Research Collection Lee Kong Chian School Of Business   Lee Kong Chian School of Business

4-2019

Managing wind-based electricity generation in the presence of storage and transmission capacity

Yangfang (Helen) ZHOU
Singapore Management University, helenzhou@smu.edu.sg

Alan SCHELLER-WOLF
Carnegie Mellon University

Nicola SECOMANDI
Carnegie Mellon University

Stephen SMITH
Carnegie Mellon University

Follow this and additional works at: https://ink.library.smu.edu.sg/lkcsb_research

Part of the Environmental Sciences Commons, and the Operations and Supply Chain Management Commons

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

Yangfang (Helen) Zhou

Lee Kong Chian School of Business, Singapore Management University, Singapore, helenzhou@smu.edu.sg

Alan Scheller-Wolf, Nicola Secomandi

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

Stephen Smith

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

We investigate the management of a merchant wind energy farm co-located with a grid-level storage facility and connected to a market through a transmission line. We formulate this problem as a Markov decision process (MDP) with stochastic wind speed and electricity prices. Consistent with most deregulated electricity markets, our model allows these prices to be negative. As this feature makes it difficult to characterize any optimal policy of our MDP, we show the optimality of a *stage- and partial-state-dependent-threshold* policy when prices can only be positive. We extend this structure when prices can also be negative to develop heuristic one (H1) that approximately solves a stochastic dynamic program. We then simplify H1 to obtain heuristic two (H2) that relies on a *price-dependent-threshold* policy and derivative-free deterministic optimization embedded within a Monte Carlo simulation of the random processes of our MDP. We conduct an extensive and data-calibrated numerical study to assess the performance of these heuristics and variants of known ones against the optimal policy, as well as to quantify the effect of negative prices on the value added by and environmental benefit of storage. We find that (i) H1 computes an optimal policy and on average is about 17 times faster to execute than directly obtaining an optimal policy; (ii) H2 has a near optimal policy (with a 2.86% average optimality gap), exhibits a two orders of magnitude average speed advantage over H1, and outperforms the remaining considered heuristics; (iii) storage brings in more value but its environmental benefit falls as negative electricity prices occur more frequently in our model.

Key words: Wind-based electricity generation; energy storage; negative electricity prices; Markov decision process; real options

1. Introduction

The last fifteen years have seen a boom in global wind-based electricity production—in the United States (US) alone wind generation capacity has grown more than eightyfold from 2000 to 2015 (Wiser and Bolinger 2015). This global trend will probably continue, as many countries have enacted policies to promote wind energy (REN21 2010). For example, the US recently renewed its Federal Renewable Energy Production Tax Credit through 2019 (DOE 2017). To support this

projected growth, wind farm merchants have been considering co-locating electricity generation and grid-scale storage facilities, such as industrial batteries. An example of such a system is the demonstration project developed by the AES Corporation in West Virginia, where a 32 megawatt (MW) lithium-ion battery supports a wind farm of 98 MW (Sustainable Business News 2014).

Co-location of grid-level storage with a merchant wind energy farm can create value by (i) stockpiling production that exceeds the capacity of the transmission lines that connect (typically remote) such farms to electricity markets; (ii) time-shifting sales when electricity prices are low or even negative (a unique feature of many electricity markets, discussed below); and (iii) enabling the purchase of electricity for future resale. Storage may also benefit the environment by reducing the curtailment of wind energy, which is a significant issue; e.g., the lack of transmission capacity decreased the generation of wind energy by 17% in the Electric Reliability Council of Texas (ERCOT) in 2009 (Wiser and Bolinger 2015) and by more than 10% in China during the 2014-15 time span (The Economist 2015).

Claiming the potential value created by grid-level storage for wind farms requires developing effective operating policies. Unfortunately, the presence of negative electricity prices complicates the management of wind-energy-production, storage, and transmission (WST) systems. For example, even in the absence of generation, when negative prices are possible the optimal policy for the merchant management of electricity storage is known only in the restricted case of fast storage with abundant transmission capacity (Zhou et al. 2016). Though still infrequent in most power markets, negative prices have been observed in the markets run by the New York Independent System Operator (NYISO 2013) and ERCOT (ERCOT 2012, Huntowski et al. 2012), as well as the Nordic Power Exchange (Sewalt and De Jong 2003) and the European Energy Exchange (Genoese et al. 2010, Fanone et al. 2013). One potential cause of such prices is the high costs of ramping conventional power plants up or down: Plant operators may try to avoid these expenses by paying others to consume their excess power (Knittel and Roberts 2005, Sewalt and De Jong 2003, Genoese et al. 2010).

In this paper we develop and evaluate heuristics for operating a WST system in the presence of negative prices and examine the effect of such prices on the value added by and environmental benefit of storage. Specifically, we model a WST system as a finite-horizon Markov decision process (MDP) with stochastic wind availability and electricity prices. In contrast to the difficulty of characterizing any optimal policy of our MDP when prices can be negative, which forces us to rely on stochastic dynamic programming for computing it, we establish the optimality of a *stage- and partial-state-dependent-threshold* policy when prices can only be positive. We develop heuristic one

(H1) based on an extension of this structure when prices can also be negative and approximate stochastic dynamic programming. Seeking a simpler heuristic, we modify H1 to obtain heuristic two (H2) that uses a *price-dependent-threshold* policy and a derivative-free deterministic optimization technique (the Nelder-Mead simplex method; Lagarias et al. 1998) embedded within a Monte Carlo simulation of the stochastic processes of our MDP.

We conduct an extensive numerical study by choosing wind speed and electricity price models that allow us to compute an optimal policy and execute H1. The former model is similar to the ones employed by Kim and Powell (2011) and Wu and Kapuscinski (2013). The latter model combines variants of the ones in Lucia and Schwartz (2002), Seifert and Uhrig-Homburg (2007), and Schneider (2012). We calibrate these models to wind data for Buffalo, NY, and price data for NYISO. Our instances involve making decisions every five minutes during one week. We find that:

(i) H1 computes an optimal policy, even when negative prices occur as frequently as 20% of the time. (Although we exhibit an example in which H1 does not obtain an optimal policy, this instance is pathological.) This observed optimality suggests that H1 may yield an optimal policy in most practical settings. On average, H1 can be executed in about 33 minutes, whereas computing an optimal policy explicitly takes approximately ten hours.

(ii) H2 has a close to optimal policy—with average and maximal optimality gaps equal to 2.86% and 6.49%, respectively—and is two orders of magnitude faster to execute than H1, with an average running time of under 15 seconds. Hence, despite the suboptimality of its policy, H2 is potentially more practical than H1. In addition, H2 outperforms, in terms of both optimality gaps of their policies and execution times, versions of heuristics from the literature that use policies based on knowledge of only price, rather than also inventory.

(iii) The value and environmental benefit of storage are fairly sensitive to changes in the parameters that determine the frequency of occurrence of negative prices in our model. Specifically, as we make negative prices occur more frequently, the value grows—due to buying negatively priced electricity—and the environmental benefit shrinks—as a result of the drop in the amount of reduced curtailment.

We review the literature in §2. We present our MDP in §3. In §4 we determine the structure of an optimal policy for this MDP for the special case when prices are always positive. In §5 we extend this structure—as an approximation—to the case when prices can also be negative to obtain H1, and present H2 and other heuristics. We discuss our numerical study in §6 and conclude in §7. Appendix A details the computation of the policies we analyze. An online appendix includes all the proofs of the results stated in the main text of this paper and an example in which H1 does not give an optimal policy.

2. Literature review

Denholm and Sioshansi (2009) and Fertig and Apt (2011) consider the interplay of generation, storage, and transmission capacity in electricity systems: The former studies how to best locate storage when transmission capacity is scarce; the latter investigates the optimal sizing of storage and transmission capacity. In contrast, we study how to manage a given WST system. However, our numerical study includes a variant of a heuristic available in Fertig and Apt (2011).

Other authors examine the use of storage for wind farms. Brown et al. (2008) focus on how to satisfy the demand 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 merchant wind-hydro system. Korpaas et al. (2003) and Harsha and Dahleh (2015) consider a wind-storage system that serves load as well as trades in a wholesale electricity market. The models of these papers ignore transmission capacity constraints, unlike our model.

Wu and Kapuscinski (2013) 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. We instead take the perspective of a merchant wind farm generator, considering the use of co-located storage. Xi et al. (2014) optimize the use of an electricity storage facility to trade in both electricity energy and ancillary markets. In contrast, we consider how storage can be used to support wind energy production. Hu et al. (2015), Kök et al. (2018), and Aflaki and Netessine (2017) examine capacity investment decisions in renewable energy technologies without storage, whereas we study operating policies to manage WST systems.

Another stream of work centers on how wind farm managers can use storage to make better bidding decisions in a market (e.g., Bathurst and Strbac 2003, Costa et al. 2008, Gonzalez et al. 2008, Löhndorf and Minner 2010, Kim and Powell 2011, and Jiang and Powell 2015). We do not consider bidding, assuming that any electricity offered to the market is accepted. This assumption is realistic: Many electricity markets in the US treat wind generators as "must-run" in normal conditions (Wiser and Bolinger 2013) and 38% of the wind capacity developed in the US in 2009 was sold through merchant/quasi-merchant agreements that do not involve bidding (Wiser and Bolinger 2013).

Our paper is also 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 (threshold) policy. Rempala (1994) and Secomandi (2010) extend this work to incorporate limits on the rates at which the commodity inventory can be adjusted. (See also Secomandi and Seppi 2014, chapter 5.) Other related work includes Mokrian and Stephen (2006),

Chen and Forsyth (2007), Boogert and de Jong (2008), Thompson et al. (2009), Lai et al. (2010), Devalkar et al. (2011), Wu et al. (2012), Nadarajah et al. (2015), Secomandi (2015), and Secomandi et al. (2015). Different from those studied by these authors, our model has a random inflow (wind). Related settings with such intake include hydropower generation (Nasakkala and Keppo 2008) and liquified natural gas (LNG) regasification (Lai et al. 2011). These systems differ from ours, which stores its output (electricity), in that they store their input (water or LNG). Thus, the operating policies in these papers feature only sell-down thresholds, whereas the policies associated with our heuristics have additional buy-and-store-up-to and generate-and-store-up-to thresholds. Boyabatlı et al. (2017) consider agricultural commodity processing with output storage. The primary focus of their work is on capacity investment, so their operating model and resulting policy are simpler than ours.

Zhou et al. (2016) is, to the best of our knowledge, the only other paper that studies the implication of negative prices on the management and value of energy storage. However, whereas it compares electricity storage and disposal strategies, the main focus of this research is to develop and evaluate heuristics to manage a WST system. Further, the work of Zhou et al. (2016) uses the electricity price model calibrated here, which was presented in an earlier version of our paper. This model modifies the one in Schneider (2012) to more realistically generate prices for markets with both negative prices and price spikes (short-lived price jumps).

Under the H1 policy the end-of-period inventory level can be a non-monotonic function of the beginning-of-period inventory availability. This feature of this policy appears to be unique in the literature on energy storage. The H2 policy differs from suitable variants of the policies presented in Graves et al. (1999), Fertig and Apt (2011), and Powell and Meisel (2016b) because its decision rules depend on both the electricity price and the inventory level, instead of only this price. The use of approximate stochastic dynamic programming to compute the H1 policy resembles the use of this technique to obtain energy storage policies in Lai et al. (2010), Lai et al. (2011), and Nadarajah et al. (2015). H2 is an example of a heuristic that applies direct search methods to tune the values of the parameters that define its policy (Powell and Meisel 2016b).

3. Model

We consider the operation of a WST system: A remote wind farm is co-located with a storage facility, both of which are connected to an electricity market via a transmission line (Figure 1). The merchant managing this system can thus buy and sell electricity in the market. We assume that the WST system is small relative to the market, so the merchant decisions do not affect market

Figure 1   System overview

prices. The merchant makes operational and trading decisions periodically over a finite horizon, i.e., at each time t in the finite set $\mathcal{T} := \{0, 1, \ldots, T-1\}$. In particular, electricity trading occurs at the beginning of a period. Further, simultaneously buying and selling the same quantity at the same price, known as a "wash trade," is illegal in commodity markets.[1] Thus, the transmission line can transmit power in only one direction in any given time period. Any electricity left in the storage facility at the terminal time T is worthless. We neglect the deterioration of the storage facility due to charging and discharging (see, e.g., Guo et al., 2017).

Parameters. We assume the storage facility is finite in *energy* capacity and *power* capacity. If we think of this facility as a warehouse for electricity, the energy capacity is analogous to its space and the power capacity represents the maximal rate at which its inventory can be modified. For the rest of this paper, any capacity should be interpreted as power capacity unless specified otherwise. We use the following parameters:

- C^{S}: Energy capacity of the storage facility (in energy units); $C^{\mathrm{S}} > 0$.
- $C^{\mathrm{C}}, C^{\mathrm{D}}$: Charging, discharging capacity (in energy units/period); $C^{\mathrm{C}}, C^{\mathrm{D}} > 0$.
- $C^{\mathrm{G}}, C^{\mathrm{T}}$: Generation, transmission capacity (in energy units/period); $C^{\mathrm{G}} \geq 0$, $C^{\mathrm{T}} > 0$. The transmission capacity represents the part of the capacity of the transmission line contracted by the wind farm. (See Duke Energy, 2017 for an example of such a contract, which is not exclusively available to wind farms.) We realistically assume that $C^{\mathrm{G}} + C^{\mathrm{D}} \geq C^{\mathrm{T}}$: In practice, the contracted transmission capacity is typically smaller than the sum of the wind farm generation capacity and the storage discharging capacity, because wind energy is intermittent (if $C^{\mathrm{G}} + C^{\mathrm{D}} < C^{\mathrm{T}}$ holds then the transmission capacity is never constraining).
- α, β: Charging, discharging efficiency of the storage facility; both parameters are in $(0, 1]$. Our model can be easily extended to include inefficient holding of electricity over time without changing the structural results in §4 and §5.1.

[1] https://www.law.cornell.edu/cfr/text/17/38.152

- τ: Transmission efficiency, i.e., the ratio of electricity flowing out of the transmission line to that flowing into this line, so that $1 - \tau$ is the line loss rate; $\tau \in (0, 1]$. Losses are incurred at the end of the transmission line in either direction.

- δ: One-period risk-free discount rate (we use risk-neutral valuation; Seppi 2002, Smith 2005); this parameter is in $(0, 1]$.

State variables. Period t is defined as the time interval $[t, t+1)$. A state variable with subscript t is known at the beginning of period t, i.e., time t, but unknown in earlier periods (Powell 2007, §5.2). The state at time t, S_t, includes the following components:

- x_t: Inventory of electricity (in energy units) in the storage facility at the beginning of period t. The domain of this variable is $\mathcal{X} := [0, C^{\mathrm{S}}]$.

- w_t: Available wind energy, or electricity that can be generated given the wind speed at time t by the wind turbines in period t (in energy units/period). This quantity is limited by the generation capacity C^{G} of the turbines; $w_t \in \mathcal{W} \subseteq [0, C^{\mathrm{G}}]$. We model wind speed using a stochastic process and convert it to w_t using the production curve of turbines. We set w_t to zero if the wind speed falls below the cut-in speed or exceeds the cut-off speed (the minimal and maximal speeds at which a given turbine can generate output; in the latter case, the issue is avoiding damage to the turbines).

- $\vec{p}_t$: Price-component pair at time t, $\vec{p}_t \in \mathcal{P} \subseteq \mathbb{R}^2$, which includes a mean-reverting component and a spike component. The electricity price at time t, P_t (\$/energy unit), is a function of time t and these two components. In §6.2.2 we specify such a function. One could model the electricity price as a function of more than two components without changing the structural results in §4 and §5.1.

The state S_t is the vector $(x_t, w_t, \vec{p}_t)$. The *given* initial state S_0 is $(x_0, w_0, \vec{p}_0)$. The state space in each period is $\mathcal{S} := \mathcal{X} \times \mathcal{W} \times \mathcal{P}$.

Decision variables (actions). At time $t \in \mathcal{T}$, the merchant observes S_t and determines the inventory change and generation (action) pair $(a_t, g_t) \in \mathbb{R} \times \mathbb{R}^+$. In particular, inventory reductions and additions, respectively, correspond to negative and positive values of a_t. We use a single inventory adjustment decision variable rather than separate such variables for inventory increases and decreases because they either involve wash trades, which are illegal, or are provably suboptimal.

Transition functions. The inventory level changes from time t to time $t + 1$ according to $x_{t+1} = x_t + a_t$. The variable w_t and vector $\vec{p}_t$ evolve to w_{t+1} and $\vec{p}_{t+1}$, respectively, according to Markovian stochastic processes (see §6.2 for examples thereof), which we assume to be independent; relaxing this assumption does not change our structural results given in §4 and §5.1.

Immediate payoff function and constraints. We define the immediate payoff function as

$$R(a_t, g_t, \vec{p}_t) := \begin{cases} -P_t(\vec{p}_t)(a_t/\alpha - g_t)/\tau & \text{if } a_t > \alpha g_t, \\ P_t(\vec{p}_t)(g_t - a_t/\alpha)\tau & \text{if } 0 \leq a_t \leq \alpha g_t, \\ P_t(\vec{p}_t)(g_t - \beta a_t)\tau & \text{if } a_t < 0, \end{cases}$$

where the first case represents purchasing cost and the last two cases express selling revenue. Given (x_t, w_t), we denote by $\Psi(x_t, w_t)$ the set of action pairs (a_t, g_t) in $\mathbb{R} \times \mathbb{R}^+$ that satisfy

$$[(a_t/\alpha - g_t)/\tau - C^{\mathrm{T}}]\mathbf{1}(a_t > \alpha g_t) \leq 0, \tag{1}$$

$$(g_t - a_t/\alpha - C^{\mathrm{T}})\mathbf{1}(0 \leq a_t \leq \alpha g_t) \leq 0, \tag{2}$$

$$(g_t - \beta a_t - C^{\mathrm{T}})\mathbf{1}(a_t < 0) \leq 0, \tag{3}$$

$$g_t \leq w_t, \tag{4}$$

$$-x_t \leq a_t \leq C^{\mathrm{S}} - x_t, \tag{5}$$

$$-C^{\mathrm{D}} \leq a_t \leq C^{\mathrm{C}}; \tag{6}$$

here (1), (2), and (3) are the transmission capacity constraints—$\mathbf{1}(\cdot)$ is the indicator function that equals one if its argument is true and zero otherwise; (4) restricts the generation to be no more than the wind energy availability (this constraint and $w_t \leq C^{\mathrm{G}}$ imply $g_t \leq C^{\mathrm{G}}$); (5) restrains the inventory change to lie between the (negative of the) energy available in the storage facility and the remaining storage *energy* capacity; and (6) imposes on this change the limits expressed by the (negative of the) discharging capacity and the charging capacity.

