced curtailment, as $V_0^{H3} - V_0^{NS}$ is the difference between the value of the system with storage under the naïve policy and that of the system without storage.

5.2. Energy value of storage

To measure the energy value of storage we again use the triple-threshold policy. For the cases both with and without storage, we compare (i) the average wind energy curtailed: the average wind energy available minus the average wind energy generated; (ii) the average energy sold: the average energy sold to the market, from generation and inventory; and (iii) the average wind energy sold: the difference between the average energy sold and the average energy bought from the market. These quantities are the sums of the corresponding expected quantities for all stages, starting in the initial stage and state at inventory level zero, divided by the number of periods (T).

To provide a basis for evaluating the performance of different policies and computing the values of storage, we need a price and a wind energy model.

6. Price model

In this section we specify the price model that is used in our numerical study in §8. We describe the price data that we use in §6.1, review the literature on commodity price models in §6.2, and discuss the price model that we use in §6.3. We discuss how we calibrate this model to this data in §6.4, and how we discretize this model in §6.5.

6.1. Price data

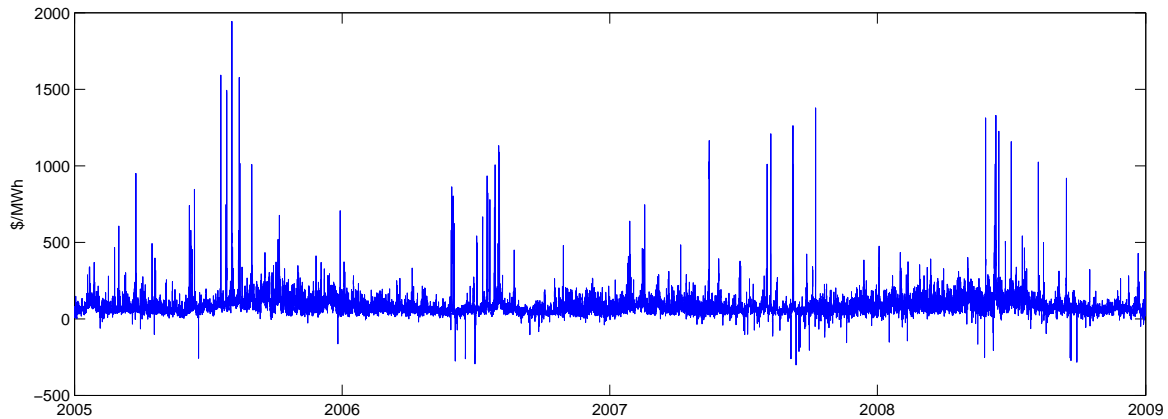
We use price data from NYISO; NYISO is among the largest and most liquid electricity markets (NYISO 2011). For each of its fifteen zones, NYISO reports the real-time price for every five minutes from 1999 until the present¹. We choose a single zone (New York City), and focus on data from 2005-2008, as this time span is recent and long enough for calibration. We take the real-time price of every five minutes in New York City and average it over every hour, and use this as the original price series for calibrating our model. In §8 we use one hour as the stage length of our MDPs; most real-time markets, including NYISO, run on a hourly interval: though prices are given every five minutes, they are cleared simultaneously once every hour.

We plot this hourly price series in Figure 3. This figure shows that these prices exhibit mean reversion, a tendency to revert back to the mean price level (approximately \$85/MWh, where MWh represents megawatt hour); and seasonality, a term we use to describe any repeated pattern on any time scale, such as hourly, daily, weekly, monthly, and yearly. The most prominent feature, however, is the variability, with a substantial number of positive and negative price jumps. The maximum price is \$1,943.5/MWh; the minimum is \$-300.3/MWh. Another important feature is the existence of negative prices: 166 negative hourly prices over four years ($166/365/24/4 \times 100\% = 0.47\%$), most of which result from negative jumps. Thus, a reasonable price model for this data should capture mean reversion, seasonality, negative prices, and both positive and negative jumps.

6.2. Electricity price models

Since electricity is a commodity, its price models are mostly based on those of general commodities from the field of financial engineering (such as Schwartz and Smith 2000, and references therein). Most of these commodity price models (mainly reduced-form, see Seppi 2002) capture mean reversion and a long-term trend, and can be tailored to describe special features of electricity prices, such as seasonality (Lucia and Schwartz 2002). Apart from reduced-form price models, there are also equilibrium models to describe the evolution of electricity prices, such as those of Barlow (2002) and Benth and Koekebakker (2008). Other models, commonly used in the electrical engineering

¹ www.nyiso.com/public/markets_operations/market_data/pricing_data/index.jsp; last retrieved, 12/2011.

Figure 3 The Average Hourly Real-time Price in New York City from 2005-2008 in NYISO

literature, include input/output hidden Markov models (Gonzalez et al. 2005) and artificial neural networks (Szkuta et al. 1999). A comprehensive review of electricity price models in the engineering field can be found in Bunn (2004).

To fully describe the features of electricity prices one must capture extreme variability. For example, on June 25, 1998, the electricity price in the Midwest region jumped to \$7,500/MWh (FERC 1998), compared to around \$40/MWh on average in that year. This type of extreme variability cannot be modeled by pure Gaussian price models. To capture it, alternatives include:

- Adding a jump process where the jump size may be modeled as a normal distribution (Cartea and Figueroa 2005), or an affine jump diffusion (Villaplana 2005) in which the jump size and intensity are affine functions of the state variables.
- Modeling volatility as a stochastic process, for instance, a GARCH model (Escribano et al. 2002, Garcia et al. 2005), or a Lévy process (Benth et al. 2007).
- Specifying—on top of “normal” regimes—regimes of “abnormal” states, a method often called regime-switching, such as in Deng (2000), Rambharat et al. (2005), Huisman and Mahieu (2003), De Jong (2006), Geman and Roncoroni (2006), and Seifert and Uhrig-Homburg (2007). Variants of this approach differ in modeling the transitions between regimes.

6.3. Our price model

To obtain a price model that describes all the features of electricity prices discussed in §6.1, we modify one of the models in Lucia and Schwartz (2002) by generalizing its seasonality component and adding a jump process (the first alternative of modeling price volatility). Specifically, in our reduced-form model, the electricity price p_t for period t is

$$p_t = \xi_t + f(t) + J_t,$$

where ξ_t is a mean reverting process; $f(t)$ a deterministic seasonality function; and J_t a jump process. A mean-reverting process is commonly used in energy price models, such as in Smith and McCardle (1999), Seppi (2002), Jaillet et al. (2004), Secomandi (2010a,b), and Devalkar et al. (2011). As in one of the models in Lucia and Schwartz (2002), we use the following discrete-time mean reverting model with zero mean-reversion level:

$$\xi_{t+1} - \xi_t = -\kappa\xi_t dt + \sigma\epsilon_t,$$

where κ is the mean-reverting rate; σ the constant volatility; and each $\epsilon_t \sim N(0, 1)$ is an i.i.d. error term. We use this spot price model rather than a log-spot price model because of the existence of negative electricity prices. The mean price level in the data is captured by the seasonality function of our model.

To obtain the seasonality function $f(t)$, we generalize one from Lucia and Schwartz (2002), which consists of three terms: a constant level; a term to describe the effect of a non-business day; and a term to capture monthly seasonality. To capture the non-business day effect, they use the term $\gamma^1 \cdot D_t^1$, where γ^1 is the magnitude of this effect, and D_t^1 is a dummy variable which equals 1 if period t belongs to a holiday or a weekend, and 0 otherwise. We add another term to this seasonality to capture the hour-of-day effect. To summarize:

$$f(t) = A + \gamma^1 \cdot D_t^1 + \sum_{i=1}^{11} \gamma^{2i} \cdot D_t^{2i} + \sum_{j=1}^{23} \gamma^{3j} \cdot D_t^{3j},$$

where A is the constant level; γ^{2i} ($i = 1, \dots, 11$) and γ^{3j} ($j = 1, \dots, 23$) represent the seasonality of month i and hour j , respectively; and D_t^{2i} ($i = 1, \dots, 11$) and D_t^{3j} ($j = 1, \dots, 23$) are dummy variables that equal one if period t is in month i and hour j , respectively.

We model the jump component J_t as a compound Bernoulli process where each jump occurs in each period with probability λ as follows:

$$J_t = \text{jump size} \cdot \{1, \text{ if there is a jump}; 0, \text{ if there is no jump}\},$$

where the jump size follows an empirical distribution, derived from our data. This jump arrival process is similar to Cartea and Figueroa (2005), but we use an empirical distribution, rather than a normal distribution because the latter does not fit the jump size extracted using the algorithm below, and also because the empirical distribution is sufficient to solve our MDPs *numerically*. Also unlike in Cartea and Figueroa (2005) where multiple jumps can occur in one period (one day in their paper), we assume that only one jump can occur per period, as our period length is one hour. Finally, unlike in Cartea and Figueroa (2005) where jumps revert back through mean reversion, for tractability jumps in our paper are spikes (Deng, 2000) that last only one period. This fits our data well; the average length of a jump in our data is 1.25 periods.

Table 1 Estimated Parameters for Each Component of the Price Model (MAE = \$10.7955/MWh)

Mean Reversion		Jump	Seasonality	
κ	σ	λ	A	γ^1
0.1679	14.1098	0.1253	65.5475	-6.0324

6.4. Price model calibration

We calibrate our price model to the hourly price series described in §6.1. To estimate the model parameters, we use the following iterative approach.

Step 0. Initialize ν_J , the set of periods where jumps are found; set it to be empty.

Step 1. Calibrate seasonality: Remove from the original hourly price $\{p_t : 1 \leq t \leq \bar{t}, \text{ where } \bar{t} \text{ is the maximum number of periods}\}$ those periods in set ν_J . Use a linear regression to estimate parameters for seasonality, and remove seasonality to obtain a new series of prices pertaining only to the mean reversion plus possible jumps, denoted by $\{p'_t : 1 \leq t \leq \bar{t}\}$.

Step 2. Calibrate mean reversion: Based on $\{p'_t : 1 \leq t \leq \bar{t}\}$, estimate mean reversion rate $\hat{\kappa}$ and standard deviation $\hat{\sigma}$. Set the mean reversion process $\{\hat{\xi}_t : 1 \leq t \leq \bar{t}\} = \{p'_t : 1 \leq t \leq \bar{t}\}$.

Step 3. Identify jumps and recompute mean reversion: For any $t < \bar{t} - 1$, if $\hat{\xi}_{t+1}$ is outside of $n_\sigma \cdot \hat{\sigma}$ around its mean (n_σ is used to control the number of jumps extracted; we experimented with it and set it as 2.7), that is if $\hat{\xi}_{t+1}$ is outside of $[(1 - \hat{\kappa}) \cdot \hat{\xi}_t \pm n_\sigma \hat{\sigma}]$, then $\hat{\xi}_{t+1}$ contains a jump. Define the jump size as $\hat{\xi}_{t+1} - (1 - \hat{\kappa})\hat{\xi}_t$, add this period t to ν_J , and update $\hat{\xi}_{t+1}$ to be $(1 - \hat{\kappa})\hat{\xi}_t$.

Step 4. Repeat Step 1 through Step 3 until the estimated parameters for both the seasonality and mean reversion have converged: the difference between one iteration and the next falls within a prespecified tolerance. (The actual estimation stopped with no estimated parameter change.)

This estimation process yields the estimated parameters for mean reversion and seasonality displayed in Table 1 and 2. The mean absolute error (MAE) is 10.7955 (\$/MWh). We assume the market price of risk (Duffie 1992) is zero, because we do not have futures prices to calibrate it. If one had such data, one could estimate it accordingly (Lucia and Schwartz 2002).

We next bin the extracted jumps to construct an empirical distribution (Table 3), and estimate the jump arrival rate λ : jump occurrences divided by periods (Table 1). This method, though simple, extracts almost all jumps that are visually apparent. More sophisticated methods could also be applied, see for example Klüppelberg et al. (2010) and Fanone and Prokopczuk (2010).

6.5. Discretization

To numerically solve our MDPs we discretize the mean reverting process to a lattice: A tree with discrete time steps that specifies attainable price levels and their probabilities for each period. Using the estimated parameters for the mean reversion model, we follow the method in Hull and

Table 2 γ^{2i} and γ^{3i} for Seasonality Component

i	1	2	3	4	5	6	7	8	9	10	11	
γ^{2i}	-0.652	-0.978	-0.421	1.607	-3.853	3.592	12.621	9.738	-3.437	-4.167	-8.456	
γ^{3i}	-1.61	-7.243	-12.669	-15.392	-12.546	-8.226	-1.898	1.712	8.899	14.514	17.293	
i	12	13	14	15	16	17	18	19	20	21	22	23
γ^{3i}	19.712	19.780	18.498	17.304	17.001	18.183	22.226	21.794	21.661	21.361	17.783	10.678

Table 3 Empirical Jump Size Distribution

Size	-300	-250	-200	-150	-100	-50	50	100	150	200	250
Prob.	0.0027	0.0027	0.0007	0.0041	0.0191	0.1802	0.4809	0.1903	0.0537	0.0196	0.0105
size	300	350	400	450	500	550	600	650	700	750	800
Prob.	0.0082	0.0048	0.0034	0.0023	0.0011	0.0018	0.0027	0.0016	0.0011	0.0009	0.0011
Size	850	900	950	1000	1050	1150	1200	1350	1450	1800	
Prob.	0.0016	0.0009	0.0005	0.0009	0.0007	0.0005	0.0002	0.0005	0.0005	0.0002	

Note. Any size that does not appear in this table has probability zero.

White (1993) to construct a trinomial price lattice: each time step Δt equals an hour; prices in all periods can be only multiples of $\sqrt{3}\sigma\Delta t$, and each price level in each period transits to three possible price levels in the next period with probabilities that are computed to match the first two moments of the original continuous time distribution. For more details, please refer to Hull and White (1993). This discretized mean reverting process and the discrete jump process comprise two dimensions in the state of our MDPs.

7. Wind energy model

In this section we specify the wind energy model that is used in our numerical study in §8. We describe the data and our model in §7.1, and the calibration and discretization of our model in §7.2 and §7.3, respectively.

7.1. Data and model

We use hourly wind speed data from Buffalo, home to one of the largest wind farms in New York state². We obtain these data from NOAA (2010), and focus on 2005-2008, the same time span as the price data. The wind speed obtained was recorded at 10 meters above ground, thus we need to convert it to wind speed at the hub height of a wind turbine. We accomplish this by using the model from Heier (2006):

$$\text{Wind speed at height } h = \text{Wind speed at 10 meters} \cdot (h/10)^{\text{constant}},$$

where this constant depends on the geographic location, the ground terrain, and the air stability. We set $h = 80$ meters, the height for the General Electric (GE) turbine model 1.5-77, because this

² www.dec.ny.gov/docs/permits_ej_operations_pdf/windpwrnys.pdf

model is among the best selling turbines in the U.S. (Wiser and Bolinger 2010). We choose the constant to be 0.2, so that our wind data’s capacity factor (the ratio of actual energy output over a long period versus nameplate capacity) is 33%, in the range of capacity factors of wind farms in New York state³. We convert wind speed to wind energy through the production curve of GE1.5-77 (see Table 5 in the online Appendix H), which specifies the amount of wind energy produced by such a turbine for any given wind speed.

We model the wind speed process Q_t as the sum of an autoregressive (AR) process ξ'_t and a seasonality function $f'(t)$:

$$Q_t = \xi'_t + f'(t).$$

Specifically, ξ'_t is an AR(1) (an AR process of order 1) as follows:

$$\xi'_{t+1} = \phi \xi'_t + \sigma' \epsilon'_t,$$

where ϕ and σ' are scalars, and each $\epsilon'_t \sim N(0, 1)$ is an i.i.d. error term. The seasonality $f'(t)$ is

$$f'(t) = A' + \gamma_4 \cos((t + \omega_4) \cdot 2\pi / (24 \cdot 365)) + \gamma_5 \cos((t + \omega_5) \cdot 2\pi / 24),$$

where A' is a constant level; γ_4 and ω_4 are the magnitude and the phase shift of daily seasonality, respectively; and γ_5 and ω_5 are the magnitude and the phase shift of the hourly seasonality, respectively.

7.2. Calibration

We calibrate the wind speed model using a nonlinear regression, with the estimated parameters shown in Table 4. We test the fit of this calibration by computing MAE in terms of power: We first compute the energy of the actual wind speed series and that of the estimated wind speed series, then sum up the differences between these two energy series, and finally divide the difference by the number of periods. The MAE is 0.145 MW; the turbine generation capacity is 1.5 MW. We also experimented with an AR(2) model; it does not fit the data any better than the AR(1) model.

Even though our analytical model allows price and wind to be correlated, statistical analysis shows that the *stochastic* component of our price series is uncorrelated with the *stochastic* component of the wind data: correlation -0.0043. (The *deterministic* seasonality components of wind and price capture the fact that wind tends to blow most strongly at night, when prices tend to be low.) This lack of correlation is not surprising, because for the considered time period wind energy consisted of only a small proportion of the electricity generation in NYISO.

³ www.windpoweringamerica.gov/pdfs/wind_maps/ny_wind_potential_chart.pdf

Table 4 Estimated Parameters for the Wind Speed Model (MAE = 0.145 MW)

ϕ	A'	σ'	γ_4	γ_5	ω_4	ω_5
0.8813	6.7770	1.6734	1.2558	1.4743	-27.8434	28.7040

7.3. Discretization

In order to use the continuous-space wind speed process in our MDPs, we discretize this process into a grid. We choose a grid rather than a trinomial tree as we have a natural set of wind speed levels: the positive integers that do not exceed 25 (meters/second), which comprise the production curve of the GE15-77 turbine. We denote this set as \mathcal{M} . To find the transition probability between levels $i, j \in \mathcal{M}$, denoted by ρ_{ij} , we match the first two moments of the discretized wind speed model and its continuous counterpart using the method in Miller and Rice (1983):

$$\min_{\rho_{ij}, \forall j \in \mathcal{M}} \left(\text{mean of the discretized} - \text{mean of the continuous} \right)^2, \quad \forall i \in \mathcal{M},$$

$$+ \left(\text{2nd moment of the discretized} - \text{2nd moment of the continuous} \right)^2.$$

We give equal weights to these moment matchings for simplicity; one can easily specify different weights. Likewise, other methods of estimating the transition probabilities can be applied, for example, those in Hoyland and Wallace (2001).

8. Numerical study

We discuss the setup of our numerical study in §8.1, examine the relevance of each feature of our problem in §8.2, and quantify the monetary and energy values of storage in §8.3.

8.1. Setup

We study a wind farm that includes 120 GE1.5-77 turbines, so its total generation capacity is $120 \times 1.5 \text{ MW} = 180 \text{ MW}$ (MW represents megawatt; 1 watt = 1 joule/ second). This wind farm is remote and connected to a wholesale market via a transmission line that has a loss of 3%, that is, $\tau = 3\%$ (DukeEnergy 2011). We vary the wind farm’s transmission capacity (in MW) from 80 MW to 180 MW in steps of 20 MW. (This transmission capacity could be leased from a transmission company, so the capacity of the entire transmission line could be larger.) Co-located with the wind farm is an industrial battery, which can be charged or discharged fully in ten hours (EPRI 2004), so the unsigned charging power capacity (K_1) and discharging power capacity (K_2) is the energy capacity divided by ten hours. We vary the battery’s energy capacity (in MWh; 1 megawatt hour = 3.6 Giga joules) from 200 MWh to 1200 MWh in steps of 200 MWh. The relative size of the WST system is consistent with Denholm and Sioshansi (2009) and Pattanariyankool and Lave (2010). We also varied the round-trip efficiency of the battery, and found that our qualitative managerial

insights remained consistent. Therefore in the results we report, the storage efficiency parameters are fixed at $\alpha = 0.85$, $\beta = 1$, and $\eta = 1$.


As mentioned in §6.1, we assume each period is one hour. We use a time horizon of one month, so the total number of periods is $31 \times 24 = 744$. A given state is four-dimensional: one for inventory, one for the mean-reverting component, one for the jump state, and one for wind speed. The discount factor δ for each stage (period) is 0.99999, corresponding to an annual risk-free interest rate of 1% with continuous compounding (recall that we use risk-neutral valuation).