Objective function. Each stage of our MDP corresponds to a time in $\mathcal{T}$. A feasible policy π is the sequence of decision rules $(A_t^\pi(S_t))_{t \in \mathcal{T}}$, where $A_t^\pi(S_t)$ maps the state S_t to the feasible action pair $(a_t^\pi(S_t), g_t^\pi(S_t))$ in stage t. Our objective is to maximize the total discounted expected cash flows over all feasible policies, which we include in the set Π:

$$\max_{\pi \in \Pi} \sum_{t \in \mathcal{T}} \delta^t \mathbb{E}\left[R(A_t^\pi(S_t^\pi), \vec{p}_t)|S_0\right], \tag{7}$$

where the expectation $\mathbb{E}$ is taken with respect to the distribution of the random state S_t^π reached by policy π in stage t. For each stage $t \in \mathcal{T}$ and state $S_t \in \mathcal{S}$, the value function of each feasible policy π, $V_t^\pi(S_t)$, satisfies the recursion

$$V_t^\pi(S_t) = R(a_t^\pi(S_t), g_t^\pi(S_t), \vec{p}_t) + \delta \mathbb{E}\left[V_{t+1}^\pi(x_t + a_t^\pi(S_t), w_{t+1}, \vec{p}_{t+1})|S_t\right] \tag{8}$$

and the optimal value function, $V_t^*(S_t)$, solves

$$V_t^*(S_t) = \max_{(a_t, g_t) \in \Psi(x_t, w_t)} \left\{R(a_t, g_t, \vec{p}_t) + \delta \mathbb{E}\left[V_{t+1}^*(x_t + a_t, w_{t+1}, \vec{p}_{t+1})|S_t\right]\right\} x, \tag{9}$$

with both $V_T^\pi(S_T)$ and $V_t^*(S_t)$ set equal to zero for each $S_T \in \mathcal{S}$.

Appendix A.1 discusses the computation of an optimal policy in the context of our numerical study, in which the dynamics of the available wind energy and the price-component pair are specified in §6.2 and the state components and the action pairs are discretized as discussed in §6.3.

4. Analysis when the electricity prices are always positive

In this section we analyze model (7). Specifically, we establish an optimal policy structure when the electricity prices can only be positive, which forms the basis of H1 and H2, our heuristics for the case when prices can also be negative.

To obtain a well defined model we make the benign Assumption 1.

ASSUMPTION 1. For each $t \in \mathcal{T}$, we have $\mathbb{E}[|P_{t'}(\vec{p}_{t'})||\vec{p}_t] < \infty$ for each $t' \in \mathcal{T}$, $t' \geq t$, and $\vec{p}_t \in \mathcal{P}$.

We state in Proposition 1 the concavity of the resulting optimal value function in the inventory level given all the other state components in each stage.

PROPOSITION 1. Suppose $P_t(\vec{p}_t) > 0$ for each $t \in \mathcal{T}$ and $\vec{p}_t \in \mathcal{P}$. For each $t \in \mathcal{T} \cup \{T\}$, $V_t^*(x_t, w_t, \vec{p}_t)$ is concave in x_t for each given $(w_t, \vec{p}_t) \in \mathcal{W} \times \mathcal{P}$.

Based on Proposition 1, we obtain the policy structure presented in Proposition 2. We denote by $(a_t^*(S_t), g_t^*(S_t))$ an optimal action pair for the optimization on the right hand side of (9). For notational convenience, we define the ending inventory level after modifying x_t by a_t as $y_t := x_t + a_t$; the optimal continuation function as $U_t^*(\cdot, w_t, \vec{p}_t) := \delta \mathbb{E}\left[V_{t+1}^*(\cdot, w_{t+1}, \vec{p}_{t+1})|w_t, \vec{p}_t\right]$; the functions

$$X_t^{(1)}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{U_t^*(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t)y_t/(\alpha\tau)\right\}, \tag{10}$$

$$X_t^{(2)}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{U_t^*(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t)\tau y_t/\alpha\right\}, \tag{11}$$

$$X_t^{(3)}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{U_t^*(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t)\beta\tau y_t\right\}, \tag{12}$$

where we implicitly employ the largest element of each defining set; and the sets

$$\Gamma^0 := \left\{(x_t, w_t) \in \mathcal{X} \times \mathcal{W} : w_t \geq C^{\mathrm{T}} + \min\left\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\right\}/\alpha\right\},$$

$$\Gamma^1 := \left\{(x_t, w_t) \in \mathcal{X} \times \mathcal{W} : C^{\mathrm{T}} \leq w_t < C^{\mathrm{T}} + \min\left\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\right\}/\alpha\right\},$$

$$\Gamma^2 := \left\{(x_t, w_t) \in \mathcal{X} \times \mathcal{W} : 0 \leq w_t < C^{\mathrm{T}}\right\}.$$

For brevity, below we write $X_t^{(\nu)}$ in lieu of $X_t^{(\nu)}(w_t, \vec{p}_t)$ for each $\nu \in \{1, 2, 3\}$.

PROPOSITION 2. *Suppose $P_t(\vec{p}_t) > 0$ for each $t \in \mathcal{T}$. For each $t \in \mathcal{T}$ we have (i) $X_t^{(1)} \leq X_t^{(2)} \leq X_t^{(3)}$, (ii) $g_t^*(S_t) = \min\{w_t, C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$, and (iii) $a_t^*(S_t) = \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}$ (store excess generation) if $(x_t, w_t) \in \Gamma^0$;*

$$
a_t^*(S_t) = \begin{cases}
\min\left\{X_t^{(1)} - x_t, \alpha\left(\tau C^{\mathrm{T}} + w_t\right), C^{\mathrm{C}}\right\} & \\
\textit{(store generated and purchased electricity up to } X_t^{(1)}) & \textit{if } x_t \in \left[0, X_t^{(1)} - \alpha w_t\right], \\
\min\left\{X_t^{(2)} - x_t, \alpha w_t\right\} & \\
\textit{(store generation up to } X_t^{(2)}) & \textit{if } x_t \in \left(X_t^{(1)} - \alpha w_t, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right)\right], \\
\alpha\left(w_t - C^{\mathrm{T}}\right) & \\
\textit{(store excess generation)} & \textit{if } x_t \in \left(X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right), C^{\mathrm{S}}\right],
\end{cases}
$$

$(x_t, w_t) \in \Gamma^1$, and $\alpha w_t < \min\left\{X_t^{(1)}, C^{\mathrm{C}}\right\}$;

$$
a_t^*(S_t) = \begin{cases}
\min\left\{X_t^{(2)} - x_t, \alpha w_t\right\} & \\
\textit{(store generation up to } X_t^{(2)}) & \textit{if } x_t \in \left[0, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right)\right], \\
\alpha\left(w_t - C^{\mathrm{T}}\right) & \\
\textit{(store excess generation)} & \textit{if } x_t \in \left(X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right), C^{\mathrm{S}}\right],
\end{cases}
$$

$(x_t, w_t) \in \Gamma^1$, $\alpha\left(w_t - C^{\mathrm{T}}\right) < X_t^{(2)}$, and either $X_t^{(1)} \leq \alpha w_t < C^{\mathrm{C}}$ or $\alpha w_t \geq C^{\mathrm{C}}$; $a_t^*(S_t) = \alpha\left(w_t - C^{\mathrm{T}}\right)$ (store excess generation) if $(x_t, w_t) \in \Gamma^1$, $\alpha\left(w_t - C^{\mathrm{T}}\right) \geq X_t^{(2)}$, and either $X_t^{(1)} \leq \alpha w_t < C^{\mathrm{C}}$ or $\alpha w_t \geq C^{\mathrm{C}}$;

$$
a_t^*(S_t) = \begin{cases}
\min\left\{X_t^{(1)} - x_t, \alpha\left(\tau C^{\mathrm{T}} + w_t\right), C^{\mathrm{C}}\right\} & \\
\textit{(store generated and purchased electricity up to } X_t^{(1)}) & \textit{if } x_t \in \left[0, X_t^{(1)} - \alpha w_t\right], \\
\min\left\{X_t^{(2)} - x_t, \alpha w_t\right\} & \\
\textit{(store generation up to } X_t^{(2)}) & \textit{if } x_t \in \left(X_t^{(1)} - \alpha w_t, X_t^{(2)}\right], \\
0 \textit{ (keep inventory unchanged)} & \textit{if } x_t \in \left(X_t^{(2)}, X_t^{(3)}\right], \\
\max\left\{X_t^{(3)} - x_t, \left(w_t - C^{\mathrm{T}}\right)/\beta, -C^{\mathrm{D}}\right\} & \\
\textit{(sell inventory down to } X_t^{(3)}) & \textit{if } x_t \in \left(X_t^{(3)}, C^{\mathrm{S}}\right],
\end{cases}
$$

$(x_t, w_t) \in \Gamma^2$, and $\alpha w_t < \min\left\{X_t^{(1)}, C^{\mathrm{C}}\right\}$; and

$$
a_t^*(S_t) = \begin{cases}
\min\left\{X_t^{(2)} - x_t, \alpha w_t\right\} & \\
\textit{(store generation up to } X_t^{(2)}) & \textit{if } x_t \in \left[0, X_t^{(2)}\right], \\
0 \textit{ (keep inventory unchanged)} & \textit{if } x_t \in \left(X_t^{(2)}, X_t^{(3)}\right], \\
\max\left\{X_t^{(3)} - x_t, \left(w_t - C^{\mathrm{T}}\right)/\beta, -C^{\mathrm{D}}\right\} & \\
\textit{(sell inventory down to } X_t^{(3)}) & \textit{if } x_t \in \left(X_t^{(3)}, C^{\mathrm{S}}\right],
\end{cases}
$$

$(x_t, w_t) \in \Gamma^2$, and either $X_t^{(1)} \leq \alpha w_t < C^{\mathrm{C}}$ or $\alpha w_t \geq C^{\mathrm{C}}$.

Figure 2 **Illustration of the optimal end-of-period inventory,** $x_t + a_t^*(S_t)$**, as a function of the beginning-of-period inventory,** x_t**, for the case** $(x_t, w_t) \in \Gamma^2$ **and** $\alpha w_t < \min\left\{X_t^{(1)}, C^C\right\}$ **in Proposition 2 (for ease of illustration we assume that the thresholds can be reached from every inventory level in each of their respective intervals)**
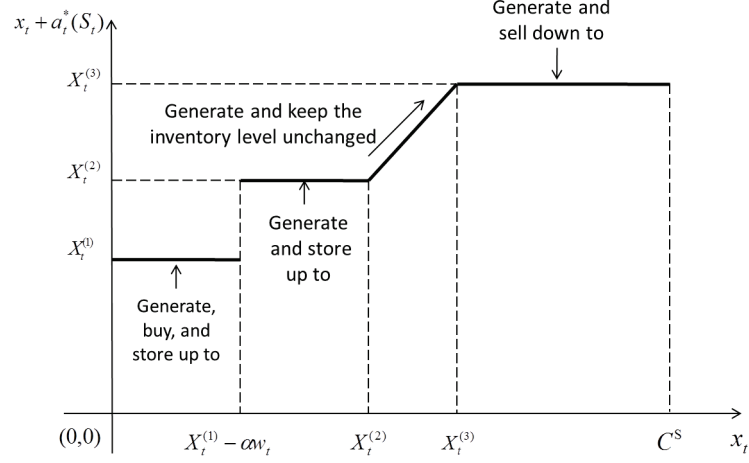
Generating to the maximal extent is optimal. When w_t is large, i.e., in Γ^0, it is optimal to sell as much as possible and store any leftover so that the ending inventory level is as close as possible to C^S. Otherwise, i.e., w_t is in Γ^1 or Γ^2, the optimal inventory change depends on the quantities $X_t^{(1)}$, $X_t^{(2)}$, and $X_t^{(3)}$. In particular, when $(x_t, w_t) \in \Gamma^2$ and $\alpha w_t < \min\left\{X_t^{(1)}, C^C\right\}$ this action can be of four distinctive types: If $x_t \leq X_t^{(1)} - \alpha w_t$, store generation and purchased electricity to bring the inventory level as close as possible to $X_t^{(1)}$; if $X_t^{(1)} - \alpha w_t < x_t \leq X_t^{(2)}$, store generated electricity without buying so that the resulting inventory level is as close as possible to $X_t^{(2)}$; if $X_t^{(2)} < x_t \leq X_t^{(3)}$, generate and keep the inventory level unchanged; if $x_t > X_t^{(3)}$, generate and sell to bring the inventory level as near as achievable to $X_t^{(3)}$. Thus, we refer to $X_t^{(1)}$, $X_t^{(2)}$, and $X_t^{(3)}$ as inventory threshold functions. Figure 2 illustrates the behavior of the inventory level at the end of a period as a function of the one at the beginning of a period corresponding to the action types just described.

The functions $X_t^{(1)}$, $X_t^{(2)}$, and $X_t^{(3)}$ in general return separate values due to the presence of both the transmission loss ($\tau < 1$) and the charging/discharging losses ($\alpha\beta < 1$), which make the immediate marginal values of the following three types of actions, or, equivalently, the respective slopes of the immediate payoff function, different from each other: Storing one unit of bought electricity, storing one unit of generation, and selling one unit of inventory. If $\tau = 1$ then $X_t^{(1)} = X_t^{(2)}$, because the marginal values of increasing the inventory using one unit of purchased or generated electricity are equal (in this case the optimal structure reduces to that in Secomandi 2010). Similarly, if $\alpha\beta = 1$ then $X_t^{(2)}$ coincides with $X_t^{(3)}$.

5. Heuristics

In this section we define four heuristics for model (7). We introduce H1 in §5.1, H2 in §5.2, and heuristics three and three-plus (H3 and H3+), which are inspired by the literature and rely on policies that use only price information, in §5.3. The H2, H3, and H3+ policies are stationary, which is a potential limitation in a non-stationary environment, such as our numerical study. We also considered heuristics with policies based on time-of-day (used, e.g., for managing hydro reservoirs; see Powell and Meisel 2016a), deterministic re-optimization, and two- and three-period lookahead (Powell and Meisel 2016b). They are not included in this paper because they perform unsatisfactorily on our instances.

5.1. Heuristic one

The idea for H1 is as follows: (i) We extend the structure of our optimal policy for the special case when prices can only be positive to obtain a feasible policy for the general case when prices can also be negative and (ii) we approximately solve a stochastic dynamic program to compute it. It can be shown that this extended structure would be optimal in the latter case if the objective function of the optimization in (9) were jointly concave in the initial inventory level and the action pair. Unfortunately this joint concavity is not true in general: Online Appendix D includes an example for which the H1 policy is not optimal.

H1 partly relies on the threshold functions

$$X_t^{(1),\text{H1}}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{ U_t^{\text{H1}}(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t) y_t / (\alpha\tau) \right\}, \tag{13}$$

$$X_t^{(2),\text{H1}}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{ U_t^{\text{H1}}(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t) \tau y_t / \alpha \right\}, \tag{14}$$

$$X_t^{(3),\text{H1}}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{ U_t^{\text{H1}}(y_t, w_t, \vec{p}_t) - P_t(\vec{p}_t) \beta\tau y_t \right\}, \tag{15}$$

$$X_t^{(4),\text{H1}}(w_t, \vec{p}_t) := \underset{y_t \in \mathcal{X}}{\arg\max} \left\{ U_t^{\text{H1}}(y_t, w_t, \vec{p}_t) \right\}, \tag{16}$$

where the largest element of each defining set is implicitly used and $U_t^{\text{H1}}(\cdot, w_t, \vec{p}_t) := \delta\mathbb{E}\left[V_{t+1}^{\text{H1}}(\cdot, w_{t+1}, \vec{p}_{t+1}) | w_t, \vec{p}_t\right]$; $V_{t+1}^{\text{H1}}(S_{t+1})$ satisfies (8) with π set to H1 (we describe the actions of the H1 policy below). Proposition 3 orders these functions. We abbreviate $X_t^{(\nu),\text{H1}}(w_t, \vec{p}_t)$ to $X_t^{(\nu),\text{H1}}$ for each $\nu \in \{1, 2, 3, 4\}$.

PROPOSITION 3. *In each $t \in \mathcal{T}$ and $(w_t, \vec{p}_t) \in \mathcal{W} \times \mathcal{P}$ we have $X_t^{(1),\text{H1}} \leq X_t^{(2),\text{H1}} \leq X_t^{(3),\text{H1}} \leq X_t^{(4),\text{H1}}$ when $P_t(\vec{p}_t) > 0$ and $X_t^{(1),\text{H1}} \geq X_t^{(2),\text{H1}} \geq X_t^{(3),\text{H1}} \geq X_t^{(4),\text{H1}}$ when $P_t(\vec{p}_t) \leq 0$.*

Proposition 3 states that in each stage and partial state (i.e., excluding the inventory level) the ordering of $X_t^{(1),\text{H1}}$ through $X_t^{(4),\text{H1}}$ depends only on the sign of the price $P_t(\vec{p}_t)$. Intuitively, this

claim holds because the objective functions in (13)-(16) differ only by a term that is linear in $P_t(\vec{p}_t)$. When $P_t(\vec{p}_t) > 0$ Proposition 3 is consistent with part (i) of Proposition 2.