We numerically evaluate an optimal policy to (3), along with all the heuristics of §4. For all policies, we discretize the inventory set evenly to 101 levels (beyond this level of discretization the results vary little). The value of each policy is the value of the objective function in (2) evaluated at initial inventory zero using the action of this policy for each stage and state. We find the actions of each policy using dynamic programming to solve each corresponding MDP as follows. To compute an optimal policy for the MDP (3), we compute the values of all feasible decision pairs (a_t, g_t) for every discrete state (x_t, w_t, \vec{p}_t) of period t and determine an optimal action. To solve for the triple-threshold policy for (4), for every (w_t, \vec{p}_t) (note that the state variable x_t is excluded) of each period t , we compute an optimal inventory action in two steps: step 1) is to find the three critical inventory levels \underline{X}_t^1 , \underline{X}_t^2 , and \overline{X}_t ; step 2) is to compute the value function of each state (x_t, w_t, \vec{p}_t) using Proposition 2. We compute the dual-threshold policy similarly. For all other policies, see §4.3. We carried out all experiments on a computer with Intel(R) Core(TM) i7-2600K CPU of 3.40GHz with 8 MB cache. On average, the threshold policies could finish within half an hour, while the optimal policy needs about eight hours. All other policies can be computed within a few minutes.

8.2. Value of negative prices, the buying option, information, and optimization

We compare the value of each policy in §4 relative to that of an optimal policy. Figure 4 reports these results for a range of transmission capacity levels given a storage energy capacity of 600 MWh (this corresponds to the proportions of storage energy capacity to generation capacity in Denholm and Sioshansi 2009).

As seen in this figure, the problem of managing a WST system is non-trivial: using the naïve policy can result in a significant loss of optimality, from 8.5% at 80 MW of transmission capacity to 18% at 180 MW of transmission capacity. This gives the value of carrying out optimization. In contrast, almost all the optimal value is captured by the triple-threshold policy: it achieves about 99.5% of the optimal value over all transmission capacity levels. This implies that ignoring negative prices is benign. Furthermore, the triple-threshold policy achieves at least 99% of the optimal value for all the other parameters considered. This near-optimality stem from three factors: 1) negative

Figure 4 Value of Heuristics versus that of an Optimal Policy


Note. Storage energy capacity equals 600 MWh.

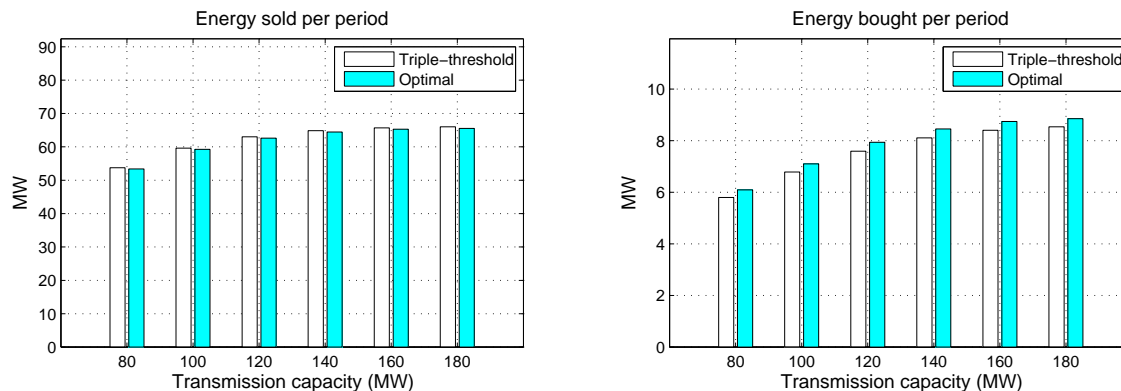
prices are relatively infrequent; 2) negative prices in consecutive periods, one major scenario where the optimal value function may be non-concave as mentioned in §4, are even less frequent; and 3) transmission capacity constrains electricity that could be bought from the market at negative prices, limiting the profitability of the optimal policy at negative prices.

But this triple-threshold policy is *not* optimal. To explore the differences in action between the triple-threshold policy and the optimal policy, we plot the energy sold and energy bought per period for both the triple-threshold policy and the optimal policy in Figure 5. As seen from this figure, the triple-threshold policy sells slightly more, but buys less energy than the optimal policy. This primarily happens when prices are negative. If we refine the triple-threshold policy so that it buys as much as possible (before generating as much as possible to store) at negative prices but follows the triple-threshold policy at positive prices, then we further close the optimality gap: given storage capacity 600 MWh for all the transmission capacity levels considered, this modified policy achieves over 99.95% of the optimal value. Thus the optimal policy looks very similar to a triple-threshold policy plus maximal buying at negative prices.

The value of the buying option, in comparison to the negative prices, is more valuable: Given storage energy capacity of 600 MWh and transmission capacity of 120 MW (corresponding to the proportion in Denholm and Sioshansi (2009)), the dual-threshold policy captures about 96% of the optimal value, so the value of the buying option is around 4%. This corroborates our analysis about the relative values of the triple-threshold policy and the dual-threshold policy in §4.2. The value of the buying option is not overly significant because of the energy loss of 15% and the transmission capacity, which limits the profitability of arbitrage.

Nevertheless, the value of information can be much more substantial: the two-period policy that ignores future information attains about 92% of the optimal value, so the value of including future

Figure 5 The Energy Sold (Left) and Energy Bought (Right) per Period for the Triple-threshold Policy and the Optimal Policy Given 600 MWh of Storage Energy Capacity



information is around 8%; the rolling horizon policy is about 89% of the optimal, so the value of including uncertainty in the optimization is about 11%. This suggests that considering the buying option does not matter, but it is important to consider future information and uncertainty. Values under other parameter settings were comparable.

8.3. The value of storage


We now analyze the monetary and energy values of storage under the triple-threshold policy.

8.3.1. The monetary value of storage. We present the monetary value of storage and then break it down into its three components (defined in §5) for a range of storage and transmission capacity levels in Figure 6. Storage can substantially increase the monetary value of a WST system: for a typical setting when the transmission capacity is 120 MW, and the storage energy capacity is 600 MWh, the monetary value of storage is around 25%, of which 4% is from arbitrage, 11% from time-shifting generation, and 10% from reducing curtailment.

Given a storage energy capacity level of 600 MWh (Figure 6 Left), the monetary value of storage decreases as transmission capacity increases, because the value of reducing curtailment plummets, even though both the value of arbitrage and the value of time-shifting generation increase (see expression (5) in §5). When transmission capacity is tight, the majority of the monetary value of storage is due to reducing curtailment; with ample transmission capacity, the majority of this value is attributable to time-shifting generation and arbitrage.

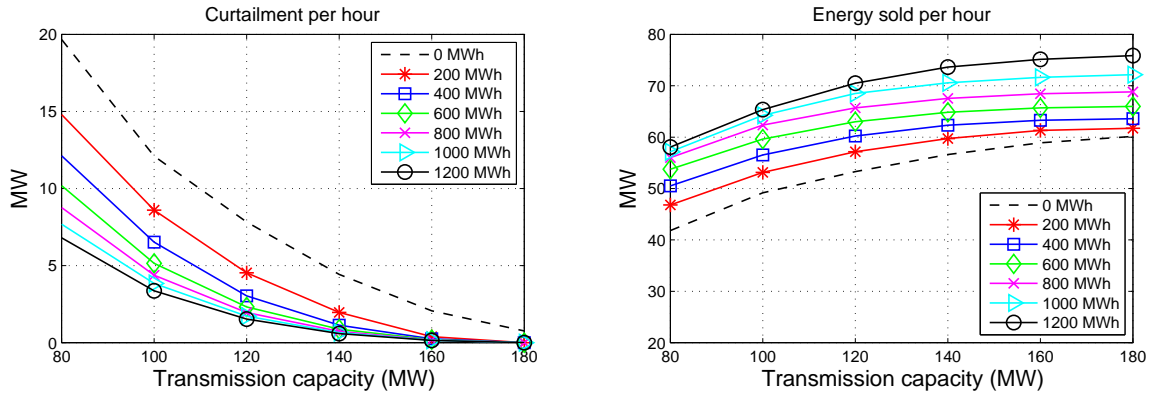
Given a transmission capacity level of 120 MW (Figure 6 Right), the monetary value of storage is increasing concave with storage energy capacity. This is because for any given transmission capacity, the value of the system without storage is constant when changing the storage energy capacity, while all three monetary values of storage (from reducing curtailment, time-shifting, and

Figure 6 The Breakdown of the Monetary Value of Storage Under the Triple-threshold Policy Given 600 MWh of Storage Energy Capacity (Left) and 120 MW of Transmission Capacity (Right)



Note. When transmission capacity is the same as the generation capacity, there may still be curtailment for the case without storage because of negative prices.

Figure 7 Curtailment Per Hour (Left) and Energy Sold Per Hour (Right) under the Triple-threshold Policy for Different Storage Energy Capacity Levels



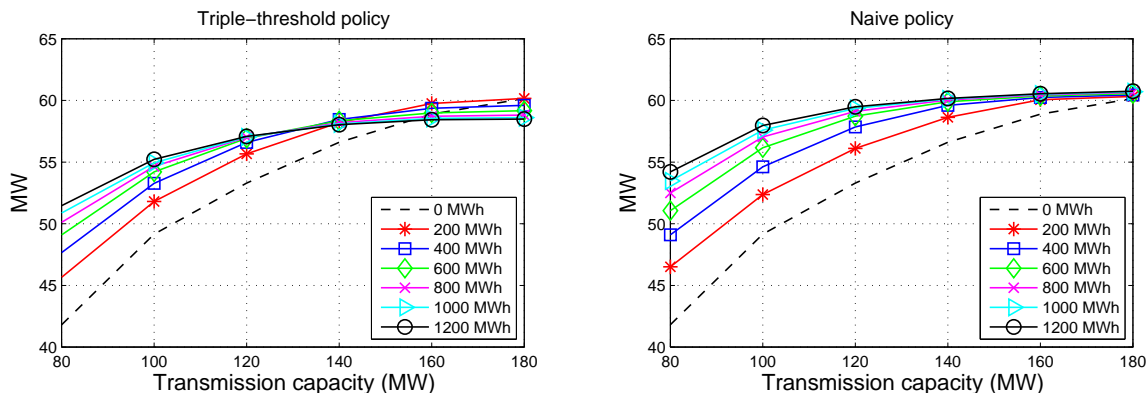
Note. Each curve corresponds to a storage energy capacity level; 0 MWh corresponds to no storage

arbitrage) increase with the storage energy capacity until they level out because of the bottleneck in transmission capacity (see expression (5) in §5).

8.3.2. The energy value of storage. Figure 7 (Left) shows that using storage can substantially reduce curtailment. For instance, when the storage energy capacity is 600 MWh and the transmission capacity is 120 MW, the average curtailment per period (i.e., per hour) is reduced by around 70% compared to the system without storage: 70% of the wind energy curtailed can be recouped through use of storage under our triple-threshold policy. Given any storage energy capacity (including the no-storage case), curtailment shrinks with transmission capacity; and given any transmission capacity, curtailment shrinks with more storage energy capacity.

Figure 7 (Right) indicates that storage can greatly increase the average energy sold per period. For instance, when the storage energy capacity is 600 MWh and the transmission capacity is 120

Figure 8 The Wind Energy Sold Per Hour for Different Storage Energy Capacity Levels Under the Triple-threshold Policy (Left) and the Naive Policy (Right)



Note. Each curve corresponds to a storage energy capacity level; 0 MWh corresponds to no storage.

MW, storage increases the average energy sold by about 18% compared to the no-storage case: 8% is due to selling more wind energy (less curtailment); and 10% is due to reselling electricity bought from the market.

However, under the triple-threshold policy, storage does not necessarily increase the average *wind* energy sold per period. As seen from Figure 8 (Left), when the transmission capacity is low, more storage energy capacity indeed increases the average wind energy sold; however, when the transmission capacity is high, more storage can actually slightly reduce average wind energy sold. This is because at high transmission capacity there is very limited curtailment for the system without storage; the benefit of reducing curtailment with storage decreases to such an extent that this benefit is surpassed by the conversion loss in the storage facility. In contrast, if the naïve policy is used, then more storage energy capacity does increase the average wind energy sold (see Figure 8 right). This observation implies that a policy that strives to maximize revenue (such as the triple-threshold policy) need not maximize the average wind (i.e. renewable) energy sold.

9. Conclusions and future work

We consider the problem of managing a WST system that trades electricity in an electricity market, and we shed light on important aspects of this system. We find that this problem should not be solved using a simple policy, as this may sacrifice significant value compared to the optimal policy. We find that even though negative prices are a unique feature of electricity markets, they do not matter much in operating a typical WST system. We show this by developing a policy that ignores negative prices, and by demonstrating that it is nearly optimal using price and wind energy models calibrated to real data for a range of round-trip efficiency, transmission capacity,

and storage energy capacity levels. We prove that this policy has a triple-threshold structure, which generalizes those known in the commodity storage literature. This structure could be helpful in developing an approximate dynamic programming algorithm (Nascimento and Powell 2008, 2009) if one wanted to use higher-dimensional price models, or if one did not want to discretize prices.

But the triple-threshold policy is not optimal. By modifying it to take advantage of negative prices, we find that it becomes virtually optimal (achieving over 99.95% of the optimal value). This suggests that the optimal policy is consistent with the structure of this enhanced policy.

Furthermore, we show that the buying option is more valuable than negative prices, claiming around 4% of the optimal value. The policy that ignores this option turns out to be a dual-threshold policy, the same as in existing literature. Moreover, the value of information is significant: future information (around 8%) and modeling uncertainty in the optimization (around 11%).

Our experiments also show that storage can substantially increase the monetary value of the system: when transmission capacity is tight, the majority of this value comes from reducing curtailment and time-shifting generation; when transmission capacity is abundant, the majority arises from time-shifting generation and arbitrage. The substantial monetary value of storage—combined with the fact that its cost is dropping rapidly (AquionEnergy 2011)—makes investing in storage potentially appealing. We also find that although storage can substantially reduce curtailment and increase the average energy sold, more storage energy capacity may actually decrease the average wind energy sold if a revenue-maximizing policy is used.

Our work can be extended in several directions. First, one could include bidding in a forward market, in addition to the spot market considered in this paper. Second, one could investigate the collective effect of multiple wind farms and storage facilities on electricity prices, at high levels of wind energy penetration, such as 20% of the total electricity generation in the U.S. (DOE 2008).

Acknowledgments

Nicola Secomandi was supported by NSF grant CMMI1129163. The authors would like to thank Jay Apt, Emily Fertig, Liang LV, and the late Lester Lave for stimulating discussions. We would also like to thank the entire review team for their valuable suggestions that significantly improved this paper.

References

- AquionEnergy. 2011. Low cost, aqueous electrolyte sodium ion energy storage. Available at: www.aquionenergy.com/sites/default/files/user_files/news-press/2011_esa_aquion_whitacre.pdf.
- Barlow, M. T. 2002. A diffusion model for electricity prices. *Mathematical Finance* **12** 287–298.
- Bathurst, G.N., G. Strbac. 2003. Value of combining energy storage and wind in short-term energy and balancing markets. *Electric Power Systems Research* **67** 1–8.

- Benth, F. E., J. Kallsen, T. Meyer-Brandis. 2007. A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance* **14** 153–169.
- Benth, F.E., S. Koekebakker. 2008. Stochastic modeling of financial electricity contracts. *Energy Economics* **30** 1116–1157.
- Boogert, A., C. de Jong. 2008. Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15** 81–98.
- Boyabatli, O. 2011. Supply management in multi-product firms with fixed proportions technology. Working paper, Singapore Management University, Singapore.
- Boyabatli, O., P. R. Kleindorfer, S. R. Koontz. 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science* **57** 1771–1787.
- Brown, P.D., J.A. Peas Lopes, M.A. Matos. 2008. Optimization of pumped storage capacity in an isolated power system with large renewable penetration. *IEEE Transactions on Power Systems* **23** 523–531.
- Bunn, D. 2004. *Modelling Prices in Competitive Electricity Markets*. John Wiley & Sons, England.
- Cahn, A. S. 1948. The warehouse problem. *Bulletin of the American Mathematical Society* **54** 1073.
- Cartea, A., M. G. Figueroa. 2005. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* **12** 313–336.
- Castronuovo, E.D., J.A.P. Lopes. 2004. On the optimization of the daily operation of a wind-hydro power plant. *IEEE Transactions on Power Systems* **19** 1599–1606.
- Charnes, A., J. Dreze, M. Merton. 1966. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* **34** 307–330.
- Chen, Z., P. Forsyth. 2007. A semi-Lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing* **30** 339–368.
- Costa, L.M., F. Bourry, J. Juban, G. Kariniotakis. 2008. Management of energy storage coordinated with wind power under electricity market conditions. *Probabilistic Methods Applied to Power Systems, 2008. PMAPS '08. Proceedings of the 10th International Conference on*. 1–8.
- De Jong, C. 2006. The nature of power spikes: A regime-switch approach. *Studies in Nonlinear Dynamics & Econometrics* **10**.
- Deng, S. 2000. Pricing electricity derivatives under alternative stochastic spot price models. *Hawaii International Conference on System Sciences* **4** 4025.
- Denholm, P., R. Sioshansi. 2009. The value of compressed air energy storage with wind in transmission-constrained electric power systems. *Energy Policy* **37** 3149–3158.
- Devalkar, S.K., R. Anupindi, A. Sinha. 2011. Integrated optimization of procurement, processing and trade of commodities in a network environment. *Operations Research* **59** 1369–1381.

-
- DOE. 2008. 20 percent wind energy by 2030: Increasing wind energy contribution to U.S. electricity supply. Available at www1.eere.energy.gov/windandhydro/pdfs/41869.pdf.
- Duffie, D. 1992. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ.
- DukeEnergy. 2011. Summary of transmission pricing. Available at www.oatiaoasis.com/PNM/PNMdocs/Pricing_Summary_2011_posted_09-06-11.pdf.
- EPRI. 2004. EPRI-DOE handbook supplement of energy storage for grid connected wind generation applications. Tech. rep., Electric Power Research Institute. Available at mydocs.epri.com/docs/public/000000000001008703.pdf.
- ERCOT. 2008. Balancing energy services market clearing prices for energy annual reports archives. Available at www.ercot.com/mktinfo/prices/mcpea.
- Escribano, A, J. I. Pena, P. Villaplana. 2002. Modeling electricity prices: international evidence. Working paper, Universidad Carlos III de Madrid, Madrid, Spain.
- Fanone, E., M. Prokopczuk. 2010. Modeling positive and negative spikes: A non-Gaussian process for intraday electricity prices. Working paper, Università degli Studi di Trieste, Trieste, Italy.
- FERC. 1998. Causes of wholesale electric pricing abnormalities in the Midwest during June 1998. Tech. rep., Federal Energy Regulatory Commission, Washington, DC.
- Fertig, E., J. Apt. 2011. Economics of compressed air energy storage to integrate wind power: A case study in ERCOT. *Energy Policy* **39** 2330–2342.
- Garcia, R.C., J. Contreras, M. van Akkeren, J.B.C. Garcia. 2005. A GARCH forecasting model to predict day-ahead electricity prices. *IEEE Transactions on Power Systems* **20** 867–874.
- Geman, H., A. Roncoroni. 2006. Understanding the fine structure of electricity prices. *Journal of Business* **79** 1225–1261.
- Genoese, F., M. Genoese, M. Wietschel. 2010. Occurrence of negative prices on the German spot market for electricity and their influence on balancing power markets. *Energy Market (EEM), 2010 7th International Conference on the European*. 1–6.
- Gonzalez, A.M., A.M.S. Roque, J. Garcia-Gonzalez. 2005. Modeling and forecasting electricity prices with input/output hidden markov models. *IEEE Transactions on Power Systems* **20** 13–24.
- Gonzalez, J.G., R. Moraga, L.M. Santos, A.M. Gonzalez. 2008. Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. *IEEE Transaction on Power Systems* **23** 460–468.
- Heier, S. 2006. *Grid integration of wind energy conversion systems*. Wiley.
- Hoyland, K., S.W. Wallace. 2001. Generating scenario trees for multistage decision problems. *Management Science* **47** 295–307.
- Huisman, R., R. Mahieu. 2003. Regime jumps in electricity prices. *Energy Economics* **25** 425–434.