The H1 policy follows:

(i) $P_t(\vec{p}_t) > 0$: If $x_t \in \left[0, X_t^{(4),\mathrm{H1}}\right]$ then use the decision rules specified in Proposition 2 replacing $X_t^{(\nu)}$ with $X_t^{(\nu),\mathrm{H1}}$ for $\nu \in \{1, 2, 3\}$, C^{S} with $X_t^{(4),\mathrm{H1}}$, and Γ^0, Γ^1, and Γ^2, respectively, with

$$\Gamma_t^{0,\mathrm{H1}}(\vec{p}_t) := \left\{ w_t \in \mathcal{W}, x_t \in \left[0, X_t^{(4)}(w_t, \vec{p}_t)\right] : w_t \geq C^{\mathrm{T}} + \min\{X_t^{(4)}(w_t, \vec{p}_t) - x_t, C^{\mathrm{C}}\}/\alpha \right\},$$
$$\Gamma_t^{1,\mathrm{H1}}(\vec{p}_t) := \left\{ w_t \in \mathcal{W}, x_t \in \left[0, X_t^{(4)}(w_t, \vec{p}_t)\right] : C^{\mathrm{T}} \leq w_t < C^{\mathrm{T}} + \min\{X_t^{(4)}(w_t, \vec{p}_t) - x_t, C^{\mathrm{C}}\}/\alpha \right\},$$
$$\Gamma_t^{2,\mathrm{H1}}(\vec{p}_t) := \left\{ w_t \in \mathcal{W}, x_t \in \left[0, X_t^{(4)}(w_t, \vec{p}_t)\right] : 0 \leq w_t < C^{\mathrm{T}} \right\}.$$

If $x_t \in \left(X_t^{(4),\mathrm{H1}}, C^{\mathrm{S}}\right]$ then we define $a_t^{\mathrm{H1}}(S_t) := \max\left\{X_t^{(3),\mathrm{H1}} - x_t, -C^{\mathrm{T}}/\beta, -C^{\mathrm{D}}\right\}$ and $g_t^{\mathrm{H1}}(S_t) := \min\{w_t, C^{\mathrm{T}} - \beta a_t^{\mathrm{H1}}(S_t)\}$, i.e., inventory is sold to reach a level that is as close as possible to $X_t^{(3),\mathrm{H1}}$ and as much wind energy as feasible is produced and sold taking into account the residual transmission capacity.

(ii) $P_t(\vec{p}_t) \leq 0$: The threshold functions $X_t^{(\nu),\mathrm{H1}}$ for $\nu \in \{1, 2, 3, 4\}$ and $Z_t(w_t, \vec{p}_t)$, which we define below and abbreviate to Z_t for notational convenience, determine the action pair as follows:

- If $\alpha\tau C^{\mathrm{T}} < X_t^{(4),\mathrm{H1}}$ then

$$a_t^{\mathrm{H1}}(S_t) := \begin{cases} \min\left\{X_t^{(4),\mathrm{H1}} - x_t, \alpha\left(\tau C^{\mathrm{T}} + w_t\right), C^{\mathrm{C}}\right\} & \\ \text{(store purchased electricity and} & \\ \quad \text{generated electricity up to } X_t^{(4),\mathrm{H1}}) & \text{if } x_t \in \left[0, X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}\right], \\ \min\left\{X_t^{(1),\mathrm{H1}} - x_t, \alpha\tau C^{\mathrm{T}}, C^{\mathrm{C}}\right\} & \\ \text{(store purchased electricity up to } X_t^{(1),\mathrm{H1}}) & \text{if } x_t \in \left(X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}, Z_t\right], \\ \max\left\{X_t^{(3),\mathrm{H1}} - x_t, -C^{\mathrm{T}}/\beta, -C^{\mathrm{D}}\right\} & \\ \text{(sell inventory down to } X_t^{(3),\mathrm{H1}}) & \text{if } x_t \in (Z_t, C^{\mathrm{S}}], \end{cases}$$

$$g_t^{\mathrm{H1}}(S_t) := \begin{cases} (a_t^{\mathrm{H1}}(S_t)/\alpha - \tau C^{\mathrm{T}})^+ \text{ (store generation)} & \text{if } x_t \in \left[0, X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}\right], \\ 0 \text{ (do not generate)} & \text{if } x_t \in (Z_t, C^{\mathrm{S}}]. \end{cases}$$

- If $\alpha\tau C^{\mathrm{T}} \geq X_t^{(4),\mathrm{H1}}$ then

$$a_t^{\mathrm{H1}}(S_t) := \begin{cases} \min\left\{X_t^{(1),\mathrm{H1}} - x_t, \alpha\tau C^{\mathrm{T}}, C^{\mathrm{C}}\right\} & \\ \text{(store purchased electricity up to } X_t^{(1),\mathrm{H1}}) & \text{if } x_t \in [0, Z_t], \\ \max\left\{X_t^{(3),\mathrm{H1}} - x_t, -C^{\mathrm{T}}/\beta, -C^{\mathrm{D}}\right\} & \\ \text{(sell inventory down to } X_t^{(3),\mathrm{H1}}) & \text{if } x_t \in (Z_t, C^{\mathrm{S}}], \end{cases}$$

$$g_t^{\mathrm{H1}}(S_t) := 0 \text{ if } x_t \in \left[0, C^{\mathrm{S}}\right] \text{ (do not generate)}.$$

The function Z_t is the maximal inventory level $x_t \in \left[X_t^{(3),\mathrm{H1}}, X_t^{(1),\mathrm{H1}}\right]$ such that the two action pairs $\left(\min\left\{X_t^{(1),\mathrm{H1}} - x_t, \alpha\tau C^{\mathrm{T}}, C^{\mathrm{C}}\right\}, 0\right)$ and $\left(\max\left\{X_t^{(3),\mathrm{H1}} - x_t, -C^{\mathrm{T}}/\beta, -C^{\mathrm{D}}\right\}, 0\right)$ result in the same

Figure 3 **Illustration of the H1 policy end-of-period inventory,** $x_t + a_t^{\mathrm{H1}}(S_t)$, **as a function of the beginning-of-period inventory,** x_t, **when** $P_t(\vec{p}_t) \leq 0$ **and** $\alpha\tau C^{\mathrm{T}} < X_t^{(4),\mathrm{H1}}$ **(for illustration purposes we make the assumption that the thresholds are attainable from every inventory level in each of their corresponding intervals)**
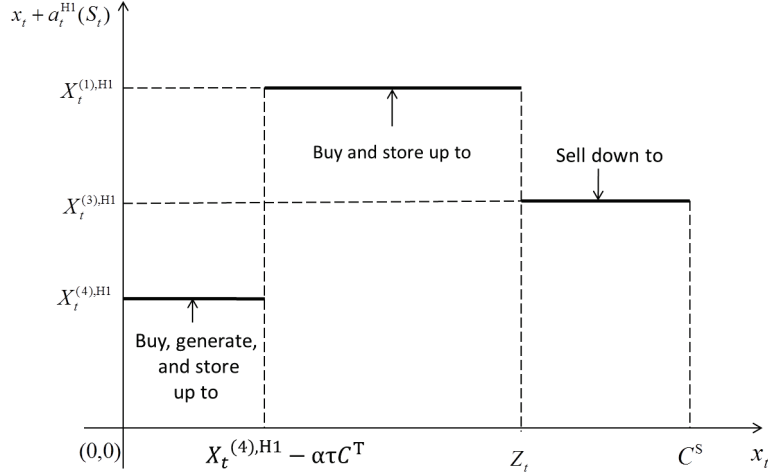
evaluation of the right hand side of (8) with π set equal to H1. Intuitively, this level is such that we are indifferent between storing purchased electricity up to $X_t^{(1),\mathrm{H1}}$ and selling inventory down to $X_t^{(3),\mathrm{H1}}$ when no wind energy is produced.

When $P_t(\vec{p}_t) \leq 0$ and $\alpha\tau C^{\mathrm{T}} < X_t^{(4),\mathrm{H1}}$, the action pair under the H1 policy can be of three separate types (see Figure 3): When $x_t \leq X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}$ buy and generate electricity so that the ending inventory level is as close as possible to $X_t^{(4),\mathrm{H1}}$; when $X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}} < x_t \leq Z_t$ purchase energy to bring the inventory level as near as viable to $X_t^{(1),\mathrm{H1}}$ without producing; and when $x_t > Z_t$ sell inventory so that the resulting level approaches $X_t^{(3),\mathrm{H1}}$ as much as is feasible and do not generate. It may be desirable for the H1 policy to sell inventory when the price in the current period is negative (e.g., if $x_t > Z_t$ in Figure 3) because a larger end-of-period inventory level is not always more appealing than a smaller one; i.e., the continuation function of this policy may fail to be monotonic in this type of inventory level. For instance, when the expected price in the next period is even more negative than the price in the current period, selling now at a loss to free space to buy and store then is potentially appealing.

Under the H1 policy, the inventory level at the end of a period can fail to be monotonic in the one at the beginning of a period. For example, in Figure 3 as the latter level increases from zero to C^S the former one first increases and then decreases. This rise occurs because (i) if the inventory level is smaller than $X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}$ then after purchasing to the maximal extent, i.e., C^{T}, it is advantageous to increase it up to $X_t^{(4),\mathrm{H1}}$ by generating provided residual charging capacity is available and (ii) if this level lies between $X_t^{(4),\mathrm{H1}} - \alpha\tau C^{\mathrm{T}}$ and Z_t then it is beneficial to purchase

electricity to bring it as close as possible to $X_t^{(1),\text{H1}}$. This drop arises for the following reasons: (i) If the inventory level is between $X_t^{(4),\text{H1}} - \alpha\tau C^\text{T}$ and Z_t then there is sufficient space to purchase and store electricity up to $X_t^{(1),\text{H1}}$ to take advantage of the negative price in the current period and (ii) if this level exceeds Z_t then the space available for storing purchased electricity is so limited that it is advantageous to sell inventory down to $X_t^{(3),\text{H1}}$ to create the opportunity to purchase electricity in the event that prices will be negative in the future.

In the discretized state and action setting of §6, we obtain the H1 policy by approximately, rather than exactly, solving a stochastic dynamic program as discussed in Appendix A.2, exploiting the fact that its threshold functions do not depend on the inventory level. Nonetheless, H1 may not be practical for real time use, because the value function of its policy needs to be evaluated for every stage and state.

5.2. Heuristic two

H2 simplifies H1 both in terms of type of policy used and the approach taken to compute the values of its policy parameters.

The H2 policy is a version of the H1 policy that employs price dependent inventory threshold functions, $X^{(\nu),\text{H2}}(\vec{p}_t)$ for $\nu \in \{1,2,3,4\}$, which we simplify by suppressing their argument. Their definitions, which are based on the scalars P^S and $X^{(1)}$ through $X^{(4)}$ with $0 \le X^{(1)} \le X^{(2)} \le X^{(3)} \le X^{(4)} \le C^\text{S}$, follow:

- If $P_t(\vec{p}_t) \le 0$ then $X^{(1),\text{H2}} = X^{(2),\text{H2}} = X^{(3),\text{H2}} = X^{(4),\text{H2}} := C^\text{S}$ (purchase as much electricity as possible and generate as much as viable);
- If $0 < P_t(\vec{p}_t) < P^\text{S}$ then $X^{(1),\text{H2}} := X^{(1)}, X^{(2),\text{H2}} := X^{(2)}, X^{(3),\text{H2}} := X^{(3)}$, and $X^{(4),\text{H2}} := X^{(4)}$ (analogous to the case when $P_t(\vec{p}_t) > 0$ for the H1 policy);
- If $P_t(\vec{p}_t) \ge P^\text{S}$ then $X^{(1),\text{H2}} = X^{(2),\text{H2}} = X^{(3),\text{H2}} := 0$ and $X^{(4),\text{H2}} := C^\text{S}$ (generate and sell to the maximal extent and then bring to market as much inventory as feasible).

The H2 policy has five parameters. We compute them using a derivative-free nonlinear optimization algorithm embedded within a wind-speed and electricity-price Monte Carlo simulation, which we describe in Appendix A.3.

5.3. Heuristics three and three-plus

H3 and its enhanced version H3+ are two heuristics whose policies use only the current price, as is common in the electricity storage literature (e.g., Graves et al. 1999, Fertig and Apt 2011, Powell and Meisel 2016a, and Powell and Meisel 2016b). The H3 policy relies on two scalars, $P^\text{D} > P^\text{C} > 0$, where the superscripted D and C stand for discharging and charging, respectively. It increases the

inventory by first generating as much as possible and then by purchasing to the maximal extent if the price is low, $P_t(\vec{p}_t) \leq P^{\mathrm{C}}$; keeps the inventory level unchanged and produces and sells as much as possible if the price is moderate, $P^{\mathrm{C}} < P_t(\vec{p}_t) < P^{\mathrm{D}}$; and generates and brings to market as much as feasible, storing any excess, if the price is high, $P_t(\vec{p}_t) \geq P^{\mathrm{D}}$. Specifically, the action pairs of the H3 policy follow:

- If $P_t(\vec{p}_t) \leq P^{\mathrm{C}}$ then $a_t^{\mathrm{H3}}(S_t) := \min\{\alpha(w_t + \tau C^{\mathrm{T}}), C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}$ and $g_t^{\mathrm{H3}}(S_t) := \min\{w_t, \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$;

- If $P^{\mathrm{C}} < P_t(\vec{p}_t) < P^{\mathrm{D}}$ then $a_t^{\mathrm{H3}}(S_t) := 0$ and $g_t^{\mathrm{H3}}(S_t) := \min\{w_t, C^{\mathrm{T}}\}$;

- If $P_t(\vec{p}_t) \geq P^{\mathrm{D}}$ then

$$a_t^{\mathrm{H3}}(S_t) := \begin{cases} -\max\{x_t, (C^{\mathrm{T}} - w_t)/\beta, -C^{\mathrm{D}}\} & \text{if } w_t \leq C^{\mathrm{T}}, \\ \alpha[g_t^{\mathrm{H3}}(S_t) - C^{\mathrm{T}}] & \text{otherwise}, \end{cases}$$

and $g_t^{\mathrm{H3}}(S_t) := \min\{w_t, C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$.

The H3+ policy differs from the H3 policy only when prices are negative, in which case the former policy purchases as much electricity as possible and then maximizes the amount generated. H3 and H3+ compute their respective policies based on an approach similar to that used by H2 (see Appendix A.3).

6. Numerical study

In this section we discuss our numerical study. We present the setup of this study in §6.1. We discuss the wind and electricity price models and their calibrations in §6.2. We explain the approach that supports the computation of the various policies in §6.3. In §6.4 we assess the performance of the heuristics presented in §5. We quantify the effect of changing the frequency of negative prices on the incremental value brought about by and environmental benefit of storage in §6.5.

6.1. Setup

Electricity is traded on the New York City real time market. This market is managed by the New York Independent System Operator (NYISO), which includes fifteen local markets (zones), and is one of the largest and most liquid electricity markets (NYISO 2013). As real time prices are set every five minutes, we specify our model using this frequency, i.e., each period corresponds to a five minute interval. The per stage discount factor δ is 0.9999999, corresponding to an annual risk-free interest rate of 1% with continuous compounding (recall that we use risk-neutral valuation). We fix the horizon to be one week, i.e., the number of five-minute periods is $12 \times 24 \times 7 = 2{,}016$.

We consider a hypothetical wind farm located in Buffalo, NY, which houses one of the largest wind farms in New York State. This wind farm consists of 120 General Electric (GE) model 1.5-77

turbines. This type of turbine is among the best selling ones in the US (Wiser and Bolinger 2013). Its capacity is 1.5 MW (1 MW = 1 million watts and 1 watt = 1 joule/second), so the generation capacity of the wind farm is 120×1.5 MW = 180 MW. We scale it given the five-minute period length to obtain C^{G}, i.e., $C^{\mathrm{G}} = 180$ MW/(12 periods/hour) = 15 MWh/period.

Co-located with our wind farm is an industrial battery. We vary its energy capacity (C^{S}, in megawatt hour, or MWh; 1 MWh = 3.6 Giga joules) from 200 MWh to 1200 MWh in steps of 200 MWh. This battery can be fully charged or discharged in ten hours (EPRI 2004). Thus, its (charging/discharging) power capacity is its energy capacity divided by ten hours, which we scale taking into account the period length to obtain C^{C} and C^{D} (in MWh/period). The base values of the battery charging/discharging efficiencies are $\alpha = 0.85$ and $\beta = 1$. We varied these values over a broad range but found that our insights remained qualitatively unchanged. Therefore, we report results only for the base values of these parameters.

A transmission line connects our wind farm to New York City. This line has a loss of 3%, i.e., $\tau = 97\%$ (Duke Energy 2017). We consider values for the portion of the capacity of this line leased to the wind farm between 80 MW and 180 MW in steps of 20 MW. The parameter C^{T} is this rented transmission capacity scaled according to the given period length.

The relative values of the considered generation, storage, and transmission capacities are consistent with those in Denholm and Sioshansi (2009) and Pattanariyankool and Lave (2010).

6.2. Wind speed and electricity price models and their calibrations

We describe the wind speed and electricity price models and their calibrations in §6.2.1 and §6.2.2, respectively.

6.2.1. Wind speed model and its calibration.
To represent wind speed evolution, we use an autoregressive of order one, AR(1), process with deterministic seasonality (Kim and Powell 2011, Wu and Kapuscinski 2013). We convert wind speed to available wind energy using the production curve of the GE 1.5-77 wind turbine (displayed in Table 1; General Electric 2013), which has cut-in and cut-off speeds of 4 and 25 meters per second (m/s), respectively.

Table 1   Production curve of the GE 1.5-77 turbine (with cut-in and cut-off speeds of 4 and 25 m/s)

Speed (m/s)	4	5	6	7	8	9	10	11	12	13	14	15-25
Power (MW)	0.043	0.131	0.25	0.416	0.64	0.924	1.181	1.359	1.436	1.481	1.494	1.5

We have available from NOAA (2010) hourly wind speed data from 2005 to 2008 for Buffalo, NY. This wind speed data was recorded at 10 meters above ground (Buffalo), whereas the height

Table 2   **Estimated parameters of the wind speed model (MAE $= 0.145$ MW in the hourly case)**

Seasonality function					Hourly AR(1) process		Five-minute AR(1) process	
$\hat{\gamma}_0$	$\hat{\gamma}_1$	$\hat{\gamma}_2$	$\hat{\omega}_1$	$\hat{\omega}_2$	$\hat{\phi}^{\mathrm{h}}$	$\hat{\sigma}^{\mathrm{h}}$	$\hat{\phi}$	$\hat{\sigma}$
6.777	1.2558	1.4743	-27.8434	28.7040	0.8813	1.6734	0.9895	0.5112

of the GE 1.5-77 turbine is 80 meters. Based on Heier (2006), we convert this data to the height of this turbine by multiplying the observed wind speeds by $(80/10)^{0.2} = 8^{0.2}$, where the choice of 0.2 yields a value of 33% for the ratio of the hypothetical sample output during 2005-2008 of our wind farm and its generation capacity, which lies within the range of reported values of analogous ratios for wind farms in NY (DOE 2017; such metrics are known as empirical capacity factors in the wind energy engineering literature).