- Hull, J., A. White. 1993. One-factor interest-rate models and the valuation of interest-rate derivative securities. *Journal of Financial and Quantitative Analysis* **28** 235–254.
- Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50** 909–921.
- Kim, J. H., W.B. Powell. 2011. Optimal energy commitments with storage and intermittent supply. *Operations Research* **59** 1526–5463.
- Klüppelberg, C., T. Meyer-Brandis, A. Schmidt. 2010. Electricity spot price modelling with a view towards extreme spike risk. *Quantitative Finance* **10** 963–974.
- Knittel, K., M. Roberts. 2005. An empirical examination of restructured electricity prices. *Energy Economics* **27** 791–817.
- Korpaas, M., A.T. Holen, R. Hildrum. 2003. Operation and sizing of energy storage for wind power plants in a market system. *International Journal of Electrical Power and Energy Systems* **25** 599–606.
- Lai, G., F. Margot, N. Secomandi. 2010a. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58** 564–582.
- Lai, G., M. X. Wang, N. Secomandi, S. Kekre, A. Scheller-Wolf. 2010b. Valuation of the real option to store liquefied natural gas at a regasification terminal. *Operations Research* **59** 602–616.
- Löhndorf, N., S. Minner. 2010. Optimal day-ahead trading and storage of renewable energies—an approximate dynamic programming approach. *Energy Systems* **1** 61–77.
- Lucia, J., E. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research* **5** 5–50.
- Miller, A.C., T.R. Rice. 1983. Discrete approximations of probability distributions. *Management Science* **29** 352–362.
- Mokrian, P., M. Stephen. 2006. A stochastic programming framework of the valuations of electricity storage. Working paper, Department of Management Science and Engineering, Stanford University.
- Nasakkala, E., J. Keppo. 2008. Hydropower with financial information. *Applied Mathematical Finance* **15** 1–27.
- Nascimento, J., W. Powell. 2010. Dynamic programming models and algorithms for the mutual fund cash balance problem. *Management Science* **56** 801–815.
- Nascimento, J. M., W. B. Powell. 2008. Optimal approximate dynamic programming algorithms for a general class of storage problems. *Working Paper, Princeton University, Department of Operations Research and Financial Engineering* .
- Nascimento, Juliana M., Warren B. Powell. 2009. An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Math. Oper. Res.* **34** 210–237.
- NOAA. 2010. Integrated surface data, hourly, global. Available at: www.ncdc.noaa.gov/oa/mpp/.

-
- NYISO. 2011. NYISO electricity pricing data. Available at www.nyiso.com.
- Pattanariyankool, S., L.B. Lave. 2010. Optimizing transmission from distant wind farms. *Energy Policy* **38** 2806–2815.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA.
- Powell, W.B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, Hoboken, NJ.
- Rambharat, R., A. Brockwell, D. Seppi. 2005. A threshold autoregressive model for wholesale electricity prices. *Journal of the Royal Statistical Society*. **54** 287–299.
- Rempala, R. 1994. Optimal strategy in a trading problem with stochastic prices. *System Modelling and Optimization, Lecture Notes in Control and Information Sciences*, vol. 197.
- REN21. 2010. Renewables 2010 global status report. Available at www.worldwatch.org/node/6481.
- Schwartz, E., J. E. Smith. 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46** 893–911.
- SECO. 2011. Wind energy transmission. Available at www.seco.cpa.state.tx.us/re_wind-transmission.htm.
- Secomandi, N. 2010a. On the pricing of natural gas pipeline capacity. *Manufacturing & Service Operations Management* **12** 393–408.
- Secomandi, N. 2010b. Optimal commodity trading with a capacitated storage asset. *Management Science* **56** 449–467.
- Seifert, J., M. Uhrig-Homburg. 2007. Modelling jumps in electricity prices: theory and empirical evidence. *Review of Derivatives Research* **10** 59–85.
- Seppi, D. J. 2002. Risk-neutral stochastic processes for commodity derivative pricing: An introduction and survey. *E. Ronn, ed. Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 3–60.
- Sewalt, M., C. de Jong. 2007. Negative prices in electricity markets. Available at www.erasmusenergy.com/articles/91/1/Negative-prices-in-electricity-markets/Page1.html.
- Smith, J. E., K. F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research* **47** 1–15.
- Szkuta, B.R., L.A. Sanabria, T.S. Dillon. 1999. Electricity price short-term forecasting using artificial neural networks. *IEEE Transactions on Power Systems* **14** 851–857.
- Thompson, M., M. Davison, H. Rasmussen. 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics* **56** 226–238.
- Villaplana, P. 2005. Pricing power derivatives: A two-factor jump-diffusion approach. Working paper, InterMoney Energia, Madrid, Spain.

- Wiser, R., M. Bolinger. 2010. 2009 wind technologies market report. Tech. rep., Lawrence Berkeley National Laboratory. Available at: eetd.lbl.gov/ea/ems/re-pubs.html.
- Wu, O., R. Kapuscinski. 2012. Curtailing intermittent generation in electrical systems. Working paper, Ross School of Business, University of Michigan, Ann Arbor, MI.
- Wu, O., D. Wang, Z. Qin. 2012. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing and Service Operations Management* **14** 449–467.
- Xi, X., R. Sioshansi, V. Marano. 2011. A stochastic dynamic programming model for co-optimization of distributed energy storage. Working paper, The Ohio State University.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill, Boston.

Online Appendix

Appendix A: Proof of Lemma 1

Proof: In period $t = T$, we have $V_T^{H1}(x_T, w_T, \vec{p}_T) = 0, \forall x_T$, thus the hypothesis holds for this case.

Suppose this hypothesis holds for all periods $k + 1, \dots, T - 1$. We next prove that for period $t = k$, $V_k^{H1}(x_k^1, w_k, \vec{p}_k) \leq V_k^{H1}(x_k^2, w_k, \vec{p}_k)$ for any $0 \leq x_k^1 < x_k^2$. We achieve this by proving that for any feasible action (a_k^1, g_k^1) in $\Psi(x_k^1, w_k)$, we can always find a feasible action (a_k^2, g_k^2) in $\Psi(x_k^2, w_k)$ such that the objective function in (4) at (a_k^2, g_k^2) is no lower than that at (a_k^1, g_k^1) .

- 1) If $a_k^1 < 0$, then $(a_k^2, g_k^2) = (a_k^1, g_k^1)$: (a_k^2, g_k^2) is feasible as (a_k^1, g_k^1) is feasible (it satisfies constraints C1-C6); the immediate payoff function for (a_k^1, g_k^1) and (a_k^2, g_k^2) are the same, but the resulting inventory level from (a_k^2, g_k^2) is higher than that from (a_k^1, g_k^1) . Because of the hypothesis in period $t = k + 1$, the objective function in (4) at (a_k^2, g_k^2) is no lower than that at (a_k^1, g_k^1) .
- 2) If $a_k^1 \geq 0$, then $(a_k^2, g_k^2) = (a_k^1 - (x_k^2 - x_k^1), [g_k^1 - (x_k^2 - x_k^1)/\alpha]^+)$: (a_k^2, g_k^2) is feasible as (a_k^1, g_k^1) is feasible. The resulting inventory level from (a_k^1, g_k^1) and (a_k^2, g_k^2) are the same ($x_k^2 + a_k^2 = x_k^2 + a_k^1 - (x_k^2 - x_k^1) = x_k^1 + a_k^1$), and the immediate payoff function from (a_k^2, g_k^2) is no lower than that at (a_k^1, g_k^1) . Thus the objective function in (4) at (a_k^2, g_k^2) is no lower than that at (a_k^1, g_k^1) . \square

Appendix B: Proof of Lemma 2

Proof: For each period t and any given state S_t , denote the maximum quantity that one can generate by $\bar{g}_t = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\}$. We prove that $g_t^{H1}(S_t) = \bar{g}_t$ by showing that for any feasible action (a_t, g_t) such that $g_t < \bar{g}_t$, we can always find a feasible solution $(a'_t, g_t + \varepsilon)$, $\varepsilon > 0$, which gives no lower objective value function in (4). Denote q_t as the quantity to sell (if $q_t \geq 0$) or buy (if $q_t < 0$) resulting from a feasible action (a_t, g_t) . Since $g_t < \bar{g}_t$, it follows that either $q_t < C$ or $x_t + a_t < \min\{1, x_t + K_2\}$, thus we consider the following cases:

- 1) If $q_t < 0$, define $\varepsilon = \min\{w_t - g_t, -q_t\}$. Thus action $(a'_t, g'_t) = (a_t, g_t + \varepsilon)$ is feasible, and gives $0 \geq q'_t = q_t + \varepsilon \geq q_t$ (buying less) and the same ending inventory. Hence the objective function at (a'_t, g'_t) is no lower than that at (a_t, g_t) .
- 2) If $0 \leq q_t < C$, define $\varepsilon = \min\{w_t - g_t, C - q_t\}$. Thus action $(a'_t, g'_t) = (a_t, g_t + \varepsilon)$ is feasible, and gives $q'_t = q_t + \varepsilon > q_t \geq 0$ (selling more) and the same ending inventory. Hence the objective function at (a'_t, g'_t) is no less than that at (a_t, g_t) .
- 3) If $q_t = C$, then $x_t + a_t < \min\{1, x_t + K_2\}$. We consider the following two cases:
 - If $a_t \geq 0$, define $\varepsilon = \min\{w_t - g_t, \min\{1 - x_t - a_t, K_2 - a_t\}/\alpha\}$, then $a'_t = a_t + \alpha\varepsilon \leq a_t + \alpha \cdot \min\{1 - x_t - a_t, K_2 - a_t\}/\alpha \leq \min\{1 - x_t, K_2\}$.
 - If $a_t < 0$, define $\varepsilon = \min\{w_t - g_t, -\beta a_t\}$, then $a_t < a'_t = a_t + \varepsilon/\beta \leq a_t - \beta a_t/\beta = 0$.

In both cases, $(a'_t, g_t + \varepsilon)$ is feasible and gives $q'_t = q_t$ and $a'_t > a_t$, hence it results in no lower objective value than (a_t, g_t) does, due to Lemma 1. \square

Appendix C: Proof of Proposition 1

Proof: We first prove finiteness. For each stage the quantity sold cannot exceed C . Thus given any S_t , it holds that $|V_t^{H1}(S_t)| \leq \sum_{k=t}^T |\mathbb{E}[p_k | w_t, \vec{p}_t]| \cdot C \leq \sum_{k=t}^T \mathbb{E}[|p_k| | w_t, \vec{p}_t] \cdot C < \infty$, where the last inequality follows from Assumption 1.

We next prove concavity by induction. For period T , $V_T^{H1}(S_T) = 0$, so the hypothesis holds. Suppose for all periods $k+1, \dots, T-1$, the hypothesis holds. For period $t = k$, according to Lemma 2, $g_k^{H1}(S_k) = \min\{w_k, C + \min\{1 - x_k, K_2\}/\alpha\}$. Substituting $g_k^{H1}(S_k)$ for g_k in (4), we obtain an optimization problem with only one decision variable, a_k , as follows:

$$V_k^{H1}(S_k) = \max_{a_k \text{ s.t. } (a_k, g_k^{H1}(S_k)) \in \Psi(x_k, w_k)} R(a_k, g_k^{H1}(S_k), p_k^+) + \delta \mathbb{E}[V_{k+1}^{H1}(S_{k+1}) | S_k]. \quad (6)$$

We next prove that the set $\mathcal{C} := \{(x_k, a_k) | x_k \in \mathcal{X}, (a_k, g_k^{H1}(S_k)) \in \Psi(x_k, w_k)\}$ is convex: given any (x_k^1, a_k^1) and (x_k^2, a_k^2) in \mathcal{C} , the linear combination $(x_k^\lambda, a_k^\lambda) = (\lambda x_k^1 + (1-\lambda)x_k^2, \lambda a_k^1 + (1-\lambda)a_k^2)$ is also in \mathcal{C} , where $\lambda \in [0, 1]$. It is easy to verify that $(x_k^\lambda, a_k^\lambda)$ is also in \mathcal{C} if: a_k^1 and a_k^2 are both positive ($(x_k^\lambda, a_k^\lambda)$ satisfies constraints C2-C5-C6, or C3-C5-C6); or both negative ($(x_k^\lambda, a_k^\lambda)$ satisfies constraint C1-C5-C6). If one is positive and the other negative, we next show that $(x_k^\lambda, a_k^\lambda)$ is also in \mathcal{C} . Without loss of generality, we assume that $a_k^1 < 0 \leq a_k^2$. Clearly, $(x_k^\lambda, a_k^\lambda)$ satisfies constraint C5 and C6 as (x_k^1, a_k^1) and (x_k^2, a_k^2) are in \mathcal{C} . If a_k^λ is negative, then $(x_k^\lambda, a_k^\lambda)$ satisfies C1: $g_k^{H1}(S_k) - a_k^\lambda \beta \leq g_k^{H1}(S_k) - a_k^1 \beta \leq C$ (the first equality follows as $0 > a_k^\lambda = \lambda a_k^1 + (1-\lambda)a_k^2 \geq \lambda a_k^1 + (1-\lambda)a_k^1 = a_k^1$). If a_k^λ is nonnegative, then $(x_k^\lambda, a_k^\lambda)$ satisfies either C2 or C3: if $0 \leq g_k^{H1}(S_k) - a_k^2/\alpha \leq C$, then $g_k^{H1}(S_k) - a_k^\lambda/\alpha \leq g_k^{H1}(S_k) < g_k^{H1}(S_k) - a_k^1 \beta \leq C$ (the second inequality following from the fact that $a_k^1 < 0$, and the last inequality from the fact that (x_k^1, a_k^1) is in \mathcal{C}), thus C2 is satisfied; if $0 \leq (a_k^2/\alpha - g_k^{H1}(S_k))/\tau \leq C$, then $(a_k^\lambda/\alpha - g_k^{H1}(S_k))/\tau \leq (a_k^2/\alpha - g_k^{H1}(S_k))/\tau \leq C$ (the first inequality following from the fact that $a_k^2 \geq a_k^\lambda$), thus C3 is satisfied. In summary, \mathcal{C} is a convex set.

We now prove that the whole objective function in (6) is concave on \mathcal{C} . Since $p_k^+ \geq 0$, it follows that $-\beta \cdot \tau \cdot p_k^+ \geq -p_k^+ \cdot \tau/\alpha \geq -p_k^+ / (\alpha \cdot \tau)$, so the immediate payoff function $R(\cdot)$ in (6) is concave over a_k . As $R(\cdot)$ is constant in x_k (x_k does not appear in the function itself), it is jointly concave on \mathcal{C} . Furthermore, according to the induction hypothesis, the value function $V_k^{H1}(x_k, w_k, \vec{p}_k)$ is concave in $x_k \in \mathcal{X}$ given any w_k and \vec{p}_k . Since $x_k = (x_k + a_k)\eta$, $V_k^{H1}(x_k, w_k, \vec{p}_k)$ is jointly concave in x_k and a_k given any w_k and \vec{p}_k , and thus $\delta \mathbb{E}[V_{k+1}^{H1}((x_k + a_k)\eta, w_{k+1}, \vec{p}_{k+1}) | w_k, \vec{p}_k]$ is also concave on \mathcal{C} given any w_k and \vec{p}_k as expectation preserves concavity. As a result, the objective function in (6) is concave on \mathcal{C} given any w_k and \vec{p}_k .

Moreover, \mathcal{X} is a convex set, and $\Psi(x_k, w_k)$ is a nonempty set for any $x_k \in \mathcal{X}$. Also since $V_k^{H1}(S_k) < \infty$, according to Theorem A.4 in Porteus (2002), the expression inside the first expectation in (4) is concave on \mathcal{X} given any w_k and \vec{p}_k , and thus $V_k^{H1}(S_k)$ is concave on \mathcal{X} given any w_k and \vec{p}_k .

By the principle of mathematical induction, the hypothesis holds for all periods $t = 1, \dots, T+1$. \square

Appendix D: Proof of Proposition 2

Proof: For each period $t \in \mathcal{T}$, according to Lemma 2, $g_t^{H1}(S_t) = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\}$. Substituting $g_t^{H1}(S_t)$ into (4) for g_t , we obtain the following optimization problem with only one decision variable, a_t :

$$V_t^{H1}(S_t) = \max_{a_t \text{ s.t. } (a_t, g_t^{H1}(S_t)) \in \Psi(x_t, w_t)} R(a_t, g_t^{H1}(S_t), p_t^+) + \delta \mathbb{E}[V_{t+1}^{H1}(S_{t+1}) | S_t]. \quad (7)$$

Define $y_t := x_t + a_t$ for all a_t such that $(a_t, g_t^{H1}(S_t)) \in \Psi(x_t, w_t)$. Let $\underline{\Psi}^1(x_t, w_t)$ denote the feasible set of y_t for all $a_t \geq g_t^{H1}(S_t) \cdot \alpha$; $\underline{\Psi}^2(x_t, w_t)$ the set of y_t for all $0 \leq a_t \leq g_t^{H1}(S_t) \cdot \alpha$; and $\overline{\Psi}(x_t, w_t)$ the set of y_t for all $a_t \leq 0$. Note the extra equality signs added to $a_t > g_t^{H1}(S_t) \cdot \alpha$ and $a_t < 0$, comparing to the cases in (1): they do not change the results of the optimal policy, but facilitate the exposition of the proof later on. We obtain $\underline{\Psi}^1(x_t, w_t)$, $\underline{\Psi}^2(x_t, w_t)$, and $\overline{\Psi}(x_t, w_t)$ by combining constraints C1, C2, C3, C5 and C6 with constraints $a_t \geq g_t^{H1} \cdot \alpha$, $0 \leq a_t \leq g_t^{H1}(S_t) \cdot \alpha$, and $a_t \leq 0$ respectively:

$$\begin{aligned} \underline{\Psi}^1(x_t, w_t) &= \{y_t | g_t^{H1}(S_t) \cdot \alpha + x_t \leq y_t \leq 1; y_t \leq x_t + \alpha(C \cdot \tau + g_t^{H1}(S_t)); K_1 + x_t \leq y_t \leq K_2 + x_t\} \\ &= \{y_t | g_t^{H1}(S_t) \cdot \alpha + x_t \leq y_t \leq 1; y_t \leq x_t + \alpha(C \cdot \tau + g_t^{H1}(S_t)); y_t \leq K_2 + x_t\}, \end{aligned} \quad (8)$$

$$\begin{aligned} \underline{\Psi}^2(x_t, w_t) &= \{y_t | x_t \leq y_t \leq 1; y_t \leq g_t^{H1}(S_t) \cdot \alpha + x_t; y_t \geq \alpha(g_t^{H1}(S_t) - C) + x_t; K_1 + x_t \leq y_t \leq K_2 + x_t\} \\ &= \{y_t | y_t \leq 1; y_t \leq g_t^{H1}(S_t) \cdot \alpha + x_t; y_t \geq \alpha(g_t^{H1}(S_t) - C) + x_t; y_t \leq K_2 + x_t\}, \end{aligned} \quad (9)$$

$$\begin{aligned} \overline{\Psi}(x_t, w_t) &= \{y_t | 0 \leq y_t \leq x_t; y_t \geq (g_t^{H1}(S_t) - C)/\beta + x_t; K_1 + x_t \leq y_t \leq K_2 + x_t\} \\ &= \{y_t | 0 \leq y_t \leq x_t; y_t \geq (g_t^{H1}(S_t) - C)/\beta + x_t; K_1 + x_t \leq y_t\}. \end{aligned} \quad (10)$$

Substituting $a_t = y_t - x_t$ into (4), (4) reduces to finding the maximum of the following three::

$$\max_{y_t \in \underline{\Psi}^1(x_t, w_t)} \left\{ (g_t^{H1}(S_t) - y_t/\alpha + x_t/\alpha)p_t^+/\tau + \delta \mathbb{E} [V_{t+1}^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}, \quad (11)$$

$$\max_{y_t \in \underline{\Psi}^2(x_t, w_t)} \left\{ (g_t^{H1}(S_t) - y_t/\alpha + x_t/\alpha)p_t^+ \cdot \tau + \delta \mathbb{E} [V_{t+1}^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}, \quad (12)$$

$$\max_{y_t \in \overline{\Psi}(x_t, w_t)} \left\{ (g_t^{H1}(S_t) - y_t/\beta + x_t/\beta)p_t^+ \cdot \tau + \delta \mathbb{E} [V_{t+1}^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}. \quad (13)$$

These three problems can be thought of as the problem of buying, the problem of generating and storing, and the problem of selling respectively. Next, we relax (11), (12) and (13) by the following three steps:

- removing the last constraint in the feasible set of (11), (12) and (13), which is equivalent to removing the charging and discharging power capacity constraints in C6;
- removing the second to last constraint in the feasible set of (11), (12) and (13), which is equivalent to removing the transmission capacity constraints C3, C2 and C1 respectively;
- in (11) setting $x_t = 0$ and $w_t = 0$, so $\underline{\Psi}^1(x_t, w_t)$ becomes $[0, 1]$;
- in (12) setting $x_t = 0$ and setting w_t to be arbitrary large (which means $g_t^{H1}(S_t) = C + \min\{1 - x_t, K_2\}/\alpha$), so that $\underline{\Psi}^2(x_t, w_t)$ becomes $[0, 1]$;
- in (13) setting $x_t = 1$, so $\overline{\Psi}(x_t, w_t)$ becomes $[0, 1]$.