Let $\bar{t}$ be the hourly time index, which ranges from 1 through $24 \times 7 = 168$ with unitary increments. The wind speed in hour $\bar{t}$ is the sum $\xi_{\bar{t}}^{\mathrm{h}} + f^{\mathrm{h}}(\bar{t})$, where $\xi_{\bar{t}}^{\mathrm{h}}$ evolves as the AR(1) process $\xi_{\bar{t}}^{\mathrm{h}} = \phi^{\mathrm{h}}\xi_{\bar{t}-1}^{\mathrm{h}} + \sigma^{\mathrm{h}}\epsilon_{\bar{t}}^{\mathrm{h}}$, with ϕ^{h} and σ^{h} scalars and $\epsilon_{\bar{t}}^{\mathrm{h}}$ an independent and identically distributed (i.i.d.) standard normal error term, and $f^{\mathrm{h}}(\bar{t})$ is the seasonality function $\gamma_0 + \gamma_1 \cos((\lceil\bar{t}/24\rceil + \omega_1)2\pi/365)) + \gamma_2 \cos((\bar{t} + \omega_2)2\pi/24)$, with $\lceil \cdot \rceil$ the ceiling function, γ_0 a constant, and γ_1-and-ω_1 and γ_2-and-ω_2, respectively, the daily and hourly magnitude-and-phase-shift parameters.

Modifying $f^{\mathrm{h}}(\bar{t})$ to obtain a seasonality function when the period length is five minutes entails only replacing the hourly index $\bar{t}$ with $\lceil(t+1)/12\rceil$, with t the five-minute period index, which varies from 0 through 2015. We denote the five-minute wind speed AR(1) process as $\xi_t = \phi\xi_{t-1} + \sigma\epsilon_t$, where ϕ, σ, and ϵ_t are analogous to ϕ^{h}, σ^{h}, and $\epsilon_{\bar{t}}^{\mathrm{h}}$, respectively. Applying this expression recursively yields $\xi_t = (\phi)^{12}\xi_{t-12} + \sigma[(\phi)^{11}\epsilon_{t-11} + (\phi)^{10}\epsilon_{t-10} + \cdots + \phi\epsilon_{t-1} + \epsilon_t]$. Matching the mean and standard deviation of the right-hand side of this expression with those of $\phi^{\mathrm{h}}\xi_{\bar{t}-1}^{\mathrm{h}} + \sigma^{\mathrm{h}}\epsilon_{\bar{t}}^{\mathrm{h}}$ gives $\phi = (\phi^{\mathrm{h}})^{1/12}$ and $\sigma = \sigma^{\mathrm{h}}\sqrt{[1 - (\phi^{\mathrm{h}})^{1/6}]/[1 - (\phi^{\mathrm{h}})^2]}$.

We calibrate the hourly model using nonlinear regression. We convert the resulting values of the parameters of its AR(1) component to obtain estimates for those of the corresponding five-minute version. Table 2 reports all these values. We measure the fit of this calibrated hourly model by computing the mean absolute error (MAE) in terms of wind-based electricity production of a single GE 1.5-77 turbine. The MAE is 0.145 MW. We also experimented with an AR(2) process and found that it did not fit the data any better than the AR(1) process.

6.2.2.   **Electricity price model and its calibration.**   Electricity prices exhibit mean reversion (a tendency to revert to a given level), spikes (jumps that last for a short duration), seasonality, and can be negative. Our price model combines a mean-reverting process from Lucia and Schwartz

(2002), a spike process as in Seifert and Uhrig-Homburg (2007), a deterministic seasonality function similar to one used by Lucia and Schwartz (2002), and an inverse hyperbolic transformation as in Schneider (2012) to accommodate negative prices.

We denote the time t mean-reverting and spike components as ξ'_t and J_t, respectively, i.e., $\vec{p}_t \equiv (\xi'_t, J_t)$. The corresponding electricity price, $P_t(\xi'_t, J_t)$, is defined as $P'_t(\xi'_t) + J_t$, where $P'_t(\xi'_t)$ is the despiked price. We model the dynamics of the spike component, J_t, as a compound Bernoulli process in which a spike occurs at time t with probability λ and its size follows an empirical distribution. We assume that the despiked price, $P'_t(\xi'_t)$, satisfies $\sinh^{-1}\left(P'_t(\xi'_t)/\ell\right) = \xi'_t + f'(t)$, where $\sinh^{-1}$ is the inverse hyperbolic sine function, ℓ is a scale parameter, and $f'(t)$ is the deterministic seasonality function. The inverse hyperbolic sine function is analogous to the natural logarithm function, a commonly used transformation of commodity prices (Lucia and Schwartz 2002), but it can be employed with negative prices. Unlike Schneider (2012), we apply this transformation to the despiked price rather than directly to the price to avoid unrealistically large spikes. We model the dynamics of the mean reverting component as the AR(1) process $\xi'_t = (1-\kappa)\xi'_{t-1} + \sigma'\epsilon'_t$, where κ is the speed of mean reversion, σ' is the volatility, and each ϵ'_t is an i.i.d. standard normal error term. This process reverts to zero, because we include the mean of the despiked and transformed price process, $\sinh^{-1}\left(P'_t(\xi'_t)/\ell\right)$, in the seasonality function. We specify this function as $f'(t) = \gamma^3 + \sum_{i=1}^{11} \gamma^{4i} D_t^{4i} + \sum_{j=1}^{5} \gamma^{5j} D_t^{5j} + \sum_{h=1}^{23} \gamma^{6h} D_t^{6h}$, where γ^3 is a constant, and γ^{4i}, γ^{5j}, and γ^{6h} are the respective coefficients of the dummy variables D_t^{4i}, D_t^{5j}, and D_t^{6h} that equal one if period t is in month i, week day j, and hour h, respectively, and zero otherwise.

We calibrate our price model to the 2005-2008 NYISO New York City zone real time prices presented in Figure 4. In addition to mean reversion and a substantial number of positive and negative spikes, this price series exhibits 1,898 negative prices during four years ($420{,}768 = (365 \times 3 + 366) \times 24 \times 12$ five-minute intervals), corresponding to a frequency of 0.45%. Albeit not apparent in Figure 4, this data displays seasonality at various time scales. Our chosen price model includes all these features.

In our price model, the despiked price in period $t+1$, $P'_{t+1}(\xi'_{t+1})$, conditional on the value of the mean-reverting component in period t, ξ'_t, has a Johnson SU distribution (Johnson 1949) with mean

$$\mathbb{E}[P'_{t+1}(\xi'_{t+1})|\xi'_t] = -\ell \exp\left[0.5(\sigma')^2\right]\sinh\left[-\xi'_t(1-\kappa) - f'(t+1)\right], \tag{17}$$

where sinh is the hyperbolic sine function. Expression (17) plays an important role in our calibration.

Figure 4   **NYISO New York City zone real time prices observed during the 2005-2008 time frame**

Given a value of ℓ, we use the following iterative approach to calibrate our price model (below the index t indicates a historical data observation period, each of which is included in set $\mathcal{T}^{\mathcal{H}} := \{1, 2, \ldots, 420{,}768\}$).

Step 0. **Initialize** *the time series of estimated spike sizes,* $\{\hat{J}_t\}_{t \in \mathcal{T}^{\mathcal{H}}}$, *to zero.*

Step 1. **Remove the identified spikes** *from the time series of observed prices,* $\{\mathsf{P}_t\}_{t \in \mathcal{T}^{\mathcal{H}}}$, *to obtain a time series of estimated despiked prices,* $\{\hat{\mathsf{P}}'_t := \mathsf{P}_t - \hat{J}_t\}_{t \in \mathcal{T}^{\mathcal{H}}}$.

Step 2. **Deseasonalize** *the time series of transformed estimated despiked prices,* $\{\sinh^{-1}(\hat{\mathsf{P}}'_t/\ell)\}_{t \in \mathcal{T}^{\mathcal{H}}}$, *to obtain the time series of estimated mean-reverting components,* $\{\hat{\xi}'_t := \sinh^{-1}(\hat{\mathsf{P}}'_t/\ell) - \hat{f}'(t)\}_{t \in \mathcal{T}^{\mathcal{H}}}$, *where the calibrated seasonality function $\hat{f}'(t)$ is obtained by applying linear regression to the former time series.*

Step 3. **Calibrate the mean-reverting process** *parameters κ and σ' on the $\{\hat{\xi}'_t\}_{t \in \mathcal{T}^{\mathcal{H}}}$ time series using simple linear regression (Lucia and Schwartz 2002) to obtain their estimates $\hat{\kappa}$ and $\hat{\sigma}'$.*

Step 4. **Identify spikes** *for each $t \in \mathcal{T}^{\mathcal{H}} \setminus \{1\}$. If $\hat{J}_{t+1}$ is zero, tag price P_{t+1} as containing a spike if $|\hat{\mathsf{P}}'_{t+1} - \mathbb{E}[P'_{t+1}(\xi'_{t+1})|\hat{\xi}'_t]|$, where the expectation is given by (17), is no less than a prespecified constant (50 in our calibration). If P_{t+1} contains a spike, update $\hat{J}_{t+1}$ to $\hat{\mathsf{P}}'_{t+1} - \mathbb{E}[P'_{t+1}(\xi'_{t+1})|\hat{\xi}'_t]$. Replace $\hat{\mathsf{P}}'_{t+1}$ with $\mathbb{E}[P'_{t+1}(\xi'_{t+1})|\hat{\xi}'_t]$ and $\hat{\xi}'_{t+1}$ with $\sinh^{-1}\left(\mathbb{E}[P'_{t+1}(\xi'_{t+1})|\hat{\xi}'_t]/\ell\right) - \hat{f}'(t+1)$.*

Step 5. Stop if the parameter estimates for the AR(1) process and seasonality function have converged. Otherwise return to step 1.

We obtain a value for the estimate $\hat{\ell}$ of the scale parameter ℓ by minimizing the sum of the absolute deviations between each of the first two (unconditional) predicted and observed price moments.

Tables 3 and 4 report the estimated parameters of the AR(1) process and the transformation function and of the seasonality function, respectively. Our estimate $\hat{\lambda}$ of the spike probability is 0.0751, which we obtain as the ratio of the number of identified spikes and the number of periods in

Table 3   Estimated parameters of the price AR(1) process and the transformation function

AR(1) process		Transformation function
$\hat{\kappa}$	$\hat{\sigma}'$	$\hat{\ell}$
0.1176	0.1770	30

Table 4   Estimated parameters of the price seasonality function

| $\hat{\gamma}^3$ | 1.3778 | | | | | | | | | | |

					i						
	1	2	3	4	5	6	7	8	9	10	11
$\hat{\gamma}^{4i}\,(\cdot 10^{-3})$	9	25.9	40.1	57	−28.9	83.5	214.6	177.4	6.4	−5.3	−72.1

			j				
	1	2	3	4	5	6
$\hat{\gamma}^{5j}\,(\cdot 10^{-3})$	−40.2	−97.6	−5.6	14	20.9	33

						h						
	1	2	3	4	5	6	7	8	9	10	11	12
$\hat{\gamma}^{6h}\,(\cdot 10^{-3})$	−73.5	−118.8	−177.6	−194	−153.8	−79.2	36.6	90.6	186.1	265.7	302.6	324.8
						h						
	13	14	15	16	17	18	19	20	21	22	23	
$\hat{\gamma}^{6h}\,(\cdot 10^{-3})$	320.7	317.1	302.7	295.5	310.1	356	337.7	339.9	313.1	243.7	145.7	

Figure 5   Empirical spike distribution $\left(\hat{\lambda} = 0.0751\right)$
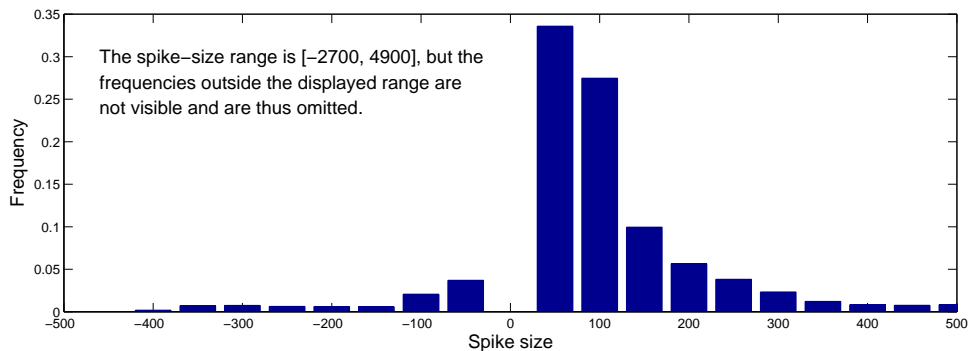
our data set. We construct the empirical spike distribution displayed in Figure 5 based on the spikes extracted by our calibration procedure. The MAE of our calibrated model is \$7.63/MWh, whereas the average of the observed prices is \$85.12/MWh. The first two (unconditional) price moments estimated on a set of 10,000 simulated price paths differ from their respective values computed on our data by less than 1%. The frequency of negative prices observed on these simulated paths is 0.44%, which compares favorably with the empirical frequency of 0.45%.

We obtain a value of 0.0147 for the correlation between the residuals of the hourly wind speeds and the averages of the residuals of the despiked electricity prices for each hour from our calibration. Therefore, for both simplicity and consistency with our model set up, we assume the wind speed and price AR(1) processes are independent in our analysis. However, our calibrated parameters

of the seasonality functions capture deterministic relationships between the observed wind speeds and electricity prices, e.g., strong wind and low electricity price at night.

6.3. Computational approach

We employ a discretized version of our MDP (7) to carry out our computations.

We discretize the calibrated five-minute wind speed AR(1) process as a period-independent grid that includes the integers between 0 and 25 (m/s), which is the cut-off speed of the GE 1.5-77 turbine. To obtain the transition probability between any two such levels on this grid, we minimize the squared difference between the first two moments of the discretized wind speed model and those of its continuous counterpart. Given a wind speed value, which is the sum of the values of its AR(1) and seasonality components, we obtain its corresponding available energy amount applying linear interpolation to the production curve displayed in Table 1.

Based on the method in Jaillet et al. (2004), we discretize the calibrated price AR(1) process as a trinomial lattice with five-minute time increments that specifies attainable despiked price levels and their transition probabilities for each stage, assuming the market price of risk (Duffie 1992) is zero. With $\xi'_0 = 0$ the constructed lattice converges to 11 levels in the 6-th stage.

We discretize the feasible inventory set using a grid with $120C^{\mathrm{S}}/200 + 1$ evenly spaced levels between 0 and the storage energy capacity C^{S} (200 MWh is the smallest value of this parameter that we consider). We use a set of feasible action pairs that is consistent with the feasible inventory and available wind energy sets to compute both our optimal policy and the inventory-change part of the H1 policy. The initial state $S_0 \equiv (x_0, w_0, \xi'_0, J_0)$ is $(0, 4.69, 0, 0)$, where 4.69 MWh corresponds to a wind speed of 7.238 m/s.

Appendix A includes the details of the computation of both the optimal policy and the H1, H2, H3, and H3+ policies for our discretized MDP. We evaluate the H2, H3, and H3+ policies by applying (8) to this version of our MDP (computing the optimal policy and the H1 policy also yields their respective values).

6.4. Performance of heuristics

Our study is based on an extensive set of instances. In addition to the six values for each of the storage energy capacity and the transmission capacity discussed in §6.1, we consider six values for the frequency of negative price occurrence, namely, 0.05%, 1%, 5%, 10%, 15%, and 20%, which we obtain by varying the probability of observing negative spikes. Thus, the total number of examined instances is $6^3 = 216$.

We assess the performance of all the heuristics based on the optimality gaps of their policies in the initial stage and state (i.e., $V_0^*(S_0) - V_0^\pi(S_0)$ with $\pi \in \{\text{H1, H2, H3, H3+}\}$) and their run times.

Table 5   **Percentage optimality gaps of all the heuristics across all the instances**

	H1	H2	H3	H3+
Average	0	2.86	11.39	6.05
Min	0	0.95	2.95	2.53
Max	0	6.49	22.37	11.16

Table 6   **CPU minutes required to compute the optimal policy and to execute all the heuristics across all the instances**

	Optimal policy	H1	H2	H3	H3+
Average	588.42	33.41	0.24	8.55	8.76
Min	57.68	9.30	0.03	2.12	2.22
Max	1449.65	60.33	0.90	15.88	16.57

Note. All experiments are run on a computer with Intel(R) Core(TM) i7-3770K 3.40GHz CPU and 8 GB RAM.

Tables 5 and 6 display these metrics, respectively. Table 6 excludes the time taken to evaluate the H2, H3, and H3+ policies.

The H1 policy is optimal on each considered instance. Furthermore, it continues to be optimal in the following two additional sets of experiments: When more negative prices are obtained by reducing the values of the negative hourly seasonality coefficients, i.e., decreasing $\hat{r}^{6h}$, for $h = 1, \ldots, 23$, in Table 4 if $\hat{r}^{6h} < 0$ instead of increasing the probability of negative spike occurrence, and when we vary the estimated mean-reversion parameter, $\hat{\kappa}$. (More precisely, the optimality gap of the H1 policy is less than 0.006% on all these instances; we attribute the presence of such a positive, yet very small, gap to numerical approximations in our computations.) Thus, while the H1 policy is not optimal in general, as illustrated in Online Appendix D, our results suggest that in practice the optimal policy may be very similar, if not identical, to the H1 policy. On average it takes about 10 hours and 33 minutes, respectively, to compute the optimal policy and to execute H1; i.e., running H1 is on average about 17 times faster than obtaining the optimal policy. This difference occurs because the computation of the H1 policy avoids looping over the possible values of one component of the state and one decision variable compared to that of the optimal policy (see Appendices A.1 and A.2).

The H2 policy is near optimal, with average and maximal optimality gaps equal to 2.86% and 6.49%, respectively, and outperforms the H3 and H3+ policies, whose respective average-and-maximal optimality gaps are 11.39%-and-22.37% and 6.05%-and-11.16%. (The optimality gaps of the H2 policy are always smaller than those of the H3 policy, and exceed those of the H3+ policy in only 1 out of the 216 instances by 0.5%.) The H2 policy outperforms the H3 policy and, overall, the H3+ policy because whereas the latter two policies make decisions solely based on the price information in the current period, the H2 policy also uses knowledge of the inventory level of the

storage facility. Our results suggest that even if it has a stationary nature, the H2 policy is a reasonable approximation for the optimal policy, which in general is non-stationary. On average, H2, which takes less than 15 seconds to execute, is faster than H1 by two orders of magnitude and both the H3 and H3+ policies, which require about 9 minutes of run time, by one order of magnitude. Even though the H2 policy depends on more parameters than the H3 and H3+ policies do, the algorithm that we employ to compute the former policy becomes frequently trapped in local optima when applied in the context of the latter two policies. Therefore, we must resort to using a slower method to obtain them (see Appendix A.3).

6.5. The effect of negative prices on the value added by and environmental benefit of storage

We define the value added by storage as the difference between the optimal value of the WST system, $V_0^*(S_0)$, and the one of the wind energy production and transmission (WT) system that has no storage (NS) and thus optimally sells as much generated energy as possible when prices are positive and curtails otherwise, $V_0^{\mathrm{NS}}(S_0)$. Figure 6 displays these values and their difference as functions of the frequency of occurrence of negative prices for 600 MWh storage energy capacity and 120 MW transmission capacity, which are typical when expressed as ratios of the generation capacity (Denholm and Sioshansi 2009, Pattanariyankool and Lave 2010; other choices of these parameters yield qualitatively similar results). When this frequency increases from 0.5% to 20%, the optimally managed WST and WT systems, respectively, gain and lose progressively more value, so that the value added by storage rises rapidly, from 30.4% to 102.4% of $V_0^{\mathrm{NS}}(S_0)$.

To obtain some understanding for this finding, we break down the total expected optimal amount of stored electricity, $\sum_{t \in \mathcal{T}} \mathbb{E}\left[(a_t^*(S_t)/\alpha)^+|S_0\right]$, into the total expected optimally (i) purchased energy net of the transmission loss, $\sum_{t \in \mathcal{T}} \mathbb{E}\left[(a_t^*(S_t)/\alpha - g_t^*(S_t))^+|S_0\right]$, and (ii) stored generation, $\sum_{t \in \mathcal{T}} \mathbb{E}\left[(a_t^*(S_t)/\alpha)^+|S_0\right]$ minus the previous quantity (all these metrics are expressed before applying the charging loss). The left panel of Figure 7 presents the values of these two components for the same configuration that underlies Figure 6: The first one increases whereas the second one decreases when negative prices become more frequent. Intuitively, it is more valuable to store electricity purchased at a negative price than wind energy generated for free. This analysis suggests that the increase in the incremental value brought about by storage stems from the increased amount of purchased energy when negative prices occur more frequently.

We also examine how the frequency of negative prices affects the environmental benefit of storage, measured as the total expected reduction of the amount of curtailed wind energy that storage enables: $\sum_{t \in \mathcal{T}} \mathbb{E}[(w_t - g_t^{\mathrm{NS}}(S_t)) - (w_t - g_t^*(S_t))|S_0] = \sum_{t \in \mathcal{T}} \mathbb{E}[g_t^*(S_t) - g_t^{\mathrm{NS}}(S_t)|S_0]$, where $g_t^{\mathrm{NS}}(S_t)$ is

Figure 6 **The optimal WST and WT system values (left panel) and their difference (right panel) for 600 MWh storage energy capacity and 120 MW transmission capacity**


Figure 7 **The two components of the total expected optimal stored electricity (left) and the reduced curtailment (right) for the setup of Figure 6**

the optimal generation in stage t and state S_t of the WT system. The right panel of Figure 7 shows that reduced curtailment decreases when negative prices become more numerous, because energy acquired in the market increasingly displaces from storage energy generated by the wind farm. Transportation inefficiencies render this use of storage environmentally unappealing even if the purchased energy were as "clean" as wind energy. This issue would be more acute if this replacement involved "dirty" energy obtained from sources such as coal or natural gas. Thus, the environmental benefit of storage erodes as negative electricity prices occur more frequently in our

model.

7. Conclusions

We examine the merchant management of a WST system modeled as an MDP with stochastic wind energy availability and electricity prices. Characterizing any optimal policy structure in the general case when these prices can be negative is difficult. However, we establish that a stage- and partial-state-dependent-threshold policy is optimal in the special case when they are always positive. We extend it as an approximation to the general case and further simplify it to a price-dependent-threshold policy. H1 and H2 use these policies, respectively. Whereas H1 requires approximately solving a stochastic dynamic program, H2 relies on a derivative-free optimization algorithm embedded in Monte Carlo simulation of the stochastic processes of our MDP. We analyze the performance of these and other known heuristics. We also investigate the value added by and the environmental benefit of storage.

Using data-calibrated models of wind speed and electricity prices that allow us to compute an optimal policy and execute H1, we find that the H1 policy is optimal, even when these prices are negative 20% of the time. This observation suggests that this policy may be optimal for most, if not all, practical instances, even though we present a pathological counter example. Executing H1 is much faster than computing an optimal policy explicitly: Their respective average computational requirements are 33 minutes and 10 hours. On average, the optimality gap of the H2 policy equals 2.86% and H2 has a run time of 15 seconds. Hence, despite the suboptimality of its policy, H2 is more practical than H1. Further, the H2 policy generally outperforms two variants of simple and known policies that rely exclusively on the electricity price, because it depends on inventory availability too. It is also faster to compute than these policies. When we amplify the frequency of negative price occurrence, the optimal amount of purchased electricity rises and the curtailment reduction drops, so that storage increases the value of the WST system but its own environmental benefit shrinks.

Appendix

A. Details about the computation of the considered policies

We describe the computation of the optimal policy in §A.1, the H1 policy in §A.2, and the H2, H3, and H3+ policies in §A.3, based on the discretization of our MDP (7) discussed in §6.3.

A.1. Computation of the optimal policy

We denote by $\mathcal{X}^{\mathrm{D}}$, $\mathcal{W}^{\mathrm{D}}$, and $\mathcal{J}^{\mathrm{D}}$ the discretized inventory, available wind energy, and spike price component sets, respectively, and by N_x, N_w, and N_J their respective cardinalities. We let the discretized set of values of the mean-reverting price component and its cardinality in stage t be Ξ'^{D}_t and $N_{\xi',t}$. We define $N_{\xi'}$ as the maximum of all the $N_{\xi',t}$'s. For each inventory and available wind energy pair $(x_t, w_t) \in \mathcal{X}^{\mathrm{D}} \times \mathcal{W}^{\mathrm{D}}$, to reduce the effect of the approximation induced by the discretization, we include in the discretized action set $\Psi^{\mathrm{D}}(x_t, w_t)$ all the action pairs from set $\Psi(x_t, w_t)$ that satisfy $a_t = x_{t+1} - x_t$ for some $x_{t+1} \in \mathcal{X}^{\mathrm{D}}$ and $g_t \in \mathcal{W}^{\mathrm{D}}$, as well as the extreme points of set $\Psi(x_t, w_t)$.

We compute the optimal policy and value function by standard backward dynamic programming on the discretized state space $\cup_{t \in \mathcal{T}} \mathcal{X}^{\mathrm{D}} \times \mathcal{W}^{\mathrm{D}} \times \Xi'^{\mathrm{D}}_t \times \mathcal{J}^{\mathrm{D}}$. To speed up computation and save memory, we exploit the state independence of the spike price component by evaluating for each partial state $(x_{t+1}, w_{t+1}, \xi'_{t+1})$ in stage $t+1$ the function $\overline{V}^*_{t+1}(x_{t+1}, w_{t+1}, \xi'_{t+1}) := \mathbb{E}\left[V^*_{t+1}(x_{t+1}, w_{t+1}, \xi'_{t+1}, J_{t+1})\right]$, where expectation is with respect to the discrete distribution of the random variable J_{t+1}. With a slight abuse of notation, we then evaluate the version $U^*_t(x_{t+1}, w_t, \xi'_t)$ of the stage t optimal continuation function as $\delta \mathbb{E}\left[\overline{V}^*_{t+1}(x_{t+1}, w_{t+1}, \xi'_{t+1}) \mid w_t, \xi'_t\right]$, where expectation is with respect to the discrete joint distribution of the random variables w_{t+1} and ξ'_{t+1} conditional on the known pair (w_t, ξ'_t). Given how we model the evolution of J_t, $U^*_t(x_{t+1}, w_t, \xi'_t, J_t)$ is identical to $U^*_t(x_{t+1}, w_t, \xi'_t)$. We thus equivalently express (9) as

$$\max_{(a_t, g_t) \in \Psi^{\mathrm{D}}(x_t, w_t)} \{R(a_t, g_t, \xi'_t, J_t) + U^*_t(x_t + a_t, w_t, \xi'_t)\}. \tag{18}$$

When solving (18), a next-stage inventory level $x_t + a_t$ may not belong to the set $\mathcal{X}^{\mathrm{D}}$ due to the inclusion in set $\Psi^{\mathrm{D}}(x_t, w_t)$ of the extreme points of set $\Psi(x_t, w_t)$. In this case, we linearly interpolate the values of the considered version of the optimal continuation function for the two inventory levels in set $\mathcal{X}^{\mathrm{D}}$ that are adjacent to this next-stage inventory level. When multiple action pairs are optimal, we select the one that results in the largest next-stage inventory level.

The number of stages and states in the backward recursion that computes the optimal policy and value function is of order $O(TN_xN_wN_{\xi'}N_J)$. The number of feasible action pairs in each stage and state is of order $O(N_xN_w)$. The total number of operations taken to execute this backward recursion is thus of order $O(TN_xN_wN_{\xi'}N_J) + O(TN_xN_w^2N_{\xi'}) + O(TN_x^2N_w^2N_{\xi'}N_J) = O(TN_x^2N_w^2N_{\xi'}N_J)$, where the first, second, and third terms in this sum correspond to computing $\overline{V}^*_{t+1}$, obtaining U^*_t from $\overline{V}^*_{t+1}$ (given that ξ'_t transitions to only three possible values of ξ'_{t+1} on the trinomial lattice), and exhaustively searching for the optimal action pair, respectively.

From stage 6 onward, when the mean-reverting price lattice converges to 11 levels, the number of states per stage for the largest value of the storage energy capacity that we consider is $(120 \times 9 + 1) \times 26 \times 11 \times 68 =$

21,023,288. We thus keep in memory the values of the optimal next-stage value function only when evaluating the optimal value function in the current stage; i.e., we discard them when this computation is complete. Likewise, we do not keep in memory the optimal actions that we compute.

A.2. Computation of the H1 policy

We determine the threshold functions of the H1 policy using Algorithm 1, a backward recursion that exploits the structure of this policy discussed in §5.1. We perform this recursion based on the same discretized state and inventory action spaces used to compute the optimal policy, but directly use the generation decision rules of this policy to determine the energy production action.

Algorithm 1 Computation of the H1 policy

1: $\overline{V}_T^{\text{H1}}(x_T, w_T, \xi_T') \leftarrow 0$, $\forall (x_T, w_T, \xi_T') \in \mathcal{X}^{\text{D}} \times \mathcal{W}^{\text{D}} \times \Xi_T'^{\text{D}}$

2: **for each** $t = T-1, \cdots, 0$ **do**

3:      **for each** $(w_t, \xi_t') \in \mathcal{W}^{\text{D}} \times \Xi_t'^{\text{D}}$ **do**

4:          **for each** $x_{t+1} \in \mathcal{X}^{\text{D}}$ **do**

5:              $U_t^{\text{H1}}(x_{t+1}, w_t, \xi_t') \leftarrow \delta \mathbb{E}\left[\overline{V}_{t+1}^{\text{H1}}(x_{t+1}, w_{t+1}, \xi_{t+1}') \mid w_t, \xi_t'\right]$.

6:          **end for**

7:      **end for**

8:      **for each** $(w_t, \xi_t', J_t) \in \mathcal{W}^{\text{D}} \times \Xi_t'^{\text{D}} \times \mathcal{J}^{\text{D}}$ **do**

9:          Compute $X_t^{(\nu),\text{H1}}(w_t, \xi_t', J_t)$ for $\nu \in \{1, 2, 3, 4\}$ by applying exhaustive search to (13)-(16) with $U_t^{\text{H1}}(y_t, w_t, \vec{p}_t)$ replaced by $U_t^{\text{H1}}(y_t, w_t, \xi_t')$.

10:      **end for**

11:      **for each** $(x_t, w_t, \xi_t', J_t) \in \mathcal{X}^{\text{D}} \times \mathcal{W}^{\text{D}} \times \Xi_t'^{\text{D}} \times \mathcal{J}^{\text{D}}$ **do**

12:          Determine $(a_t^{\text{H1}}(x_t, w_t, \xi_t', J_t), g_t^{\text{H1}}(x_t, w_t, \xi_t', J_t))$ as explained in §5.1.

13:          $V_t^{\text{H1}}(x_t, w_t, \xi_t', J_t) \leftarrow R(a_t^{\text{H1}}(x_t, w_t, \xi_t', J_t), g_t^{\text{H1}}(x_t, w_t, \xi_t', J_t), \xi_t', J_t) + U_t^{\text{H1}}(x_t + a_t^{\text{H1}}(x_t, w_t, \xi_t', J_t), w_t, \xi_t')$.

14:      **end for**

15:      **for each** $(x_t, w_t, \xi_t') \in \mathcal{X}^{\text{D}} \times \mathcal{W}^{\text{D}} \times \Xi_t'^{\text{D}}$ **do**

16:          $\overline{V}_t^{\text{H1}}(x_t, w_t, \xi_t') \leftarrow \mathbb{E}\left[V_t^{\text{H1}}(x_t, w_t, \xi_t', J_t)\right]$.

17:      **end for**

18: **end for**

We define the functions $U_t^{\text{H1}}(x_{t+1}, w_t, \xi_t')$ and $\overline{V}_{t+1}^{\text{H1}}(x_{t+1}, w_{t+1}, \xi_{t+1}')$ for the H1 policy analogously to how we specify the versions of these functions for the optimal policy. Specifically, for each stage t we evaluate the function $U_t^{\text{H1}}(x_{t+1}, w_t, \xi_t')$ as $\delta \mathbb{E}\left[\overline{V}_{t+1}^{\text{H1}}(x_{t+1}, w_{t+1}, \xi_{t+1}') \mid w_t, \xi_t'\right]$ and the function $\overline{V}_{t+1}^{\text{H1}}(x_{t+1}, w_{t+1}, \xi_{t+1}')$ as $\mathbb{E}\left[V_{t+1}^{\text{H1}}(x_{t+1}, w_{t+1}, \xi_{t+1}', J_{t+1}) | w_{t+1}, \xi_{t+1}'\right]$ in lines 3-7 and 15-17, respectively, of Algorithm 1. We compute each function $X_t^{(\nu),\text{H1}}(w_t, \xi_t', J_t)$, with $\nu \in \{1, 2, 3, 4\}$, via an exhaustive search in lines 8-10 of this algorithm.

We determine the action pair $(a_t^{\mathrm{H1}}(x_t, w_t, \xi_t', J_t), g_t^{\mathrm{H1}}(x_t, w_t, \xi_t', J_t))$ in stage t and state (x_t, w_t, ξ_t', J_t) based on the structure of the H1 policy discussed in §5.1 and evaluate the value function $V_t^{\mathrm{H1}}(x_t, w_t, \xi_t', J_t)$ in this stage and state in lines 11-14 of Algorithm 1. We use linear interpolation when needed.

The total number of operations of Algorithm 1 is of order $O(TN_x N_w N_{\xi'} N_J) + O(TN_x N_w^2 N_{\xi'}) = O(TN_x N_w N_{\xi'}(N_J + N_w))$, which amounts to a reduction of order $O(N_x N_w N_J/(N_J + N_w))$ compared to the computation of the optimal policy.

A.3. Computation of the H2, H3, and H3+ policies

We use the following three steps to obtain the parameters of the H2 policy:

Step 1. Generate N sample paths of the available wind energy and electricity price from stage 1 through $T - 1$, starting from the stage 0 values w_0 and ξ_0 (recall that the price spike component random variables are i.i.d.).

Step 2. For each sample path n, use the Nelder-Mead simplex method, a derivative-free nonlinear optimization method (Lagarias et al. 1998), to choose values for the decision variables $P^{S,n}$ and $X^{(\nu),n}$ for $\nu \in \{1, 2, 3, 4\}$ with the goal of maximizing the total discounted cash flows that result from using the decision rules of the H2 policy corresponding to these variables on this sample path. Specifically, we apply the Nelder-Mead method to the following sample path model:

$$\max_{P^{S,n}, (X^{(\nu),n}, \nu \in \{1,2,3,4\})} \sum_{t \in \mathcal{T}} \delta^t R(A_t^{\mathrm{H2}}(S_t^{\mathrm{H2},n}; P^{S,n}, X^{(\nu),n}, \nu \in \{1,2,3,4\}), \xi_t'^{,n}, J_t^n),$$
$$\text{s.t. } P^{S,n} \geq 0, 0 \leq X^{(1),n} \leq X^{(2),n} \leq X^{(3),n} \leq X^{(4),n} \leq C^{\mathrm{S}},$$

where $S_t^{\mathrm{H2},n}$, $\xi_t'^{,n}$, and J_t^n are analogous to S_t^{H2}, ξ_t', and J_t but are specific to sample path n; we make explicit the dependence of the H2 policy decision rules on the decision variables; and we evaluate the objective function using a backward recursion that relies on the discretized inventory set (see §6.3). Let $P^{S,n,\mathrm{NM}}$ and $X^{(\nu),n,\mathrm{NM}}$, with $\nu \in \{1, 2, 3, 4\}$, denote the solution for this model found by the Nelder-Mead algorithm (we employ the one available in the GNU Scientific Library[2]).

Step 3. Set P^{S} and $X^{(\nu)}$, with $\nu \in \{1, 2, 3, 4\}$, equal to the averages of the values $P^{S,n,\mathrm{NM}}$'s and $X^{(\nu),n,\mathrm{NM}}$'s for $\nu \in \{1, 2, 3, 4\}$ obtained across the N sample paths, i.e., $P^{\mathrm{S}} = \sum_{n=1}^{\mathsf{N}} P^{S,n,\mathrm{NM}}/\mathsf{N}$ and $X^{(\nu)} = \sum_{n=1}^{\mathsf{N}} X^{(\nu),n,\mathrm{NM}}/\mathsf{N}$, with $\nu \in \{1, 2, 3, 4\}$.

In our numerical study we let N be equal to 5, because using larger values for N yields marginally smaller average optimality gaps (given the strong seasonality in the calibrated wind speed and electricity price models) but much longer execution times. It is impossible to perform a computational complexity analysis of our approach to determine the parameters of the H2 policy because a convergence proof for the Nelder-Meld simplex method is not available for a general function of two or more variables (Lagarias et al. 1998).

We use steps analogous to Steps 1-3 above to determine the values of the parameters of the H3 and H3+ policies, except that we rely on exhaustive search based on a grid of possible values rather than the Nelder-Mead algorithm, because it frequently becomes trapped in local minima for these policies.

---

[2] https://www.gnu.org/software/gsl/

References

Aflaki, S., S. Netessine. 2017. Strategic investment in renewable energy sources: The effect of supply intermittency. *Manufacturing & Service Operations Management* **19**(3) 489–507.

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) 1–8.

Boogert, A., C. de Jong. 2008. Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15**(3) 81–98.

Boyabatlı, Onur, J. Nguyen, T. Wang. 2017. Capacity management in agricultural commodity processing and application in the palm industry. *Manufacturing & Service Operations Management* **19**(4) 551–567.

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**(2) 523–531.

Cahn, A. S. 1948. The warehouse problem. *Bulletin of the American Mathematical Society* **54**(11) 1073.

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**(3) 1599–1606.

Charnes, A., J. Drèze, M. Miller. 1966. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* **34**(2) 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**(1) 339–368.

Costa, L.M., F. Bourry, J. Juban, G. Kariniotakis. 2008. Management of energy storage coordinated with wind power under electricity market conditions. *PMAPS '08. Proceedings of the 10th International Conference on Probabilistic Methods Applied to Power Systems, 2008*. 1–8.

Denholm, P., R. Sioshansi. 2009. The value of compressed air energy storage with wind in transmission-constrained electric power systems. *Energy Policy* **37**(8) 3149–3158.

Devalkar, S.K., R. Anupindi, A. Sinha. 2011. Integrated optimization of procurement, processing, and trade of commodities. *Operations Research* **59**(6) 1369–1381.

DOE. 2017. New York wind resource map and potential wind capacity. Available at `http://apps2.eere.energy.gov/wind/windexchange/wind_resource_maps.asp?stateab=ny`; last accessed in January, 2017.

Duffie, D. 1992. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ.

Duke Energy. 2017. Joint open access transmission tariff of Duke Energy Carolinas, LLC, Duke Energy Florida, LLC, and Duke Energy Progress, LLC. Available at `http://www.ferc.duke-energy.com/Tariffs/Joint_OATT.pdf`; last accessed in July, 2017.

EPRI. 2004. EPRI-DOE handbook supplement of energy storage for grid connected wind generation applications. Tech. rep., Electric Power Research Institute. Available at `http://www.sandia.gov/ess/publications/EPRI-DOE%20ESHB%20Wind%20Supplement.pdf`; last accessed in May, 2018.

ERCOT. 2012. Balancing energy services market clearing prices for energy annual reports archives. Available at `www.ercot.com/mktinfo/prices/mcpea`.

Fanone, E., A. Gamba, M. Prokopczuk. 2013. The case of negative day-ahead electricity prices. *Energy Economics* **35** 22–34.

Fertig, E., J. Apt. 2011. Economics of compressed air energy storage to integrate wind power: A case study in ERCOT. *Energy Policy* **39**(5) 2330–2342.

General Electric. 2013. Production curve of GE 1.5-77 turbines. Available at `www.inl.gov/wind/software/powercurves/pc_ge_wind.xls`; last accessed in 2013.

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. *2010 7th International Conference on the European Energy Market (EEM)*. 1–6.

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**(2) 460–468.

Graves, F., T. Jenkin, D. Murphy. 1999. Opportunities for electricity storage in deregulating markets. *The Electricity Journal* **12**(8) 46–56.

Guo, S., S. Hu, G.C. Souza. 2017. Batteries as energy storage for energy shifting. Working paper, Kelly School of Business, Bloomington, IN.

Harsha, P., M. Dahleh. 2015. Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable energy. *IEEE Transactions on Power Systems* **30**(3) 1164–1181.

Heier, S. 2006. *Grid Integration of Wind Energy Conversion Systems*. John Wiley & Sons. Hoboken, New Jersey.

Hu, S., G. Souza, M. Ferguson, W. Wang. 2015. Capacity investment in renewable energy technology with supply intermittency: Data granularity matters *Manufacturing & Service Operations Management* **17**(4) 480–494.

Huntowski, F., A. Patterson, M. Schnitzer. 2012. Negative electricity prices and the production tax credit. Available at `www.hks.harvard.edu/hepg/Papers/2012/Negative_Electricity_Prices_and_the_Production_Tax_Credit_0912.pdf`.

Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50**(7) 909–921.

Jiang, D.R., W.B. Powell. 2015. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing* **27**(3) 525–543.

Johnson, N.L. 1949. Systems of frequency curves generated by methods of translation. *Biometrika* **36**(1/2) 149–176.

Kim, J.H., W.B. Powell. 2011. Optimal energy commitments with storage and intermittent supply. *Operations Research* **59**(6) 1526–5463.

Knittel, K., M. Roberts. 2005. An empirical examination of restructured electricity prices. *Energy Economics* **27**(5) 791–817.

Kök, A.G., K. Shang, Ş. Yücel. 2018. Impact of electricity pricing policies on renewable energy investments and carbon emissions. *Management Science* **64**(1) 131–148.

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**(8) 599–606.

Lagarias, J.C., J.A. Reeds, M.H. Wright, P.E. Wright. 1998. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM Journal on optimization* **9**(1) 112–147.

Lai, G., F. Margot, N. Secomandi. 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58**(3) 564–582.

Lai, G., M.X. Wang, S. Kekre, A. Scheller-Wolf, N. Secomandi. 2011. Valuation of storage at a liquefied natural gas terminal. *Operations Research* **59**(3) 602–616.

Löhndorf, N., S. Minner. 2010. Optimal day-ahead trading and storage of renewable energies—an approximate dynamic programming approach. *Energy Systems* **1**(1) 61–77.

Lucia, J., E. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research* **5**(1) 5–50.

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.

Nadarajah, S., F. Margot, N. Secomandi. 2015. Relaxations of approximate linear programs for the real option management of commodity storage. *Management Science* **61**(12) 3054–3076.

Nasakkala, E., J. Keppo. 2008. Hydropower with financial information. *Applied Mathematical Finance* **15**(5-6) 503–529.

NOAA. 2010. Integrated surface data, hourly, global. Available at `https://www.ncdc.noaa.gov/data-access`; last accessed in May, 2018.

NYISO. 2013. NYISO electricty pricing data. Available at `www.nyiso.com/public/markets_operations/market_data/pricing_data/index.jsp`; last retrieved in June, 2013.

Pattanariyankool, S., L.B. Lave. 2010. Optimizing transmission from distant wind farms. *Energy Policy* **38**(6) 2806–2815.

Powell, W.B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, Hoboken, NJ.

Powell, W.B., S. Meisel. 2016a. Tutorial on stochastic optimization in energy—Part I: Modeling and policies. *IEEE Transactions on Power Systems* **31**(2) 1459–1467.

Powell, W.B., S. Meisel. 2016b. Tutorial on stochastic optimization in energy—Part II: An energy storage illustration. *IEEE Transactions on Power Systems* **31**(2) 1468–1475.

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 `http://www.ren21.net/Portals/0/documents/activities/gsr/REN21_GSR_2010_full_revised%20Sept2010.pdf`; last accessed in May, 2018.

Schneider, S. 2012. Power spot price models with negative prices. *The Journal of Energy Markets* **4**(4) 77–102.

Secomandi, N. 2010. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.

Secomandi, N. 2015. Merchant commodity storage practice revisited. *Operations Research* **63**(5) 1131–1143.

Secomandi, N., G. Lai, F. Margot, A. Scheller-Wolf, D.J. Seppi. 2015. Merchant commodity storage and term structure model error. *Manufacturing & Service Operations Management* **17**(3) 302–320.

Secomandi, N., D.J. Seppi. 2014. Real options and merchant operations of energy and other commodities. *Foundations and Trends in Technology, Information and Operations Management* **6**(3-4) 161–331.

Seifert, J., M. Uhrig-Homburg. 2007. Modelling jumps in electricity prices: theory and empirical evidence. *Review of Derivatives Research* **10**(1) 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., Cyriel De Jong. 2003. Negative prices in electricity markets. *Commodities Now* **7** 74–77.

Smith, J.E. 2005. Alternative approaches for solving real-options problems. *Decision Analysis* **2**(2) 89–102.

Sustainable Business News. 2014. AES batteries to replace dirty peak power plants. Available at `http://www.greenbiz.com/blog/2014/03/11/aes-sells-batteries-replace-peak-power-plants?mkt_tok=3RkMMJWWfF9wsRolsqTLZKXonjHpfsX56ukkW662lMI%2F0ER3fOvrPUfGjI4CS8JgI%2BSLDwEYGJlv6SgFSLHEMa5qw7gMXRQ%3D`; last accessed in May, 2018.

The Economist. 2015. Need a weatherman. Available at: `http://www.economist.com/news/china/21660164-though-wind-generation-growing-fast-much-remains-wrong-industry-need-weatherman`; last accesssed in April, 2018.

Thompson, M., M. Davison, H. Rasmussen. 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics* **56**(3) 226–238.

Wiser, R., M. Bolinger. 2013. 2013 Wind technologies market report. Tech. rep., Lawrence Berkeley National Laboratory. Available at: `https://www.energy.gov/eere/wind/downloads/2013-wind-technologies-market-report`; last accessed in May, 2018.

Wiser, R., M. Bolinger. 2015. 2015 Wind technologies market report. Tech. rep., U.S. Department of Energy. Available at: `https://energy.gov/sites/prod/files/2016/08/f33/2015-Wind-Technologies-Market-Report-08162016.pdf`; last accessed in May, 2018.

Wu, O., R. Kapuscinski. 2013. Curtailing intermittent generation in electrical systems. *Manufacturing & Service Operations Management* **15**(4) 578–595.

Wu, O., D. Wang, Z. Qin. 2012. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management* **14**(3) 455–471.

Xi, X., R. Sioshansi, V. Marano. 2014. A stochastic dynamic programming model for co-optimization of distributed energy storage. *Energy Systems* **5**(3) 475–505.

Zhou, Y., A. Scheller-Wolf, N. Secomandi, S.F. Smith. 2016. Electricity trading and negative prices: Storage vs. disposal. *Management Science* **62**(3) 880–898.

Online Appendix

Appendix A: Proof of Proposition 1

We use Lemma 1 in the proof of Proposition 1.

LEMMA 1. *For each $t \in \mathcal{T} \cup \{T\}$ and $S_t \in \mathcal{S}$, it holds that $|V_t^*(S_t)| < \infty$.*

Proof: When $t = T$ the claimed property is true as $V_T^*(S_T) \equiv 0$. For each stage $t \in \mathcal{T}$ the quantity sold or bought cannot exceed C^{T}/τ. Thus, it holds that $|V_t^*(S_t)| \leq \sum_{k=t}^{T-1} |\mathbb{E}[P_k(\vec{p}_k)|\vec{p}_t]|C^{\mathrm{T}}/\tau \leq \sum_{k=t}^{T-1} \mathbb{E}[|P_k(\vec{p}_k)||\vec{p}_t]C^{\mathrm{T}}/\tau < \infty$, where the last inequality follows from Assumption 1. $\square$

We convert model (9) to an equivalent one with linear constraints. We define the following new decision variables:

- a_t^{IG}: Amount of inventory injection from energy generated;
- a_t^{IB}: Amount of inventory injection from energy bought;
- a_t^{W}: Amount of inventory withdrawal;
- g_t^{I}: Amount of energy generated and injected into the storage facility; and
- g_t^{S}: Amount of energy generated for sale;

We define $\Psi^1(x_t, w_t)$ as the set of decision variable tuples $(a_t^{\mathrm{IG}}, a_t^{\mathrm{IB}}, a_t^{\mathrm{W}}, g_t^{\mathrm{I}}, g_t^{\mathrm{S}})$ in $\mathbb{R}_+^5$ that satisfy

$$g_t^{\mathrm{S}} + \beta a_t^{\mathrm{W}} \leq C^{\mathrm{T}}, \tag{19}$$

$$a_t^{\mathrm{IB}}/(\alpha\tau) \leq C^{\mathrm{T}}, \tag{20}$$

$$a_t^{\mathrm{IG}} = \alpha g_t^{\mathrm{I}}, \tag{21}$$

$$g_t^{\mathrm{I}} + g_t^{\mathrm{S}} \leq w_t, \tag{22}$$

$$a_t^{\mathrm{W}} \leq x_t, \tag{23}$$

$$a_t^{\mathrm{IG}} + a_t^{\mathrm{IB}} \leq C^{\mathrm{S}} - x_t, \tag{24}$$

$$a_t^{\mathrm{W}} \leq C^{\mathrm{D}}, \tag{25}$$

$$a_t^{\mathrm{IG}} + a_t^{\mathrm{IB}} \leq C^{\mathrm{C}}, \tag{26}$$

$$g_t^{\mathrm{I}}, g_t^{\mathrm{S}}, a_t^{\mathrm{IG}}, a_t^{\mathrm{IB}}, a_t^{\mathrm{W}} \geq 0, \tag{27}$$

where (19) and (20) are the transmission capacity constraints; (21) expresses the relationship between the two decision variables a_t^{IG} and g_t^{I}; (22) is the wind energy availability constraint; (23) is the available inventory constraint; (24) is the storage space constraint; (25) and (26) are the discharging and charging capacity constraints, respectively; and (27) imposes the nonnegativity of all the decision variables.

Consider the model

$$\max_{(a_t^{\mathbf{G}}, a_t^{\mathbf{B}}, a_t^{\mathbf{W}}, g_t^{\mathbf{I}}, g_t^{\mathbf{S}}) \in \Psi^1(x_t, w_t)} \left\{ P_t(\vec{p}_t)(g_t^{\mathrm{S}} + \beta a_t^{\mathrm{W}})\tau - P_t(\vec{p}_t)a_t^{\mathrm{IB}}/(\alpha\tau) \right. \\ \left. + \delta \mathbb{E}\left[V_{t+1}^*(x_t + a_t^{\mathrm{IG}} + a_t^{\mathrm{IB}} - a_t^{\mathrm{W}}, w_{t+1}, \vec{p}_{t+1})|S_t \right] \right\}. \tag{28}$$

We show that this model is equivalent to the one on the right hand side of (9). To do so we only need to establish that a feasible solution for model (28) that has both a_t^{W} and $a_t^{\mathrm{IG}} + a_t^{\mathrm{IB}}$ strictly positive cannot be optimal. Suppose there exists an optimal solution $(\bar{a}_t^{IG}, \bar{a}_t^{IB}, \bar{a}_t^W, \bar{g}_t^I, \bar{g}_t^S)$ to (28) such that $\bar{a}_t^W > 0$ and $\bar{a}_t^{IG} + \bar{a}_t^{IB} > 0$. Consider the following cases:

- Case 1: If $\overline{a}_t^{IB} > 0$, let $\check{a}_t^{IG} := \overline{a}_t^{IG}, \check{a}_t^{IB} := \overline{a}_t^{IB} - \Delta^{(1)}, \check{a}_t^W := \overline{a}_t^W - \Delta^{(1)}, \check{g}_t^I := \overline{g}_t^I$, and $\check{g}_t^S = \overline{g}_t^S$, where $\Delta^{(1)} \in (0, \min\{\overline{a}_t^{IB}, \overline{a}_t^W\}]$.

- Case 2: If $\overline{a}_t^{IB} = 0$ and $\overline{a}_t^{IG} > 0$, let $\check{a}_t^{IG} := \overline{a}_t^{IG} - \Delta^{(2)}, \check{a}_t^{IB} := \overline{a}_t^{IB}, \check{a}_t^W := \overline{a}_t^W - \Delta^{(2)}, \check{g}_t^I := \overline{g}_t^I - \Delta^{(2)}/\alpha$, and $\check{g}_t^S := \overline{g}_t^S + \Delta^{(2)}\beta$, where $\Delta^{(2)} \in (0, \min\{\overline{a}_t^{IG}, \overline{a}_t^W\}]$.

It is easy to verify that $(\check{a}_t^{IG}, \check{a}_t^{IB}, \check{a}_t^W, \check{g}_t^I, \check{g}_t^S)$ belongs to $\Psi^1(x_t, w_t)$ and yields a value of the objective function in (28) that in Case 1 is larger than and in Case 2 is equal to the one corresponding to $(\overline{a}_t^{IG}, \overline{a}_t^{IB}, \overline{a}_t^W, \overline{g}_t^I, \overline{g}_t^S)$, which contradicts the assumed optimality of $(\overline{a}_t^{IG}, \overline{a}_t^{IB}, \overline{a}_t^W, \overline{g}_t^I, \overline{g}_t^S)$. Therefore, the quantity $V_t^*(x_t, w_t, \vec{p}_t)$ equals the optimal objective function value of model (28).

We establish the claimed property using induction. This property holds trivially when $t = T$, given that $V_T^*(x_T, w_T, \vec{p}_T) \equiv 0$. Make the induction hypothesis that it is also true for each stage $k + 1$ through $T - 1$. Consider stage k. Fix $(w_k, \vec{p}_k) \in \mathcal{W} \times \mathcal{P}$. The set $\mathcal{C} := \{(x_k, a_k^{IG}, a_k^{IB}, a_k^W, g_k^I, g_k^S) | x_k \in \mathcal{X}, (a_k^{IG}, a_k^{IB}, a_k^W, g_k^I, g_k^S) \in \Psi^1(x_k, w_k)\}$ is convex, because $\mathcal{X} \equiv [0, C^S]$ is convex and $\Psi^1(x_k, w_k)$ is polyhedral and, hence, convex. The first two terms of the objective function of model (28) are linear in (g_k^S, a_k^W) and a_k^{IB}, respectively, and they do not depend on x_k, a_k^{IG}, and g_k^I. The induction hypothesis implies that $V_{k+1}^* (x_k + a_k^{IG} + a_k^{IB} - a_k^W, w_{k+1}, \vec{p}_{k+1})$ is jointly concave in $(x_k, a_k^{IG}, a_k^{IB}, a_k^W)$ for each given w_{k+1} and $\vec{p}_{k+1}$. It follows that $\delta \mathbb{E} \left[V_{k+1}^* \left(x_k + a_k^{IG} + a_k^{IB} - a_k^W, w_{k+1}, \vec{p}_{k+1} \right) | w_k, \vec{p}_k \right]$ is jointly concave in $(x_k, a_k^{IG}, a_k^{IB}, a_k^W)$. Further, this function is independent of g_k^I and g_k^S. Thus, the objective function of model (28) is jointly concave on $\mathcal{C}$. The set $\Psi^1(x_k, w_k)$ is nonempty for every $x_k \in \mathcal{X}$. We have $V_k^*(S_k) < \infty$ by Lemma 1. Theorem A.4 in Porteus (2002) implies that $V_k^*(S_k)$ is concave on $\mathcal{X}$. By the principle of mathematical induction, the claimed property holds for each stage 0 through T. $\blacksquare$

Appendix B: Proof of Proposition 2

We use Lemma 2 in the proof of Proposition 2.

LEMMA 2. *Suppose $P_t(\vec{p}_t) > 0$ for each $t \in \mathcal{T}$ and $\vec{p}_t \in \mathcal{P}$. For each $t \in \mathcal{T} \cup \{T\}$, $w_t \in \mathcal{W}$, and $\vec{p}_t \in \mathcal{P}$, the function $V_t^*(x_t, w_t, \vec{p}_t)$ is non-decreasing in x_t on $\mathcal{X}$.*

Proof: We use an induction argument. The claim holds trivially for $t = T$ as $V_T^*(x_T, w_T, \vec{p}_T) \equiv 0$. Make the induction hypothesis that the claimed property is also true for stages $k + 1$ through $T - 1$. Consider stage k. Fix $(w_k, \vec{p}_k) \in \mathcal{W} \times \mathcal{P}$. Pick $x_k^{(1)}$ and $x_k^{(2)} \in \mathcal{X}$ such that $x_k^{(1)} < x_k^{(2)}$. We show that for every action $(a_k^{(1)}, g_k^{(1)})$ in $\Psi(x_k^{(1)}, w_k)$, we can find an action $(a_k^{(2)}, g_k^{(2)})$ in $\Psi(x_k^{(2)}, w_k)$ such that the value of the objective function in (9) corresponding to $(x_k^{(2)}, w_k, \vec{p}_k, a_k^{(2)}, g_k^{(2)})$ is no lower than the one associated with $(x_k^{(1)}, w_k, \vec{p}_k, a_k^{(1)}, g_k^{(1)})$, from which it follows that $V_k^*(x_k^{(1)}, w_k, \vec{p}_k) \leq V_k^*(x_k^{(2)}, w_k, \vec{p}_k)$. We consider the following cases:

1) If $a_k^{(1)} < 0$, let $(a_k^{(2)}, g_k^{(2)})$ be equal to $(a_k^{(1)}, g_k^{(1)})$. This pair belongs to $\Psi(x_k^{(2)}, w_k)$ given that the solution $(a_k^{(1)}, g_k^{(1)})$ satisfies constraints (1)-(4) and (6), which do not depend on x_k, as well as (5), because $a_k^{(1)}$ is negative, and $x_k^{(2)}$ is no smaller than $x_k^{(1)}$. The equivalence $R(a_k^{(1)}, g_k^{(1)}, \vec{p}_k) \equiv R(a_k^{(2)}, g_k^{(2)}, \vec{p}_k)$, the inequality $x_k^{(2)} + a_k^{(2)} > x_k^{(1)} + a_k^{(1)}$, and the induction hypothesis applied to stage $k + 1$ imply that the evaluation of the objective function in (9) at $(x_k^{(2)}, w_k, \vec{p}_k, a_k^{(2)}, g_k^{(2)})$ is no lower than the one at $(x_k^{(1)}, w_k, \vec{p}_k, a_k^{(1)}, g_k^{(1)})$.

2) If $0 \le a_k^{(1)} \le x_k^{(2)} - x_k^{(1)}$, let $(a_k^{(2)}, g_k^{(2)})$ be equal to $(0, (g_k^{(1)} - a_k^{(1)}/\alpha)^+)$. Given $(a_k^{(1)}, g_k^{(1)}) \in \Psi(x_k^{(1)}, w_k)$, the pair $(a_k^{(2)}, g_k^{(2)})$ belongs to $\Psi(x_k^{(2)}, w_k)$: (1), (3), (5), and (6) hold given that $a_k^{(2)} = 0$; (4) is satisfied since $g_k^{(2)} = (g_k^{(1)} - a_k^{(1)}/\alpha)^+ \le g_k^{(1)} \le w_t$; (2) is true both if $g_k^{(2)} = (g_k^{(1)} - a_k^{(1)}/\alpha)^+ = 0$ and if $g_k^{(2)} = g_k^{(1)} - a_k^{(1)}/\alpha > 0$ (because $g_k^{(2)} - a_k^{(2)}/\alpha = g_k^{(2)} = g_k^{(1)} - a_k^{(1)}/\alpha \le C^T$, where the inequality is due to the fact $(a_k^{(1)}, g_k^{(1)}) \in \Psi(x_k^{(1)}, w_k)$). The inequalities $R(a_k^{(2)}, g_k^{(2)}, \vec{p}_k) \ge R(a_k^{(1)}, g_k^{(1)}, \vec{p}_k)$ and $x_k^{(2)} + a_k^{(2)} \equiv x_k^{(2)} \ge x_k^{(1)} + a_k^{(1)}$ and the induction hypothesis applied to stage $k+1$ imply that the evaluation of the objective function in (9) at $(x_k^{(2)}, w_k, \vec{p}_k, a_k^{(2)}, g_k^{(2)})$ is no lower than the one at $(x_k^{(1)}, w_k, \vec{p}_k, a_k^{(1)}, g_k^{(1)})$.

3) If $x_k^{(2)} - x_k^{(1)} < a_k^{(1)}$, we consider the following two possibilities:

- If $g_k^{(1)} > (x_k^{(2)} - x_k^{(1)})/\alpha$, let $(a_k^{(2)}, g_k^{(2)})$ be equal to $(a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}), g_k^{(1)} - (x_k^{(2)} - x_k^{(1)})/\alpha)$. Given $(a_k^{(1)}, g_k^{(1)}) \in \Psi(x_k^{(1)}, w_k)$, the pair $(a_k^{(2)}, g_k^{(2)})$ belongs to $\Psi(x_k^{(2)}, w_k)$: (1) and (2) are true because $g_k^{(2)} - a_k^{(2)}/\alpha = g_k^{(1)} - a_k^{(1)}/\alpha$; (3) holds since $a_k^{(2)} > 0$; (4) is satisfied given that $g_k^{(2)} = g_k^{(1)} - (x_k^{(2)} - x_k^{(1)})/\alpha \le g_k^{(1)}$; (5) is fulfilled as a result of $0 < a_k^{(2)} = a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) \le C^S - x_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) = C^S - x_k^{(2)}$; and (6) is valid by virtue of $0 < a_k^{(2)} < a_k^{(1)}$. The equivalences $R(a_k^{(1)}, g_k^{(1)}, \vec{p}_k) \equiv R(a_k^{(2)}, g_k^{(2)}, \vec{p}_k)$, which follows from the conditions $g_k^{(2)} - a_k^{(2)}/\alpha = g_k^{(1)} - a_k^{(1)}/\alpha$, $a_k^{(1)} > 0$, $a_k^{(2)} > 0$, and $x_k^{(2)} + a_k^{(2)} \equiv x_k^{(2)} + a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) \equiv x_k^{(1)} + a_k^{(1)}$ imply that the evaluation of the objective function in (9) at $(x_k^{(2)}, w_k, \vec{p}_k, a_k^{(2)}, g_k^{(2)})$ equals the one at $(x_k^{(1)}, w_k, \vec{p}_k, a_k^{(1)}, g_k^{(1)})$.

- If $g_k^{(1)} \le (x_k^{(2)} - x_k^{(1)})/\alpha$, let $(a_k^{(2)}, g_k^{(2)})$ be equal to $(a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}), 0)$. Given $(a_k^{(1)}, g_k^{(1)}) \in \Psi(x_k^{(1)}, w_k)$, the pair $(a_k^{(2)}, g_k^{(2)})$ belongs to $\Psi(x_k^{(2)}, w_k)$: (2) is satisfied since $a_k^{(2)} > 0$ and $g_k^{(2)} = 0$; (3) holds because $a_k^{(2)} > 0$; (4) is valid given that $g_k^{(2)} = 0$; (5) is true as a consequence of $0 < a_k^{(2)} = a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) \le C^S - x_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) = C^S - x_k^{(2)}$; (6) is fulfilled as a result of $0 < a_k^{(2)} < a_k^{(1)}$; and (1) stands by virtue of $(a_k^{(2)}/\alpha - g_k^{(2)})/\tau = a_k^{(2)}/\alpha/\tau = (a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}))/\alpha/\tau = a_k^{(1)}/\alpha/\tau - (x_k^{(2)} - x_k^{(1)})/\alpha/\tau \le a_k^{(1)}/\alpha/\tau - g_k^{(1)}/\tau \le C^T$ (where the last inequality follows from $(a_k^{(1)}, g_k^{(1)}) \in \Psi(x_k^{(1)}, w_k)$). The inequality $R(a_k^{(2)}, g_k^{(2)}, \vec{p}_k) \ge R(a_k^{(1)}, g_k^{(1)}, \vec{p}_k)$, which results from $(a_k^{(2)}/\alpha - g_k^{(2)})/\tau \le a_k^{(1)}/\alpha/\tau - g_k^{(1)}/\tau$, $a_k^{(1)} > 0$, $a_k^{(2)} > 0$, and the equivalence $x_k^{(2)} + a_k^{(2)} \equiv x_k^{(2)} + a_k^{(1)} - (x_k^{(2)} - x_k^{(1)}) \equiv x_k^{(1)} + a_k^{(1)}$ imply that the evaluation of the objective function in (9) at $(x_k^{(2)}, w_k, \vec{p}_k, a_k^{(2)}, g_k^{(2)})$ is no lower than the one at $(x_k^{(1)}, w_k, \vec{p}_k, a_k^{(1)}, g_k^{(1)})$.

By the principle of mathematical induction, the claimed property holds for each stage $t \in \mathcal{T} \cup \{T\}$. $\square$

Part (i). Fix $(w_t, \vec{p}_t) \in \mathcal{W} \times \mathcal{P}$. If $\tau = 1$ then $X_t^{(1)} = X_t^{(2)}$, because the objective functions in (10) and (11) are identical. Similarly, if $\alpha\beta\tau = 1$ then $X_t^{(2)} = X_t^{(3)}$. Suppose $\tau \ne 1$ and $\alpha\beta \ne 1$. The inequalities

$$-P_t(\vec{p}_t)X_t^{(1)}/(\alpha\tau) + U_t^*(X_t^{(1)}) \ge -X_t^{(2)}P_t(\vec{p}_t)/(\alpha\tau) + U_t^*(X_t^{(2)}),$$
$$-P_t(\vec{p}_t)X_t^{(2)}\tau/\alpha + U_t^*(X_t^{(2)}) \ge -X_t^{(1)}P_t(\vec{p}_t)\tau/\alpha + U_t^*(X_t^{(1)}),$$

where $U_t^*(\cdot)$ is used in lieu of $U_t^*(\cdot, w_t, \vec{p}_t)$, hold given the definitions of $X_t^{(1)}$ and $X_t^{(2)}$, respectively. It follows that

$$P_t(\vec{p}_t)[\tau/\alpha - 1/(\alpha\tau)]X_t^{(1)} \ge P_t(\vec{p}_t)[\tau/\alpha - 1/(\alpha\tau)]X_t^{(2)},$$

which implies $X_t^{(1)} \le X_t^{(2)}$ because $P_t(\vec{p}_t)[\tau/\alpha - 1/(\alpha\tau)] < 0$. The inequalities

$$-P_t(\vec{p}_t)\tau X_t^{(2)}/\alpha + U_t^*(X_t^{(2)}) \ge -P_t(\vec{p}_t)\tau X_t^{(3)}/\alpha + U_t^*(X_t^{(3)}),$$

$$-P_t(\vec{p}_t)\beta\tau X_t^{(3)} + U_t^*(X_t^{(3)}) \geq -P_t(\vec{p}_t)\beta\tau X_t^{(2)} + U_t^*(X_t^{(2)}),$$

which are true due to the definitions of $X_t^{(2)}$ and $X_t^{(3)}$, respectively, lead to

$$P_t(\vec{p}_t)\left[\beta\tau - \tau/\alpha\right] X_t^{(2)} \geq P_t(\vec{p}_t)\left[\beta\tau - \tau/\alpha\right] X_t^{(3)},$$

which yields $X_t^{(2)} \leq X_t^{(3)}$ since $P_t(\vec{p}_t)\left[\beta\tau - \tau/\alpha\right] < 0$.

Part (ii). Fix $(w_t, \vec{p}_t) \in \mathcal{W} \times \mathcal{P}$. We denote the maximal quantity that one can generate by $\bar{g}_t := \min\{w_t, C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$. We establish that $g_t^*(S_t) = \bar{g}_t$ by showing that for every feasible action (a_t, g_t) with $g_t < \bar{g}_t$ we can find a feasible pair (a_t', g_t') with $g_t' > g_t$ that yields a value for the objective function in (9) no smaller than the one associated with (a_t, g_t). We define

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

which expresses the negative of the amount bought in the first case and the quantity sold in the other two cases associated with the feasible action (a_t, g_t). We consider the following cases:

1) If $q(a_t, g_t) < 0$, we define $\Delta := \min\{w_t - g_t, -q(a_t, g_t)\}$ and $(a_t', g_t') := (a_t, g_t + \Delta)$. This action pair is feasible. Because it buys less than (a_t, g_t) does, the valuation at (a_t', g_t') of the objective function in (9) is larger than the one at (a_t, g_t).

2) If $0 \leq q(a_t, g_t) < C^{\mathrm{T}}$, we define $\Delta := \min\{w_t - g_t, C^{\mathrm{T}} - q(a_t, g_t)\}$ and $(a_t', g_t') := (a_t, g_t + \Delta)$. This action pair is feasible. Given that it sells more than (a_t, g_t) does, it results in a value of the objective function in (9) that exceeds the one corresponding to (a_t, g_t).

3) If $q(a_t, g_t) = C^{\mathrm{T}}$, we consider the following cases:

   - If $a_t \geq 0$, we define $\Delta := \min\{w_t - g_t, \min\{C^{\mathrm{S}} - x_t - a_t, C^{\mathrm{C}} - a_t\}/\alpha\}$, $a_t' := a_t + \alpha\Delta$, and $g_t' := g_t + \Delta$.
   - If $a_t < 0$, we define $\Delta := \min\{w_t - g_t, -\beta a_t\}$, $a_t' := a_t + \Delta/\beta$, and $g_t' := g_t + \Delta$.

   In both cases, the action (a_t', g_t') is feasible. Further, it satisfies both $q(a_t', g_t') = q(a_t, g_t)$ and $a_t' > a_t$, so by Lemma 2 it results in a larger value of the objective function in (9) than (a_t, g_t) does.

Part (iii). Fix $\vec{p}_t \in \mathcal{P}$. We characterize $a_t^*(S_t)$ for each of the regions Γ^0, Γ^1, and Γ^2, which partition $\mathcal{X} \times \mathcal{W}$.

Region Γ^0. Recall that $\Gamma^0 \equiv \{(x_t, w_t) \in \mathcal{X} \times \mathcal{W} : w_t \geq C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$. The quantity $g_t^*(S_t) \equiv \min\{w_t, C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha\}$ simplifies to $C^{\mathrm{T}} + \min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}/\alpha$ and the only feasible inventory change action is to store the amount by which the optimal generation exceeds the transmission capacity; that is, $a_t^*(S_t)$ equals $\min\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\}$.

To perform the analysis for regions Γ^1 and Γ^2, for each pair $(x_t, w_t) \in \mathcal{X} \times \mathcal{W}$ we include in set $\mathcal{Y}(x_t, w_t)$ the feasible inventory levels y_t's that satisfy (1)-(3) and (5)-(6) with a_t and g_t replaced by $y_t - x_t$ and $g_t^*(S_t)$, respectively (constraint (4) is redundant). We equivalently express the optimization model in (9) as

$$\max_{y_t \in \mathcal{Y}(x_t, w_t)} \{R(y_t - x_t, g_t^*(S_t), \vec{p}_t) + U_t^*(y_t, w_t, \vec{p}_t)\}. \tag{29}$$

We denote an optimal solution to this model as $y_t^*(S_t)$. In the ensuing analysis we characterize this quantity. The difference $y_t^*(S_t) - x_t$ gives $a_t^*(S_t)$.

Region Γ^1. Because $\Gamma^1 \equiv \left\{ (x_t, w_t) \in \mathcal{X} \times \mathcal{W} : C^{\mathrm{T}} \leq w_t < C^{\mathrm{T}} + \min\left\{ C^{\mathrm{S}} - x_t, C^{\mathrm{C}} \right\} / \alpha \right\}$, the quantity $g_t^*(S_t)$ simplifies from $\min\left\{ w_t, C^{\mathrm{T}} + \min\left\{ C^{\mathrm{S}} - x_t, C^{\mathrm{C}} \right\} / \alpha \right\}$ to w_t and at least the excess generation $w_t - C^{\mathrm{T}}$ must be injected into storage, so that selling inventory is impossible; that is, the inventory level x_t increases by at least $\alpha\left(w_t - C^{\mathrm{T}} \right)$. We consider the model

$$\max_{y_t \in \left[x_t, \min\left\{ x_t + g_t^*(S_t), C^{\mathrm{S}} \right\} \right]} \left\{ R\left(y_t - x_t, g_t^*(S_t), \vec{p}_t \right) + U_t^*\left(y_t, w_t, \vec{p}_t \right) \right\},$$

which ignores both the ability to purchase or sell and the constraints posed by the transmission and charging capacity limits when the generation action in stage t and state S_t equals $g_t^*(S_t)$. We denote by $y_t^{\diamond, \Gamma^1}(S_t)$ an optimal solution to this model. Proposition 1 and arguments analogous to the ones that lead to Theorem 1 in Secomandi (2010) imply

$$y_t^{\diamond, \Gamma^1}(S_t) = \begin{cases} \min\left\{ x_t + \alpha w_t, X_t^{(2)} \right\} & \text{if } x_t \in \left[0, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}} \right) \right] \text{ and } \alpha\left(w_t - C^{\mathrm{T}} \right) \leq X_t^{(2)}, \\ x_t + \alpha\left(w_t - C^{\mathrm{T}} \right) & \text{otherwise.} \end{cases}$$

It follows from Proposition 1 that if we ignore the buying decision then an optimal solution to model (29) corresponds to x_t modified to attain an inventory level that is as close as possible to $y_t^{\diamond, \Gamma^1}(S_t)$ when considering the constraints associated with the transmission and charging capacity limits. We express this solution as $\overline{y}_t^{*, \Gamma^1}(S_t)$. Further taking into account the ability to purchase allows us to obtain $y_t^*(S_t)$ from $\overline{y}_t^{*, \Gamma^1}(S_t)$. In particular, if energy can be and is purchased then the optimally generated amount, w_t, must be used to increase the inventory level by αw_t. Consequently, we consider the following exhaustive cases:

- If $\alpha w_t < C^{\mathrm{C}}$ then buying electricity is feasible. We examine the following situations:

  — $\alpha w_t < X_t^{(1)}$, which implies $\alpha\left(w_t - C^{\mathrm{T}} \right) < X_t^{(2)}$:

  * If $x_t \in \left[0, X_t^{(1)} - \alpha w_t \right]$ then $\overline{y}_t^{*, \Gamma^1}(S_t)$ equals $x_t + \alpha w_t$. Proposition 1 implies that optimizing the purchase decision at this inventory level given the residual charging capacity results in raising this level to one that is as close as possible to $X_t^{(1)}$; that is, $y_t^*(S_t)$ is as near as feasible to $X_t^{(1)}$. The action $a_t^*(S_t)$ is thus

$$\min\left\{ X_t^{(1)} - x_t, \alpha\left(\tau C^{\mathrm{T}} + w_t \right), C^{\mathrm{C}} \right\}.$$

  * If $x_t \in \left(X_t^{(1)} - \alpha w_t, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}} \right) \right]$ then $\overline{y}_t^{*, \Gamma^1}(S_t)$ equals $x_t + \alpha w_t$ and it follows from Proposition 1 that further considering the ability to buy does not add any benefit. That is, $y_t^*(S_t)$ coincides with $\overline{y}_t^{*, \Gamma^1}(S_t)$ and $a_t^*(S_t)$ is

$$\min\left\{ X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}} \right\} = \min\left\{ X_t^{(2)} - x_t, \alpha w_t \right\}.$$

  * If $x_t \in \left(X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}} \right), C^{\mathrm{S}} \right]$ then $\overline{y}_t^{*, \Gamma^1}(S_t)$ is equal to $x_t + \alpha\left(w_t - C^{\mathrm{T}} \right)$. By Proposition 1, $y_t^*(S_t)$ is identical to $\overline{y}_t^{*, \Gamma^1}(S_t)$ so that $a_t^*(S_t)$ is $\alpha\left(w_t - C^{\mathrm{T}} \right)$.

  — $\alpha w_t \geq X_t^{(1)}$: Purchasing electricity does not add any value so that $y_t^*(S_t) = \overline{y}_t^{*, \Gamma^1}(S_t)$. The cases to consider are

  * If $\alpha\left(w_t - C^{\mathrm{T}} \right) < X_t^{(2)}$ and
    · $x_t \in \left[0, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}} \right) \right]$ then $a_t^*(S_t)$ equals

$$\min\left\{ X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}} \right\} = \min\left\{ X_t^{(2)} - x_t, \alpha w_t \right\};$$

$\cdot$ $x_t \in \left(X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right), C^{\mathrm{S}}\right]$ then $a_t^*(S_t)$ equals $\alpha\left(w_t - C^{\mathrm{T}}\right)$.

$*$ If $\alpha\left(w_t - C^{\mathrm{T}}\right) \geq X_t^{(2)}$ then $a_t^*(S_t)$ equals $\alpha\left(w_t - C^{\mathrm{T}}\right)$.

- If $\alpha w_t \geq C^{\mathrm{C}}$ then buying electricity is infeasible and, hence, $y_t^*(S_t) = \overline{y}_t^{*,\Gamma_1}(S_t)$. The only situations to analyze are

$-\alpha\left(w_t - C^{\mathrm{T}}\right) < X_t^{(2)}$:

$*$ If $x_t \in \left[0, X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right)\right]$ then $a_t^*(S_t)$ equals

$$\min\left\{X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}}\right\} = \min\left\{X_t^{(2)} - x_t, \alpha w_t\right\}.$$

$*$ If $x_t \in \left(X_t^{(2)} - \alpha\left(w_t - C^{\mathrm{T}}\right), C^{\mathrm{S}}\right]$ then $a_t^*(S_t)$ equals $\alpha\left(w_t - C^{\mathrm{T}}\right)$.

$-\alpha\left(w_t - C^{\mathrm{T}}\right) \geq X_t^{(2)}$: $a_t^*(S_t)$ equals $\alpha\left(w_t - C^{\mathrm{T}}\right)$.

Region Γ^2. Given $\Gamma^2 \equiv \left\{(x_t, w_t) \in \mathcal{X} \times \mathcal{W} : 0 \leq w_t < C^{\mathrm{T}}\right\}$, the quantity $g_t^*(S_t)$ reduces from $\min\left\{w_t, C^{\mathrm{T}} + \min\left\{C^{\mathrm{S}} - x_t, C^{\mathrm{C}}\right\}/\alpha\right\}$ to w_t. We examine the model

$$\max_{y_t \in \left[0, \min\left\{x_t + g_t^*(S_t), C^{\mathrm{S}}\right\}\right]} \left\{R\left(y_t - x_t, g_t^*(S_t), \vec{p}_t\right) + U_t^*\left(y_t, w_t, \vec{p}_t\right)\right\},$$

which neglects both the decision to buy and the transmission, charging, and discharging capacity constraints when the generation action in stage t and state S_t is $g_t^*(S_t)$. We let $y_t^{\diamond,\Gamma^2}(S_t)$ be an optimal solution to this model. Applying Proposition 1 and proceeding in a manner analogous to the proof of Theorem 1 in Secomandi (2010) yield

$$y_t^{\diamond,\Gamma^2}(S_t) = \begin{cases} \min\left\{x_t + \alpha w_t, X_t^{(2)}\right\} & \text{if } x_t \in \left[0, X_t^{(2)}\right], \\ x_t & \text{if } x_t \in \left(X_t^{(2)}, X_t^{(3)}\right], \\ X_t^{(3)} & \text{if } x_t \in \left(X_t^{(3)}, C^{\mathrm{S}}\right]. \end{cases}$$

Proposition 1 implies that if we omit the purchase decision then an optimal solution to model (29), which we denote as $\overline{y}_t^{*,\Gamma_2}(S_t)$, is x_t altered to reach an inventory level that is as near as feasible to $y_t^{\diamond,\Gamma^2}(S_t)$ when taking into account the transmission, charging, and discharging capacity constraints. Given $\overline{y}_t^{*,\Gamma_2}(S_t)$, factoring the decision to buy into the analysis leads us to find $y_t^*(S_t)$. Particularly, if energy can be and is acquired from the market then the optimal generation, w_t, must be injected into storage, which increases the inventory level by αw_t. We thus examine the following comprehensive cases:

- If $\alpha w_t < C^{\mathrm{C}}$ buying electricity is possible. A logic similar to the one used in the case $\alpha w_t < C^{\mathrm{C}}$ for the region Γ^1 implies that purchasing is valuable if and only if $x_t \in \left[0, X_t^{(1)} - \alpha w_t\right]$ when $\alpha w_t < X_t^{(1)}$, in which case $y_t^*(S_t) > \overline{y}_t^{*,\Gamma_2}(S_t)$, and does not add any value when $\alpha w_t \geq X_t^{(1)}$, in which case $y_t^*(S_t) = \overline{y}_t^{*,\Gamma_2}(S_t)$. Hence, we consider the following possibilities:

$-\alpha w_t < X_t^{(1)}$:

$*$ If $x_t \in \left[0, X_t^{(1)} - \alpha w_t\right]$ then $a_t^*(S_t)$ equals $\min\left\{X_t^{(1)} - x_t, \alpha\left(\tau C^{\mathrm{T}} + w_t\right), C^{\mathrm{C}}\right\}$.

$*$ If $x_t \in \left(X_t^{(1)} - \alpha w_t, X_t^{(2)}\right]$ then $a_t^*(S_t)$ equals

$$\min\left\{X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}}\right\} = \min\left\{X_t^{(2)} - x_t, \alpha w_t\right\}.$$

$*$ If $x_t \in \left(X_t^{(2)}, X_t^{(3)}\right]$ then $a_t^*(S_t)$ equals 0.

 * If $x_t \in \left(X_t^{(3)}, C^{\mathrm{S}} \right]$ then $a_t^*(S_t)$ equals $\max \left\{ X_t^{(3)} - x_t, \left(w_t - C^{\mathrm{T}} \right) / \beta, -C^{\mathrm{D}} \right\}$.

— $\alpha w_t \geq X_t^{(1)}$:

 * If $x_t \in \left[0, X_t^{(2)} \right]$ then $a_t^*(S_t)$ equals

$$\min \left\{ X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}} \right\} = \min \left\{ X_t^{(2)} - x_t, \alpha w_t \right\}.$$

 * If $x_t \in \left(X_t^{(2)}, X_t^{(3)} \right]$ then $a_t^*(S_t)$ equals 0.
 * If $x_t \in \left(X_t^{(3)}, C^{\mathrm{S}} \right]$ then $a_t^*(S_t)$ equals $\max \left\{ X_t^{(3)} - x_t, \left(w_t - C^{\mathrm{T}} \right) / \beta, -C^{\mathrm{D}} \right\}$.

- If $\alpha w_t \geq C^{\mathrm{C}}$ buying electricity is not possible so that $y_t^*(S_t) = \overline{y}_t^{*,\Gamma_2}(S_t)$. Hence, we analyze the following

cases:

— If $x_t \in \left[0, X_t^{(2)} \right]$ then $a_t^*(S_t)$ equals

$$\min \left\{ X_t^{(2)} - x_t, \alpha w_t, C^{\mathrm{C}} \right\} = \min \left\{ X_t^{(2)} - x_t, \alpha w_t \right\}.$$

— If $x_t \in \left(X_t^{(2)}, X_t^{(3)} \right]$ then $a_t^*(S_t)$ equals 0.
— If $x_t \in \left(X_t^{(3)}, C^{\mathrm{S}} \right]$ then $a_t^*(S_t)$ equals $\max \left\{ X_t^{(3)} - x_t, \left(w_t - C^{\mathrm{T}} \right) / \beta, -C^{\mathrm{D}} \right\}$.

Appendix C: Proof of Proposition 3

If $P_t(\vec{p}_t) > 0$ then the inequalities $X_t^{(1),\mathrm{H1}} \leq X_t^{(2),\mathrm{H1}} \leq X_t^{(3),\mathrm{H1}}$ can be established using the proof of part (i) of Proposition 2 with $X_t^{(\nu)}$ replaced by $X_t^{(\nu),\mathrm{H1}}$ for $\nu \in \{1,2,3\}$. The inequality $X_t^{(3),\mathrm{H1}} \leq X_t^{(4),\mathrm{H1}}$ follows from the inequalities

$$-P_t(\vec{p}_t)\beta\tau X_t^{(3),\mathrm{H1}} + U_t^{\mathrm{H1}}(X_t^{(3),\mathrm{H1}}) \geq -P_t(\vec{p}_t)\beta\tau X_t^{(4),\mathrm{H1}} + U_t^{\mathrm{H1}}(X_t^{(4),\mathrm{H1}}),$$
$$U_t^{\mathrm{H1}}(X_t^{(4),\mathrm{H1}}) \geq U_t^{\mathrm{H1}}(X_t^{(3),\mathrm{H1}}),$$

which are implied by the definitions of $X_t^{(3),\mathrm{H1}}$ and $X_t^{(4),\mathrm{H1}}$, respectively.

If $P_t(\vec{p}_t) = 0$ then $X_t^{(1),\mathrm{H1}} = X_t^{(2),\mathrm{H1}} = X_t^{(3),\mathrm{H1}} = X_t^{(4),\mathrm{H1}}$ because the objective functions are identical in (13), (14), (15), and (16).

If $P_t(\vec{p}_t) < 0$ then the validity of the claimed property can be shown using an analogous logic.

Appendix D: Suboptimality of the H1 policy

We show by example that the H1 policy can fail to be optimal. For simplicity, we assume the storage energy capacity (space) is one ($C^{\mathrm{S}} = 1$); the charging and discharging (power) capacities are both equal to one (i.e., they are never constraining); the generation capacity is one ($C^{\mathrm{G}} = 1$); the transmission capacity is 0.4 ($C^{\mathrm{T}} = 0.4$); and there are no charging and transmission losses ($\alpha = \tau = 1$) but there is a 50% discharging loss ($\beta = 0.5$). The time horizon consists of four periods ($T = 4$). The price in the first period equals -0.3. The prices in the last three periods follow two paths, labeled as A and B, respectively: $(4, 3, 8)$ with probability 47% and $(-3.6, -5, -2)$ with probability 53%. The available wind energy is deterministic in each period and its corresponding path for the four periods is $(0, 0.3, 0.5, 0)$. The discount factor is one ($\delta = 1$).

We compute optimal value functions for stages 1 through 3 for the two price paths. We include the price path labels A and B as superscripts in our optimal value function notation, also omitting the exogenous

Figure 7 **The optimal and H1 stage 0 inventory targets for the example in Appendix D**

state components, i.e., the available wind energy and price, from this and related notation. These functions are

$$V_3^{*,\mathrm{A}}(x_3) = \begin{cases} 4x_3 & \text{if } x_3 \in [0,0.8) \\ 3.2 & \text{if } x_3 \in [0.8,1] \end{cases}, \qquad V_3^{*,\mathrm{B}}(x_3) = \begin{cases} 0.8 & \text{if } x_3 \in [0,0.6) \\ 2 - 2x_3 & \text{if } x_3 \in [0.6,1] \end{cases},$$

$$V_2^{*,\mathrm{A}}(x_2) = \begin{cases} 2.3 + 3x_2 & \text{if } x_2 \in [0,0.7) \\ 4.4 & \text{if } x_2 \in [0.7,1] \end{cases}, \qquad V_2^{*,\mathrm{B}}(x_2) = \begin{cases} 2.8 & \text{if } x_2 \in [0,0.2) \\ 3.2 - 2x_2 & \text{if } x_2 \in [0.2,0.6) \\ 5(1 - x_2) & \text{if } x_2 \in [0.6,1] \end{cases},$$

$$V_1^{*,\mathrm{A}}(x_1) = \begin{cases} 3.5 + 3x_1 & \text{if } x_1 \in [0,0.7) \\ 4.2 + 2x_1 & \text{if } x_1 \in [0.7,0.9) \\ 6 & \text{if } x_1 \in [0.9,1] \end{cases}, \qquad V_1^{*,\mathrm{B}}(x_1) = \begin{cases} 3.84 - 2x_1 & \text{if } x_1 \in [0,0.2) \\ 4.16 - 3.6x_1 & \text{if } x_1 \in [0.2,5/9) \\ 3.16 - 1.8x_1 & \text{if } x_1 \in [5/9,1] \end{cases}.$$

Thus, the optimal stage 0 continuation function is

$$U_0^*(y_0) = 0.47 V_1^{*,\mathrm{A}}(y_0) + 0.53 V_1^{*,\mathrm{B}}(y_0) = \begin{cases} 3.6802 + 0.35 y_0 & \text{if } y_0 \in [0,0.2) \\ 3.8498 - 0.498 y_0 & \text{if } y_0 \in [0.2,5/9) \\ 3.3198 + 0.456 y_0 & \text{if } y_0 \in [5/9,0.7) \\ 3.6488 - 0.014 y_0 & \text{if } y_0 \in [0.7,0.9) \\ 4.4948 - 0.954 y_0 & \text{if } y_0 \in [0.9,1) \end{cases}.$$

Given the initial inventory level x_0, the optimal stage 0 inventory change action solves

$$\max\{-(-0.3)0.5a_0\mathbf{1}(a_0 \le 0) - (-0.3)a_0\mathbf{1}(a_t > 0) + U_0^*(x_0 + a_0)\} \tag{30}$$

$$\text{s.t. } (-0.5a_0 - 0.4)\mathbf{1}(a_0 \le 0) \le 0, (a_0 - 0.4)\mathbf{1}(a_0 > 0) \le 0, -x_0 \le a_0 \le 1 - x_0.$$

The left panel of Figure 7 shows the dependence on x_0 of the resulting inventory target, i.e., the ending inventory level ignoring the transmission capacity constraint: 0.2 for $x_0 \in [0, 79/303) \cup [0.84, 1]$ and 0.9 for $x_0 \in [79/303, 0.84]$.

The H1 continuation function in stage 0 coincides with the optimal one, i.e., $U_0^{\mathrm{H1}}(y_0) \equiv U_0^*(y_0)$. We use it to compute the following thresholds ($X_0^{(2),\mathrm{H1}}$ is not relevant as $P_0 < 0$):

$$X_0^{(1),\mathrm{H1}} = \underset{y_0 \in [0,1]}{\arg\max} \left\{ U_0^{\mathrm{H1}}(y_0) - (-0.3)y_0 \right\} = 0.9,$$

$$X_0^{(3),\mathrm{H1}} = \underset{y_0 \in [0,1]}{\arg\max} \left\{ U_0^{\mathrm{H1}}(y_0) - (-0.3)0.5y_0 \right\} = 0.2,$$

$$X_0^{(4),\text{H1}} = \arg\max_{y_0 \in [0,1]} \left\{ U_0^{\text{H1}}(y_0) \right\} = 0.2.$$

The quantity Z_0 is the largest inventory level x_0 in $[0.2, 0.9]$ such that the two action pairs $(\min\{0.9 - x_0, 0.4, 1\}, 0) = (\min\{0.9 - x_0, 0.4\}, 0)$ and $(\max\{0.2 - x_0, -0.8, -1\}, 0) = (0.2 - x_0, 0)$ give the same evaluation of the right hand side of (8) for $\pi = \text{H1}$. Depending on whether a candidate value z_0 for Z_0 satisfies $0.9 - z_0 < 0.4$ or $0.9 - z_0 \geq 0.4$, z_0 solves

$$-0.5(-0.3)(0.2 - z_0)\mathbf{1}(0.2 - z_0 \leq 0) + U_0^{\text{H1}}(0.2) = -(-0.3)(0.9 - z_0)\mathbf{1}(0 < 0.9 - z_0 < 0.4) + U_0^{\text{H1}}(0.9)$$

or

$$-0.5(-0.3)(0.2 - z_0)\mathbf{1}(0.2 - z_0 \leq 0) + U_0^{\text{H1}}(0.2) = -(-0.3)(0.4)\mathbf{1}(0.9 - z_0 \geq 0.4) + U_0^{\text{H1}}(z_0 + 0.4).$$

The solutions to these equalities are 0.84 and $79/303$, respectively, so that Z_0 equals 0.84. The right panel of Figure 7 displays the H1 stage 0 inventory target: 0.9 for $x_0 \in [0, 0.84)$ and 0.2 for $x_0 \in [0.84, 1]$.

The optimal and H1 stage 0 inventory targets differ if x_0 belongs to $[0, 79/303)$. The largest corresponding H1 policy optimality gap is 1.87%. The H1 policy would be optimal if the objective function in (30) were jointly concave in (x_0, a_0), whereas it is not.

References

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

Secomandi, N. 2010. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.