Meanwhile, we remove constant terms from (11), (12) and (13), and obtain the following problems:

$$\max_{y_t \in [0, 1]} \left\{ -y_t p_t^+ / (\alpha \cdot \tau) + \delta \mathbb{E} [V_{t+1}^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}, \quad (14)$$

$$\max_{y_t \in [0, 1]} \left\{ -y_t p_t^+ \cdot \tau / \alpha + \delta \mathbb{E} [V_t^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}, \quad (15)$$

$$\max_{y_t \in [0, 1]} \left\{ -y_t p_t^+ \cdot \tau \cdot \beta + \delta \mathbb{E} [V_t^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t] \right\}. \quad (16)$$

Denote the optimal solutions to (14), (15) and (16) by \underline{X}_t^1 , \underline{X}_t^2 and \overline{X}_t respectively, which are all in $[0, 1]$. We next show that $\underline{X}_t^1 \leq \underline{X}_t^2 \leq \overline{X}_t$.

Define $U_t(y_t, w_t, \vec{p}_t) := \delta \mathbb{E} [V_{t+1}^{H1}(y_t \cdot \eta, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t]$ and denote its derivative with respect to y_t over $[0, 1]$ by $U_t'(y_t, w_t, \vec{p}_t)$. (In the case when $U_t(y_t, w_t, \vec{p}_t)$ is piecewise linear, such as when both the price and

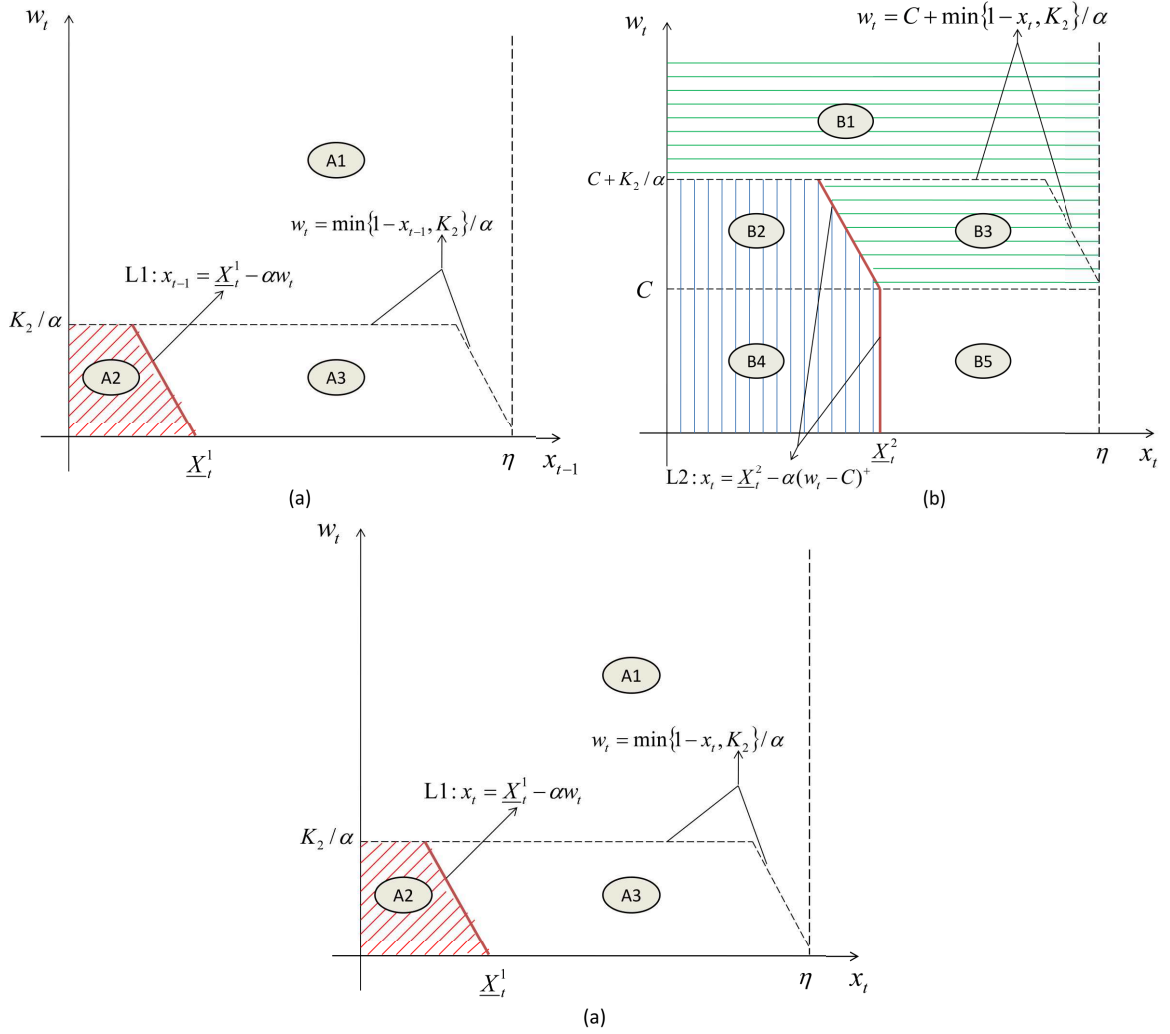


Figure 9 The Proof of Proposition 2: Use a Special Case as an Illustration

wind energy processes follow discrete distributions, define $U_t'(y_t, w_t, \vec{p}_t)$ as the right derivative at $y_t = 0$, and the left derivative over $y_t \in (0, 1]$.) According to Proposition 1, $V_t^{H1}(y_t \cdot \eta, w_t, \vec{p}_t)$ is concave in $y_t \cdot \eta$ given any w_t and \vec{p}_t , thus $U_t(y_t, w_t, \vec{p}_t)$ is concave in y_t given any w_t and \vec{p}_t , and hence $U_t'(y_t, w_t, \vec{p}_t)$ is non-increasing in y_t . Since $-p_t^+ / (\alpha \cdot \tau) \leq -p_t^+ \cdot \tau / \alpha \leq -p_t^+ \cdot \tau \cdot \beta$, it follows that $\underline{X}_t^1 \leq \underline{X}_t^2 \leq \bar{X}_t$.

These three thresholds characterize the optimal policies for (11), (12), and (13) respectively. The objective functions in (11), (12), and (13) are concave in y_t and their relaxed problems in (14), (15), and (16) achieve their global maxima at \underline{X}_t^1 , \underline{X}_t^2 , and \bar{X}_t , so the optimal action for (11), (12), and (13) is to move as close as possible to \underline{X}_t^1 , \underline{X}_t^2 , and \bar{X}_t in their corresponding feasible set.

We next show the optimal solution for each of the three problems on the (x_t, w_t) plane, as in Figure 9.

D.1. The optimal action for (11): the problem of buying

For (11), observe that the feasible set for y_t is $\underline{\Psi}^1(x_t, w_t) = [x_t + g_t^{H1}(S_t) \cdot \alpha, \min\{x_t + K_2, x_t + \alpha(C\tau + g_t^{H1}(S_t)), 1\}]$. This set can be either empty or nonempty in the following two cases:

(i1) if $w_t > \min\{1 - x_t, K_2\}/\alpha$ (see region A1 in Figure 9(a)): $g_t^{H1}(S_t) = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\} > \min\{1 - x_t, K_2\}/\alpha$, so the left end of $\underline{\Psi}^1(x_t, w_t)$ is greater than its right end, indicating that $\underline{\Psi}^1(x_t, w_t)$ is empty. Intuitively, since the amount generated exceeds what can be stored, the optimal action for (11) for this case is not to buy.

(i2) if $w_t \leq \min\{1 - x_t, K_2\}/\alpha$: $g_t^{H1}(S_t) = w_t$, and the feasible set becomes $[x_t + w_t \cdot \alpha, \min\{x_t + K_2, x_t + \alpha(C\tau + w_t), 1\}]$. Relative to the left end point of this feasible set, \underline{X}_t^1 can be either to this point's right, or to its left. Recall that the optimal action for (11) is to bring the inventory level as close as possible to \underline{X}_t^1 , so for each of these two cases, we can obtain the optimal action for (11) as follows: on the right ($x_t + w_t \cdot \alpha \leq \underline{X}_t^1$; see region A2 in Figure 9(a)), the optimal y_t for (11) equals $\min\{\underline{X}_t^1, x_t + K_2, x_t + \alpha(C\tau + w_t), 1\}$, i.e., a_t equals $\min\{\underline{X}_t^1 - x_t, K_2, \alpha(C\tau + w_t), 1 - x_t\} = \min\{\underline{X}_t^1 - x_t, K_2, \alpha(C\tau + w_t)\}$. When \underline{X}_t^1 is on the left ($\underline{X}_t^1 < x_t + w_t \cdot \alpha$; see region A3 in Figure 9(a)), the optimal y_t for (11) equals $x_t + w_t \cdot \alpha$, i.e., $a_t = w_t \cdot \alpha$;

D.2. The optimal action for (12): the problem of generating and storing

For problem (12), we consider the following three cases:

(ii1) if $w_t \geq C + \min\{1 - x_t, K_2\}/\alpha$ (see region B1 in Figure 9(b)): $g_t^{H1}(S_t) = C + \min\{1 - x_t, K_2\}/\alpha$, then it is straightforward that $a_t^{H1}(S_t)$ is $\alpha(g_t^{H1}(S_t) - C)^+$, i.e., generate as much electricity as the transmission lines can transport and the storage facility can charge.

(ii2) if $C \leq w_t < C + \min\{1 - x_t, K_2\}/\alpha$: $g_t^{H1}(S_t) = w_t$. Since the generated amount exceeds the transmission capacity, then the minimum quantity that we have to store in the inventory is the excess ($g_t^{H1}(S_t) - C$), thus the optimal a_t for this case is at least $\alpha(g_t^{H1}(S_t) - C) = \alpha(w_t - C)$. Observe in (12) that the feasible set for y_t is $[x_t + \alpha(w_t - C), x_t + \min\{\alpha \cdot w_t, K_2, 1 - x_t\}]$. Relative to the left end point of this set, \underline{X}_t^2 can either lie on its right, or on its left. For each of these two cases, we can obtain the optimal action for (12) as follows (recall that the optimal action for (12) is to bring the inventory level as close as possible to \underline{X}_t^2): when \underline{X}_t^2 is on the right ($x_t + \alpha(w_t - C) \leq \underline{X}_t^2$; see region B2 in Figure 9(b)), the optimal y_t for (12) for this case equals $\min\{\underline{X}_t^2, x_t + \min\{\alpha \cdot w_t, K_2, 1 - x_t\}\}$, so the corresponding a_t equals $\min\{\underline{X}_t^2 - x_t, \min\{\alpha \cdot w_t, K_2, 1 - x_t\}\} = \min\{\underline{X}_t^2 - x_t, \alpha \cdot w_t, K_2\}$. When \underline{X}_t^2 lies on the left ($\underline{X}_t^2 < x_t + \alpha(w_t - C)$; see region B3 in Figure 9(b)), the optimal y_t for (12) for this case equals $x_t + \alpha(w_t - C)$, i.e., $a_t = \alpha(w_t - C)$.

(ii3) $w_t < C$: $g_t^{H1}(S_t) = w_t$. The feasible set for y_t is $[x_t, x_t + \min\{\alpha \cdot w_t, K_2, 1 - x_t\}]$. Similar to the argument in (ii2), \underline{X}_t^2 can fall either on the left or right of the end point of this set. If $x_t \in [0, \underline{X}_t^2]$ (see region B4 in Figure 9(b)), then the optimal action for this case is to generate and store as much as possible to reach \underline{X}_t^2 , i.e., $a_t = \min\{\underline{X}_t^2 - x_t, \alpha \cdot w_t, K_2\}$; if $x_t \in (\underline{X}_t^2, \eta]$ (see region B5 in Figure 9(b)), the optimal action is to keep the inventory unchanged by selling all generated electricity, i.e. $a_t = 0$.

Note that the optimal action in B1 and B3 for (12) are the same: generate as much possible, sell quantity C to the market, and then store the rest. Thus, we combine these two regions and express their unified formula for the optimal action: if (x_t, w_t) satisfies either of the following two conditions: 1) $w_t \geq C + \min\{1 - x_t, K_2\}/\alpha$; or 2) $C \leq w_t < C + \min\{1 - x_t, K_2\}/\alpha$ and $x_t > \underline{X}_t^2 - \alpha(w_t - C)^+$, then the optimal action a_t for this case equals $\alpha(g_t^{H1}(S_t) - C)^+$.

Also note that the optimal action in B2 and B4 for (12) are the same: store generated electricity as much as possible to reach \underline{X}_t^2 , and sell the rest of the generated electricity. Thus, we can combine these two regions and express their unified formula for the optimal action: if (x_t, w_t) satisfies the following condition $x_t + \alpha(w_t - C)^+ \leq \underline{X}_t^2$ and $w_t < C + \min\{1 - x_t, K_2\}/\alpha$, then the optimal action a_t for this case equals $\min\{\underline{X}_t^2 - x_t, \alpha \cdot w_t, K_2\}$.

D.3. The optimal action for (13): the problem of selling

For (13), we consider the following two cases:

(iii1) if $w_t > C$ (see region C1 in Figure 9(c)): $g_t^{H1}(S_t) = \min\{w_t, C + \min\{1 - x_t, K_2\}/\alpha\} > C$, the quantity generated exceeds transmission capacity, thus we cannot sell.

(iii2) if $w_t \leq C$: $g_t^{H1}(S_t) = w_t$. Recall that the optimal action for (13) is to sell to bring the inventory level as close as possible to \bar{X}_t in its corresponding feasible set, so if $x_t \in [0, \bar{X}_t)$ (see region C2 in Figure 9(c)), the optimal action for problem (13) is to sell nothing from inventory. Otherwise, if $x_t \in [\bar{X}_t, \eta]$ (see region C3 in Figure 9(c)), the optimal action for problem (13) is to sell down as much as possible to inventory level \bar{X}_t , i.e., $a_t = \max\{\bar{X}_t - x_t, (w_t - C)^-/\beta, K_1\}$.

D.4. The optimal action for (7)

We next find the optimal solution to the original problem (7) by combining the optimal solution for (11), (12), and (13) for each region mentioned in §D.1, §D.2, and §D.3.

A) If (x_t, w_t) satisfies either of the following two conditions: 1) $w_t \geq C + \min\{1 - x_t, K_2\}/\alpha$; or 2) $C \leq w_t < C + \min\{1 - x_t, K_2\}/\alpha$ and $x_t > \underline{X}_t^2 - \alpha(w_t - C)^+$, then the optimal solution to (12) is $a_t^{H1}(S_t) = \alpha(g_t^{H1}(S_t) - C)^+$, which is also the optimal solution to (7). This is because the optimal solutions to (11) and (13) are also feasible solutions to (12).

B) If (x_t, w_t) satisfies the first of the following two conditions, but not the second: 1) $x_t + \alpha(w_t - C)^+ \leq \underline{X}_t^2$, $w_t < C + \min\{1 - x_t, K_2\}/\alpha$, but not 2) $w_t \leq \min\{1 - x_t, K_2\}/\alpha$ and $x_t + w_t \cdot \alpha \leq \underline{X}_t^1$, then the optimal solution to (12) is $a_t^{H1}(S_t) = \min\{\underline{X}_t^2 - x_t, \alpha \cdot w_t, K_2\}$, which is also the optimal solution to (7). This is because the optimal solutions to (11) and (13) are also feasible solutions to (12).

C) If $w_t \leq \min\{1 - x_t, K_2\}/\alpha$ and $x_t + w_t \cdot \alpha \leq \underline{X}_t^1$, then the optimal solution to (11) is $a_t^{H1}(S_t) = \min\{\underline{X}_t^1 - x_t, \alpha(C\tau + w_t), K_2\}$, which is also the optimal solution to (7). This is because the optimal solutions to (12) and (13) are also feasible solutions to (11).

D) If $w_t \leq C$ and $\underline{X}_t^2 \leq x_t \leq \bar{X}_t$, then the optimal solution to (7) is $a_t^{H1}(S_t) = 0$, because it is the optimal solution to (11), (12), and (13).

E) If $w_t \leq C$ and $x_t > \bar{X}_t$, then the optimal solution to (13) is $a_t^{H1}(S_t) = \max\{\bar{X}_t - x_t, (w_t - C)^-/\beta, K_1\}$, which is also the optimal solution to (7). This is because the optimal solution to (11) and (12) is also a feasible solution to (13). \square

Appendix E: Proof of Lemma 3

Proof: When $p_t \geq 0$, the proof is the same as that for Lemma 2. When $p_t < 0$, both the immediate payoff function $R(\cdot)$ and the continuation value function in the modified MDP (3) (without buying) reach their maxima when the ending inventory level reaches the highest level possible: the immediate payoff function reaches its maximum if the quantity to sell is zero, a decision leaving the highest ending inventory level possible; the continuation value function in the modified MDP (3) (without buying) is non-decreasing in the ending inventory level because the value function in period t is non-decreasing in inventory level. As a result, the optimal action is to sell nothing and store as much as possible, that is $a_t^{H2}(S_t) = g_t^{H2}(S_t) \cdot \alpha$. Adding these extra two constraints to the feasible set, it is easy to show that $g_t^{H2}(S_t) = \min\{w_t, (1 - x_t)/\alpha, K_2/\alpha\}$. Note that there may be other optimal solutions, but we focus on the optimal solution that gives the highest ending inventory level.

Appendix F: Proof of Proposition 3

Proof: We prove concavity by induction. For period T , $V_T^{H2}(S_T) = 0$, so the hypothesis holds. Suppose for all periods $k + 1, \dots, T - 1$, the hypothesis holds. For period $t = k$, if $p_k \geq 0$, the proof is a special case of that of Proposition 1. If $p_k < 0$, according to Lemma 3, $g_k^{H2}(S_k) = \min\{w_k, (1 - x_k)/\alpha, K_2/\alpha\}$, and thus $a_k^{H2}(S_k) = \alpha \cdot \min\{w_k, (1 - x_k)/\alpha, K_2/\alpha\}$. Substituting both $a_k^{H2}(S_k)$ and $g_k^{H2}(S_k)$ into (3), the objective function is concave in x_k given w_k, \vec{p}_k : the immediate payoff function is zero and thus concave in x_k ; the continuation value function $\mathbb{E}[\delta V_{k+1}^{H2}((x_k + a_k^{H2}(S_k))\eta, w_{k+1}, \vec{p}_{k+1}) | w_k, \vec{p}_k]$ is concave in x_k given any \vec{p}_k and w_k because $V_{k+1}^{H2}((x_k + a_k^{H2}(S_k))\eta, w_k, \vec{p}_k)$ is concave in $(x_k + a_k^{H2}(S_k))\eta$ given any \vec{p}_k and w_k according to the hypothesis, and expectation preserves concavity. Therefore, $V_k^{H2}(x_k, w_k, \vec{p}_k)$ is concave on \mathcal{X} given any w_k, \vec{p}_k . \square

Appendix G: Proof of Proposition 4

Proof: 1) this follows directly from Lemma 3; 2) the proof is straightforward from that of Proposition 2 and thus is omitted. \square

Appendix H: Production curve

The following table is adapted from www.inl.gov/wind/software/powercurves/pc_ge_wind.xls

Table 5 Production Curve of GE1.5MW

Speed (m/s)	0	1	2	3	4	5	6	7	8	9	10	11	12
Power (MW)	0	0	0	0	0.043	0.131	0.25	0.416	0.64	0.924	1.181	1.359	1.436
Speed (m/s)	13	14	15	16	17	18	19	20	21	22	23	24	25
Power (MW)	1.481	1.494	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5

References

Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA.