

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection Lee Kong Chian School Of
Business

Lee Kong Chian School of Business

3-2016

Electricity Trading and Negative Prices: Storage vs. Disposal

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 [Operations and Supply Chain Management Commons](#)

Citation

ZHOU, Yangfang (Helen); Scheller-Wolf, Alan; Secomandi, Nicola; and Smith, Stephen. Electricity Trading and Negative Prices: Storage vs. Disposal. (2016). *Management Science*. 62, (3), 880-898.

Available at: https://ink.library.smu.edu.sg/lkcsb_research/3483

This Journal Article is brought to you for free and open access by the Lee Kong Chian School of Business at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Lee Kong Chian School Of Business by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Electricity Trading and Negative Prices: Storage vs. Disposal

Yangfang 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

Tepper Working Paper 2012-E1

First Version: Jan 2012; Revised: April 2013; June 2014; December 2014

Electricity cannot yet be stored on a large scale, but technological advances leading to cheaper and more efficient industrial batteries make grid-level storage of electricity surpluses a natural choice. Because electricity prices can be negative, it is unclear how the presence of negative prices might affect the storage policy structure known to be optimal when prices are only non-negative, or even how important it is to consider negative prices when managing an industrial battery. For fast storage (a storage facility that can both be fully emptied and filled up in one decision period), we show analytically that negative prices can substantially alter the optimal storage policy structure, e.g., all else being equal, it can be optimal to empty an almost empty storage facility and fill up an almost full one. For more typical slow grid-level electricity storage, we numerically establish that ignoring negative prices could result in a considerable loss of value when negative prices occur more than 5% of the time. Negative prices raise another possibility: rather than storing surpluses, a merchant might buy negatively priced electricity surpluses and dispose of them, e.g., using load banks. We find that the value of such disposal strategy is substantial, e.g., about 118 \$/kW-year when negative prices occur 10% of the time, but smaller than that of the storage strategy, e.g., about 391 \$/kW-year using a typical battery. However, devices for disposal are much cheaper than those for storage. Our results thus have ramifications for merchants as well as policy makers.

Key words: inventory; electricity storage; electricity disposal; Markov decision process; asset pricing models; negative prices

1. Introduction

In a commodity market, surpluses occur when supply outstrips demand. Because electricity supply and demand must be matched in real time, dealing with electricity surpluses is particularly critical. As storing surpluses for future resale is the most common strategy for commodities (Williams and Wright 1991), it is also a natural one for electricity merchants who trade electricity in a market. So, even though electricity storage has not yet been deployed on a large scale (only around 2.3% of

electricity consumed in the U.S. currently is satisfied from storage; Gyuk et al. 2013, p. 4 and EIA 2012), several studies have recognized the potential of this strategy (EPRI 2004, Eyer and Corey 2010, The Economist 2012, Akhil et al. 2013).

In contrast to the prices of other commodities, electricity prices can be negative. Negative prices have been observed in electricity markets both in the U.S. (Huntowski et al. 2012), including the New York Independent System Operator (NYISO), PJM, the California ISO, and the Electric Reliability Council of Texas (ERCOT), and in other countries, such as the Nordic Power Exchange (Sewalt and de Jong 2003), the European Energy Exchange (Genoese et al. 2010, Nicolosi 2010, Brandstätt et al. 2011), and the Australian Energy Market Operator (AEMO 2009). In particular, between 2008 and 2011 negative prices occurred in ERCOT around 10% of the time (Huntowski et al. 2012).

Negative prices can be caused by a mix of factors: (i) technological limits on adjusting the generation levels of coal or nuclear power plants, or the high costs of these adjustments, can lead the managers of such plants to pay others to purchase their excess power when electricity demand is low, e.g., at night (Sewalt and de Jong 2003, Knittel and Roberts 2005, Genoese et al. 2010, Nicolosi 2010, Brandstätt et al. 2011, Brown 2012). (ii) Lack of transmission capacity can cause excess local electricity supply, e.g., in ERCOT, where wind energy in the western zone cannot be transmitted outside of this region (Brown 2012). (iii) The Production Tax Credit received by wind-based electricity generators in the U.S., currently valued at 23 \$/MWh (megawatt-hour) (DSIRE 2014), can induce wind generators to bid a negative price and still generate a positive revenue from a sale (Fink et al. 2009, Brown 2012, Huntowski et al. 2012). (iv) Short-term supply gluts can arise from policies that prioritize wind-based electricity generation, such as the one in Germany, which disallow curtailing it except for reliability reasons (Genoese et al. 2010, Nicolosi 2010, Brandstätt et al. 2011). It is widely believed that with the increasing use of wind power, negative prices will likely become more frequent and larger in magnitude in the future (Genoese et al. 2010, Nicolosi 2010, Brandstätt et al. 2011).

It is unclear whether the presence of negative prices may alter the optimal threshold structure of the finite-horizon merchant storage policy known to be optimal for the case when prices can only be non-negative, or even how important it is to consider negative prices in such a storage policy. We study these questions by modeling the problem of managing electricity storage with potentially negative prices as a Markov decision process (MDP).

For the case of fast storage (the case when both filling up an empty storage facility and emptying such a full facility take one decision period), we show analytically that negative prices can

make the optimal storage policy structure significantly different from those available in the literature (Charnes et al. 1966, Rempala 1994, Secomandi 2010) when prices are always non-negative. Analogous to these known threshold structures, we show that for every stage and state, the initial inventory set can be divided into three regions based on the type of optimal action: one region in which it is optimal to empty the facility, one in which it is optimal to fill up the facility, and one in which it is optimal to do nothing. But in contrast to these typical structures, our three regions can be ordered in different sequences, and thus (i) the optimal next-stage inventory level can fail to increase in the current-stage inventory level; i.e., a high current-stage inventory level can result in a low next-stage inventory level (and vice versa); and (ii) the optimal action can bring the current-stage inventory level *farther* away from a “target” band delineated by two thresholds if this inventory level is outside of this band. Our optimal policy structure subsumes the optimal policy structure of Charnes et al. (1966) for managing fast storage, i.e., when prices can only be non-negative, our policy simplifies to theirs.

For the case of slow storage—such as industrial batteries, which are more typical for grid-level electricity storage applications (EPRI 2004, Eyer and Corey 2010, Akhil et al. 2013, Gyuk et al. 2013)—we numerically quantify the importance of considering negative prices using an existing electricity price model calibrated to NYISO historical price data by Zhou et al. (2014). We find that ignoring negative prices when determining a storage policy could result in a considerable loss of value when negative prices occur more than 5% of the time, as they already do in some markets, e.g., ERCOT in 2008-2011 (Huntowski et al. 2012).

The presence of negative prices enables a merchant strategy different from storing electricity surpluses in an efficient battery: a merchant might buy negatively priced electricity surpluses and dispose of them, e.g., using specialized load banks, which are designed to mimic real load applied to power sources and can be used to consume electricity for general purposes (Emerson 2014). As Huntowski et al. (2012) point out, “negative prices could incent developers to build high electricity consuming elements to use negative-price hours in wind-rich regions, for the sole purpose of being paid to waste electricity.” It is unclear how valuable this “wasteful” strategy is, or how its value might compare with that of the storage strategy. Using the same calibrated price model as in our analysis of the storage strategy, we demonstrate that when negative prices occur 10% of the time the value of the disposal strategy is substantial—around 118 \$/kW-year—but less than that of the storage strategy, e.g., around 391 \$/kW-year using a battery with a typical 80% round-trip efficiency (the ratio of electricity withdrawn to that injected; Eyer and Corey 2010). However, because devices for disposing electricity (e.g., load banks) are much cheaper than efficient batteries

(Emerson 2014, Coffman Electrical Equipment Co 2014, Akhil et al. 2013), a merchant might find the disposal strategy more profitable than the storage strategy. Our findings also highlight for policy makers the relative values of the storage and disposal strategies, suggesting the need for additional research on assessing the potential impact of these strategies on social welfare.

We review the extant literature in §2. We model the storage strategy and derive the fast-storage optimal policy structure in §3. We model the disposal strategy in §4. Our numerical analysis is carried out in §5. We conclude in §6. Appendix A includes the proofs of the results presented in §3.2.

2. Literature review

Our work is related to the commodity and energy storage literature. A classical problem studied in this literature is the warehouse problem introduced by Cahn (1948): given a warehouse with limited space, what is the optimal inventory trading policy under seasonal (deterministic) variability in the commodity price? Dreyfus (1957) shows that if the commodity price is positive the optimal inventory trading decisions for a given time and price are of the same type for every inventory level: either fill up the warehouse, empty it, or do nothing. Charnes et al. (1966) extend this result to stochastic, but still positive, prices. The warehouse in our paper is an electricity storage facility.

Rempala (1994) imposes a limit on the rate at which the inventory can be increased, and shows the optimality of a threshold-type policy. Secomandi (2010) extends the model of Charnes et al. (1966) to include both upward and downward inventory adjustment limits and establishes that a double-threshold policy is optimal. Threshold-type policies continue to be optimal in Kaminski et al. (2008) for a continuous time version of the problem, and Devalkar et al. (2011) for the case of a commodity processor that faces procurement and processing capacity constraints and can convert a single input commodity into multiple output commodities. Other related work can be found in the commodity and energy real options literature (see, e.g., Smith and McCardle 1999, Geman 2005, and Secomandi and Seppi 2014), including Chen and Forsyth (2007), Boogert and de Jong (2008), Thompson et al. (2009), Lai et al. (2010a,b), and Wu et al. (2012). All of these papers assume that the commodity price is positive. As discussed in §1, the optimal storage policy structure that we establish for the fast storage case differs considerably from those of Charnes et al. (1966), Rempala (1994), and Secomandi (2010).

Several electricity storage papers assume perfect information on future electricity prices, including Graves et al. (1999), Figueiredo et al. (2006), Walawalkar et al. (2007), Sioshansi et al. (2009), and Hittinger et al. (2012). This assumption yields a linear program, the optimal solution of which

gives the optimal storage decisions for a given price path. In contrast, we model price as a stochastic process, and derive the optimal fast storage policy. Two papers on electricity storage that also model price uncertainty are Mokrian and Stephen (2006) and Xi et al. (2014): Mokrian and Stephen (2006) assess the value of different electricity storage facilities in an electricity market; Xi et al. (2014) co-optimize multiple usages of storage, including energy usage and backup service. However, neither of these papers considers the case when prices can be negative, nor do they derive any optimal policy structure.

Zhou et al. (2014) consider the possibility of negative electricity prices when jointly optimizing wind-based electricity generation and storage. In contrast, we focus on examining the trading of electricity surpluses using the storage and disposal strategies in the presence of negative prices.

3. Model for the storage strategy and its analysis

We model the storage strategy in §3.1. We analyze the optimal policy structure for the case of fast storage in §3.2.

3.1. Model

We consider a merchant using a storage strategy to manage electricity surpluses in an electricity wholesale market. Consistent with the literature, e.g., Walawalkar et al. (2007), Sioshansi et al. (2009), and Hittinger et al. (2012), we do not consider bidding in a forward market. The merchant trades electricity during a finite horizon in each period $t \in \mathcal{T} := \{1, \dots, T\}$; in the terminal period $T + 1$, any electricity left in the storage facility is worthless.

Let $\vec{p}_t := (\xi_t, J_t)$ be a two-dimensional price-component vector, where $\vec{p}_t \in \mathcal{P} \subseteq \mathfrak{R}^2$, ξ_t is a mean-reverting component, and J_t is a spike component. The two price components evolve over time according to exogenous and mutually independent stochastic processes. The electricity price (\$/MWh) in period t given \vec{p}_t , is the known function $P_t(\vec{p}_t) : \mathcal{T} \times \mathcal{P} \rightarrow \mathfrak{R}$. We defer to §5.1 the specific discussion of this price function and the price-component processes; their specifications do not affect our structural analysis in §3.2 as long as they satisfy the assumption that the merchant is a price taker whose trading decisions do not affect market prices, which holds throughout our paper.

We assume the merchant carries out the storage strategy with an industrial battery or a flywheel (a device used to store rotational energy), two common types of electricity storage (EPRI 2004, Akhil et al. 2013, Gyuk et al. 2013). However, our model can be easily modified to represent other storage facilities, such as compressed-air energy storage and pumped-hydro storage (EPRI 2004, The Economist 2012, Akhil et al. 2013). The merchant's storage facility is finite in energy capacity

(“the maximum amount of energy that the system can deliver to the load without being recharged,” Eyer and Corey, 2010); without loss of generality, we normalize the energy capacity to 1 (energy unit). We let \mathcal{X} be the set of feasible energy (inventory) levels in each period and define this set as the interval $[0, 1]$.

A fraction $1 - \eta$ of energy inventoried in the storage facility dissipates during one period ($1 - \eta$ is the self-discharging rate of the storage facility when on stand by; equivalently, η , in $(0, 1]$, is the storing efficiency). We assume that this inventory loss occurs at the end of each period; this is a reasonable assumption as the storing efficiency is close to 1 for many types of storage facilities (EPRI 2004), including all those listed above.

We denote as C (energy units/period) the maximum amount of energy that can be purchased from or sold to the market in one period, due to battery-architecture or flow constraints. This quantity is also referred to as the charging/discharging power capacity if one ignores the energy lost when charging or discharging the storage facility. We let α and β (both in $(0, 1]$) represent the fraction of energy lost when charging and discharging the storage facility, respectively, that is, the charging and discharging efficiencies. The quantities $\alpha \cdot C$ and C/β are thus the net charging power capacity and the gross discharging power capacity, respectively, and are analogous to the inventory adjustment capacity in the commodity storage literature. The round-trip efficiency (the ratio of the quantity of electricity withdrawn to that injected) is denoted as r and is the product $\alpha \cdot \beta$.

For a given period length, different types of storage facilities can be modeled by varying the value of the power capacity C . The case $\alpha \cdot C < 1$ represents slow storage and the case $\alpha \cdot C \geq 1$ corresponds to fast storage. With a period length of five minutes—the period length of the real-time market in NYISO (NYISO 2011) used in our numerical study in §5—examples of slow storage include industrial batteries, which usually take a few hours to fully charge/discharge (EPRI 2004, Akhil et al. 2013); examples of fast storage include flywheels that can be charged/discharged fully within minutes or even seconds (EPRI 2004, Akhil et al. 2013).

The action (decision) for each period t is denoted by a_t and represents the inventory change between periods t and $t + 1$ before accounting for the inventory loss: $a_t < 0$ is the inventory decrease due to selling, so the quantity sold to the market is $-a_t \cdot \beta$; $a_t \geq 0$ is the inventory increase due to buying, so the quantity bought from the market is a_t/α . The immediate payoff function $R(a_t, \vec{p}_t) : \Re \times \mathcal{P} \rightarrow \Re$ from performing action a_t when the price-component vector is \vec{p}_t is defined as follows:

$$R(a_t, \vec{p}_t) := \begin{cases} -P_t(\vec{p}_t) \cdot a_t \cdot \beta, & \text{if } a_t < 0, \\ -P_t(\vec{p}_t) \cdot a_t/\alpha, & \text{if } a_t \geq 0. \end{cases} \quad (1)$$

This definition relies on the assumption of a price taking merchant whose trading activity has no effect on the market price.

We denote by x_t the inventory (electricity, in energy units) in the merchant's storage facility at the beginning of period t . For simplicity and without loss of generality, we consider values of x_t in set \mathcal{X} , even though, due to the inventory loss, the inventory level at the start of periods 2 through T cannot exceed η before the action a_t is performed (but $x_t + a_t$ can equal 1). We define the feasible action set for inventory level $x_t \in \mathcal{X}$ as

$$\Psi(x_t) := \{a_t \in \mathfrak{R} : a_t \geq -x_t, a_t \leq 1 - x_t, a_t \geq -C/\beta, a_t \leq \alpha \cdot C\},$$

where the first two constraints that define this set stipulate that the inventory change cannot exceed the available energy in the storage facility and the remaining space of this facility, respectively, and the third and fourth constraints in this definition enforce the flow capacity limits.

In each period $t \in \mathcal{T}$, the sequence of events is as follows:

(i) At the beginning of period t , the merchant observes the inventory level x_t and the price-component vector \vec{p}_t , decides the amount of electricity a_t/α or $-\beta \cdot a_t$ to buy from or sell to the market, and incurs the trading cash flow $R(a_t, \vec{p}_t)$.

(ii) Electricity flows from the market to the storage facility in the case of buying or vice versa in the case of selling and the charging or discharging loss $(1 - \alpha)a_t/\alpha$ or $-(1 - \beta)a_t$, respectively, occurs: the inventory level x_t changes to $x_t + a_t$.

(iii) At the end of period t , the inventory loss takes place, so that the inventory level at the start of period $t + 1$ equals $\eta(x_t + a_t)$.

We formulate the merchant's storage model as a finite-horizon MDP. Each stage of this MDP corresponds to a time period in set \mathcal{T} . The state variables in each stage t are x_t and \vec{p}_t . Denote by Π the set of feasible policies and by $A_t^\pi(x_t, \vec{p}_t)$ the decision rule of feasible policy π in stage t and state (x_t, \vec{p}_t) . The objective is to find a feasible policy that maximizes the stage 1 market value of the merchant's cash flows incurred during the finite horizon:

$$\max_{\pi \in \Pi} \sum_{t \in \mathcal{T}} \delta^{t-1} \mathbb{E} [R(A_t^\pi(x_t^\pi, \vec{p}_t), P_t(\vec{p}_t)) | x_1, \vec{p}_1], \quad (2)$$

where $\delta \in (0, 1]$ is the risk-free discount factor; \mathbb{E} is expectation with respect to \vec{p}_t and x_t^π , the latter of which is the inventory level achieved in stage t by policy π (we use a risk-neutral probability measure for \vec{p}_t , Seppi 2002, which then induces a joint distribution for x_t^π and \vec{p}_t); and x_1 and \vec{p}_1 are the given initial (stage 1) inventory level and price-component vector.

Let $V_t(x_t, \vec{p}_t)$ denote the value function in stage $t \in \mathcal{T}$ and state $(x_t, \vec{p}_t) \in \mathcal{X} \times \mathcal{P}$. This function satisfies the Bellman equation

$$V_t(x_t, \vec{p}_t) = \max_{a_t \in \Psi(x_t)} R(a_t, P_t(\vec{p}_t)) + \delta \mathbb{E}_t [V_{t+1}(\eta(x_t + a_t), \vec{p}_{t+1})], \quad (3)$$

where $\mathbb{E}_t[\cdot]$ is shorthand notation for $\mathbb{E}[\cdot | \vec{p}_t]$ and we define $V_{T+1}(x_{T+1}, \vec{p}_{T+1}) := 0$ for all $(x_{T+1}, \vec{p}_{T+1}) \in \mathcal{X} \times \mathcal{P}$. Solving the optimization on the right hand side of (3) yields an optimal action in stage t and state (x_t, \vec{p}_t) . However, doing so is in general difficult, because for our continuous-state model the continuation function $\delta \mathbb{E}_t [V_{t+1}(\cdot, \vec{p}_{t+1})]$ is difficult to characterize and compute exactly. Specifically, when storage is slow, in general this function is neither (quasi-) concave nor (quasi-) convex, and even characterizing the structure of an optimal policy, let alone computing such policy, is difficult. Thus, in our numerical investigation in §5, where we focus on slow storage, we numerically solve a discrete-state version of model (3) by standard backward dynamic programming.

3.2. Structural analysis for the case of fast storage

In contrast to the difficulty of characterizing the structure of an optimal slow-storage policy, in §3.2.1 we are able to establish the optimal policy structure of model (2) for the fast storage case, which we illustrate in §3.2.2 using examples. Despite its limitation to the fast storage case, our analysis brings to light the impact of negative prices on the structure of the optimal storage policy of Charnes et al. (1966), who also consider fast storage but only for prices that are non-negative. It also allows us to further contrast the resulting structure against the optimal slow-storage policy structures of Rempala (1994) and Secomandi (2010), who, like Charnes et al. (1966), also assume non-negative prices.

3.2.1. Optimal policy structure We first split the optimization in (3) into two optimizations: one allows only selling and the other allows only buying. We then find the optimal solution to each of these two optimizations. For the selling optimization, the optimal decision is to sell to empty the storage facility for all the inventory levels below a threshold that depends on the stage and price components, and to do nothing for all the inventory levels above it. For the buying optimization, the optimal decision is to buy to fill up the storage facility for all the inventory levels above a threshold that depends on the stage and price components, and to do nothing for all the inventory levels below it. Finally, we combine the optimal solutions of these two optimizations to derive the optimal solution for (3).

We define by y_t the inventory level at the end of period t , after both performing the feasible action a_t given the inventory level at the start of period t , x_t , and incurring the inventory loss:

$y_t := \eta(x_t + a_t)$; y_t thus belongs to the set $[0, \eta]$ since a_t takes values in set $\Psi(x_t)$, which reduces to $[-x_t, 1 - x_t]$ when storage is fast. Substituting $a_t \equiv y_t/\eta - x_t$ into the objective function of (3) and maximizing over $y_t \in [0, \eta]$ rather than over $a_t \in \Psi(x_t)$, we obtain

$$V_t(x_t, \vec{p}_t) = \max_{y_t \in [0, \eta]} R(y_t/\eta - x_t, \vec{p}_t) + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})]. \quad (4)$$

The function $V_t(x_t, \vec{p}_t)$ can be equivalently expressed as

$$V_t(x_t, \vec{p}_t) = \max \{V_t^S(x_t, \vec{p}_t), V_t^B(x_t, \vec{p}_t)\}, \quad (5)$$

where $V_t^S(x_t, \vec{p}_t)$ and $V_t^B(x_t, \vec{p}_t)$ are the value functions attainable by optimally selling and buying, respectively, in stage t and state (x_t, \vec{p}_t) :

$$V_t^S(x_t, \vec{p}_t) := \max_{y_t \in [0, \eta x_t]} -y_t \cdot P_t(\vec{p}_t) \cdot \beta/\eta + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t) \cdot \beta, \quad (6)$$

$$V_t^B(x_t, \vec{p}_t) := \max_{y_t \in [\eta x_t, \eta]} -y_t \cdot P_t(\vec{p}_t)/(\alpha\eta) + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t)/\alpha. \quad (7)$$

We can solve the maximization in (3) by solving the two maximizations in (6) and (7) and picking the best of their solutions. After removing their respective constant terms (given x_t and \vec{p}_t), the optimizations in (6) and (7) reduce to

$$\max_{y_t \in [0, \eta x_t]} -y_t \cdot P_t(\vec{p}_t) \cdot \beta/\eta + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})], \quad (8)$$

$$\max_{y_t \in [\eta x_t, \eta]} -y_t \cdot P_t(\vec{p}_t)/(\alpha\eta) + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})]. \quad (9)$$

To avoid trivial cases, we make a standard assumption about the expected future electricity price, which holds throughout this paper.

ASSUMPTION 1. *For every $t \in \mathcal{T} \setminus \{T\}$, $\mathbb{E}[|P_\tau(\vec{p}_\tau)| | \vec{p}_t] < \infty$ for all $\tau \in \mathcal{T}$ and $\tau > t$ and $\vec{p}_t \in \mathcal{P}$.*

Lemma 1 states the convexity of the value functions of model (3) when storage is fast.

LEMMA 1. *For every $t \in \mathcal{T}$, for the fast storage case it holds that $|V_t(x_t, \vec{p}_t)| < \infty$ and $V_t(x_t, \vec{p}_t)$ is convex in $x_t \in \mathcal{X}$ given any $\vec{p}_t \in \mathcal{P}$.*

We denote the optimal solutions to (8) and (9) by $y_t^S(x_t, \vec{p}_t)$ and $y_t^B(x_t, \vec{p}_t)$, respectively (taking them to be the smallest and largest optimal solutions to these optimizations when these models admit multiple optimal solutions, respectively). Lemma 1 enables us to characterize the quantities $y_t^S(x_t, \vec{p}_t)$ and $y_t^B(x_t, \vec{p}_t)$ in Lemma 2, which states that the optimal action for the selling optimization (8) is either to sell to empty the storage facility or to do nothing, and the one for the buying optimization (9) is either to buy to fill up the storage facility or to do nothing.

LEMMA 2. *Given $t \in \mathcal{T}$, $x_t \in \mathcal{X}$, and $\vec{p}_t \in \mathcal{P}$, the value of $y_t^S(x_t, \vec{p}_t)$ can be either 0 or ηx_t and the value taken by $y_t^B(x_t, \vec{p}_t)$ can be either ηx_t or η .*

Lemma 2 suggests that the quantities $y_t^S(x_t, \vec{p}_t)$ and $y_t^B(x_t, \vec{p}_t)$ have a threshold structure in the inventory x_t given the stage t and price-component vector \vec{p}_t . Lemma 3 states the existence of such structures with inventory threshold functions $X_t^S(\vec{p}_t)$ and $X_t^B(\vec{p}_t)$ of the stage t and the price-component vector \vec{p}_t , and characterizes the value taken by these functions using the possible values of $y_t^S(1, \vec{p}_t)$ and $y_t^B(0, \vec{p}_t)$, both of which, by Lemma 2, are 0 or η . To facilitate our analysis, we define the objective functions of optimizations (8) and (9) as $w_t^S(y_t, \vec{p}_t)$ and $w_t^B(y_t, \vec{p}_t)$, respectively:

$$\begin{aligned} w_t^S(y_t, \vec{p}_t) &:= -y_t \cdot P_t(\vec{p}_t) \cdot \beta / \eta + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})], \\ w_t^B(y_t, \vec{p}_t) &:= -y_t \cdot P_t(\vec{p}_t) / (\alpha \eta) + \delta \mathbb{E}_t [V_{t+1}(y_t, \vec{p}_{t+1})]. \end{aligned}$$

LEMMA 3. *For the fast storage case there exist inventory threshold functions $X_t^S(\vec{p}_t)$ and $X_t^B(\vec{p}_t)$: $\mathcal{T} \times \mathcal{P} \rightarrow \mathcal{X}$ such that for every $t \in \mathcal{T}$, $x_t \in \mathcal{X}$, and $\vec{p}_t \in \mathcal{P}$ it holds that*

$$\begin{aligned} y_t^S(x_t, \vec{p}_t) &= \begin{cases} 0, & \forall x_t \in [0, X_t^S(\vec{p}_t)], \\ \eta x_t, & \forall x_t \in (X_t^S(\vec{p}_t), 1], \end{cases} \\ y_t^B(x_t, \vec{p}_t) &= \begin{cases} \eta x_t, & \forall x_t \in [0, X_t^B(\vec{p}_t)), \\ \eta, & \forall x_t \in [X_t^B(\vec{p}_t), 1]. \end{cases} \end{aligned}$$

The values taken by the functions $X_t^S(\vec{p}_t)$ and $X_t^B(\vec{p}_t)$ for t and \vec{p}_t depend on the two corresponding possible values of $y_t^S(1, \vec{p}_t)$ and $y_t^B(0, \vec{p}_t)$, as follows:

$$X_t^S(\vec{p}_t) = \begin{cases} 1, & \text{if } y_t^S(1, \vec{p}_t) = 0, \\ \max y / \eta \in \mathcal{X} \text{ such that } w_t^S(y_t, \vec{p}_t) = w_t^S(0, \vec{p}_t), & \text{if } y_t^S(1, \vec{p}_t) = \eta, \end{cases} \quad (10)$$

$$X_t^B(\vec{p}_t) = \begin{cases} 0, & \text{if } y_t^B(0, \vec{p}_t) = \eta, \\ \min y / \eta \in \mathcal{X} \text{ such that } w_t^B(y_t, \vec{p}_t) = w_t^B(\eta, \vec{p}_t), & \text{if } y_t^B(0, \vec{p}_t) = 0. \end{cases} \quad (11)$$

Given a stage t and a price-component vector \vec{p}_t , Lemma 3 states that when restricting the type of feasible actions to selling or doing nothing the value taken by the threshold function $X_t^S(\vec{p}_t)$ splits the feasible inventory level \mathcal{X} into two regions: a sell-to-empty region on its left, and a do-nothing region on its right; likewise, when restricting the type of feasible actions to buying or doing nothing, the value taken by the threshold function $X_t^B(\vec{p}_t)$ splits the feasible inventory level \mathcal{X} into two regions: a do-nothing region on its left, and a buy-to-fill-up region on its right. The value taken by the function $X_t^S(\vec{p}_t)$ can be smaller than, equal to, or larger than the value taken by the function $X_t^B(\vec{p}_t)$, as illustrated in §3.2.2.

The inventory threshold functions $X_t^S(\vec{p}_t)$ and $X_t^B(\vec{p}_t)$ are useful in characterizing the optimal policy structure stated in Proposition 1, under the realistic assumption that the round-trip efficiency r is strictly less than 1, i.e., $\alpha \neq 1$ or $\beta \neq 1$. (We consider the case $r = 1$ in our discussion of

this proposition below its statement, and further when discussing Example 1 in §3.2.2.) When the value of $X_t^S(\vec{p}_t)$ is equal to or larger than the value of $X_t^B(\vec{p}_t)$, this structure can also depend on one of the two additional inventory threshold functions $Z_t^{(1)}(\vec{p}_t)$ and $Z_t^{(2)}(\vec{p}_t)$, which we specify in this proposition. We denote as $a_t^*(x_t, \vec{p}_t)$ the optimal action in stage t and state (x_t, \vec{p}_t) for model (3).

PROPOSITION 1. *Suppose that the round-trip efficiency r is strictly less than 1 and storage is fast. For every stage $t \in \mathcal{T}$ and price component vector $\vec{p}_t \in \mathcal{P}$ the feasible inventory set \mathcal{X} can be partitioned into no more than three regions where it is respectively optimal to sell to empty the storage facility, buy to fill up the storage facility, and do nothing, as specified in Cases 1-3 below.*

Case 1: $0 \leq X_t^S(\vec{p}_t) < X_t^B(\vec{p}_t) \leq 1$. It holds that

$$a_t^*(x_t, \vec{p}_t) = \begin{cases} -x_t, & \forall x_t \in [0, X_t^S(\vec{p}_t)], \\ 0, & \forall x_t \in (X_t^S(\vec{p}_t), X_t^B(\vec{p}_t)), \\ 1 - x_t, & \forall x_t \in [X_t^B(\vec{p}_t), 1]. \end{cases}$$

Case 2: $X_t^S(\vec{p}_t) = 1$ and $X_t^B(\vec{p}_t) = 0$.

2(i) If $V_t^S(x_t, \vec{p}_t) \geq V_t^B(x_t, \vec{p}_t)$ for all $x_t \in \mathcal{X}$ then

$$a_t^*(x_t, \vec{p}_t) = -x_t, \quad \forall x_t \in \mathcal{X};$$

2(ii) if $V_t^S(x_t, \vec{p}_t) \leq V_t^B(x_t, \vec{p}_t)$ for all $x_t \in \mathcal{X}$ and $V_t^S(x_t, \vec{p}_t) < V_t^B(x_t, \vec{p}_t)$ for at least one $x_t \in \mathcal{X}$ then

$$a_t^*(x_t, \vec{p}_t) = 1 - x_t, \quad \forall x_t \in \mathcal{X};$$

2(iii) otherwise there exists an inventory threshold function $Z_t^{(1)}(\vec{p}_t)$ defined on $\mathcal{T} \times \mathcal{P}$ that returns a value in the interior of the feasible inventory set \mathcal{X} where $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ cross and

$$a_t^*(x_t, \vec{p}_t) = \begin{cases} 1 - x_t, & \forall x_t \in [0, Z_t^{(1)}(\vec{p}_t)], \\ -x_t, & \forall x_t \in (Z_t^{(1)}(\vec{p}_t), 1]. \end{cases}$$

Case 3: $0 < X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t) \leq 1$ or $0 \leq X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t) < 1$.

3(i) If $X_t^S(\vec{p}_t) = 1$ and $V_t^S(x_t, \vec{p}_t) \geq V_t^B(x_t, \vec{p}_t)$ for all $x_t \in \mathcal{X}$ then

$$a_t^*(x_t, \vec{p}_t) = -x_t, \quad \forall x_t \in \mathcal{X};$$

3(ii) if $X_t^B(\vec{p}_t) = 0$ and $V_t^S(x_t, \vec{p}_t) \leq V_t^B(x_t, \vec{p}_t)$ for all $x_t \in \mathcal{X}$ then

$$a_t^*(x_t, \vec{p}_t) = 1 - x_t, \quad \forall x_t \in \mathcal{X};$$




Figure 1 Fast storage optimal policy structure for Cases 1, 2(iii), and 3(iii) in Proposition 1

3(iii) otherwise there exists an inventory threshold function $Z_t^{(2)}(\vec{p}_t)$ defined on $\mathcal{T} \times \mathcal{P}$ that returns a value in one of the sets 3(iia): $[X_t^B(\vec{p}_t), 1 = X_t^S(\vec{p}_t)]$, or 3(iib): $(0 = X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$, or 3(iic): $[X_t^B(\vec{p}_t) \neq 0, X_t^S(\vec{p}_t) \neq 1]$ where $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ cross or meet and

$$a_t^*(x_t, \vec{p}_t) = \begin{cases} -x_t, & \forall x_t \in [0, Z_t^{(2)}(\vec{p}_t)], \\ 1 - x_t, & \forall x_t \in (Z_t^{(2)}(\vec{p}_t), 1]. \end{cases}$$

If $P_t(\vec{p}_t) < 0$ then Cases 1 and 3(iii) are impossible. If $P_t(\vec{p}_t) = 0$ then Cases 1, 2(ii), 2(iii), and 3(iii) are impossible. If $P_t(\vec{p}_t) > 0$ then Case 2 is impossible.

Panel (a) of Figure 1 illustrates the optimal policy structure for Case 1 in Proposition 1: the feasible inventory set \mathcal{X} is partitioned into three ordered regions where it is respectively optimal to sell to empty the storage facility, do nothing, and buy to fill up the storage facility. However, some of these regions can be empty—for instance, it is possible that $X_t^S(\vec{p}_t)$ equals 0 and $X_t^B(\vec{p}_t)$ equals 1, in which case the optimal decision is to do nothing for all feasible inventory levels, i.e., the entire feasible inventory set \mathcal{X} is a do-nothing region and the other two regions are empty.

Panels (b) and (c) of Figure 1, respectively, illustrate the optimal policy structure for Cases 2(iii) and 3(iii) in Proposition 1: the feasible inventory set \mathcal{X} is partitioned into two ordered regions where it is respectively optimal to buy to fill up the storage facility and sell to empty the storage facility, with the former region preceding the latter region in Case 2(iii) and vice versa in Case 3(iii), and these regions are separated by the value taken by the inventory threshold function $Z_t^{(1)}(\vec{p}_t)$ in Case 2(iii) and the inventory threshold function $Z_t^{(2)}(\vec{p}_t)$ in Case 3(iii).

The fast storage optimal policy structure stated in Proposition 1 generalizes the optimal structure established by Charnes et al. (1966) when storage is fast, the commodity price cannot be negative, and the buying and selling prices are linear functions (through the origin) of the electricity price in each stage. In this case the type of optimal action in a given stage and state is of the same type




Figure 2 Optimal policy structure of Charnes et al. (1966), when prices cannot be negative

for all the feasible inventory levels, i.e., either do nothing, buy to fill up the storage facility, or sell to empty it, as illustrated in panels (a), (b), and (c), respectively, of Figure 2.

In the optimal policy structure of Charnes et al. (1966), for fast storage, and in the optimal policy structures of Rempala (1994) and Secomandi (2010), for slow storage, the inventory level that results from performing an optimal action is a non-decreasing function of the feasible inventory level, which is not true in Case 2(iii) of the optimal policy structure stated in Proposition 1 (see panel (b) of Figure 1). Moreover, the optimal policy structure given in Proposition 1 may bring the feasible inventory level *farther* away from the interval $[X_t^S(\vec{p}_t), X_t^B(\vec{p}_t)]$ in Case 1 (see panel (a) of Figure 1). Specifically, this situation occurs if the feasible inventory level falls outside of this interval. The optimal policy structures of Rempala (1994) and Secomandi (2010) bring the feasible inventory level as close as possible to an interval delineated by the values of two base-stock target functions of the stage and price (components).

Mathematically, these differences are due to the *linearity* in inventory of the value functions in the model of Charnes et al. (1966) and the *concavity* of these functions in the models of Rempala (1994) and Secomandi (2010), which contrast the *convexity* in inventory of the value functions of model (3) when storage is fast. Negative prices and a round-trip efficiency strictly less than 1 are necessary for the strict convexity in inventory of these value functions: if prices cannot be negative then the conditions that lead to the optimal policy structure of Charnes et al. (1966) are satisfied, and these value functions are linear in inventory; if the round-trip efficiency is equal to 1 then a simple induction argument shows that these value functions are linear in inventory with slope equal to the electricity price given a stage and price-component vector in this stage, and the optimal policy structure of these authors emerges in this case as well (see the discussion following Example 1 in §3.2.2).




Figure 3 The stage 1 value function and inventory level after performing the optimal action for Example 1

3.2.2. Examples Examples 1-3 illustrate how negative prices together with a round-trip efficiency strictly smaller than 1 give rise to the policy structures displayed in Figure 1. Specifically, Example 1 corresponds to panel (b) of Figure 1 (Case 2(iii) in Proposition 1), Example 2 to panel (a) of Figure 1 (Case 1 in Proposition 1), and Example 3 to panel (c) of Figure 1 (Case 3(iii) in Proposition 1).

EXAMPLE 1 (CASE 2(III)). The time horizon consists of three periods ($T = 3$). The prices in periods 1, 2, and 3, respectively, are deterministic and equal to -4 , -3 , and 0 (their units of measurement are suppressed to ease the exposition both here and in the remainder of this section). For simplicity, we assume no time discounting ($\delta = 1$) and no loss in charging or storing ($\alpha = \eta = 1$). However, half of the electricity is lost in discharging ($\beta = 0.5$). The value function for period 3 is equal to 0, because the price in this period is 0. The negative price in period 2 and the zero value function in period 3 imply that the optimal decision rule in period 2 is to buy to fill up the storage facility at all inventory levels. Omitting the price-component vector argument throughout this example (because the price evolves deterministically), the value function for period 2 is thus equal to the buying value function in this period: $V_2(x_2) = 3(1 - x_2)$. The value function for period 1 is

$$\begin{aligned}
V_1(x_1) &= \max \{V_1^S(x_1), V_1^B(x_1)\} \\
&= \max \left\{ \max_{y_1 \in [0, x_1]} 2y_1 + 3(1 - y_1) - 2x_1, \max_{y_1 \in [x_1, 1]} 4y_1 + 3(1 - y_1) - 4x_1 \right\} \\
&= \max \{3 - 2x_1, 4 - 4x_1\} \\
&= (4 - 4x_1) \cdot 1\{x_1 \in [0, 0.5]\} + (3 - 2x_1) \cdot 1\{x_1 \in [0.5, 1]\},
\end{aligned}$$

where $1\{\cdot\}$ is the indicator function, which equals 1 if its argument is true, and 0 otherwise. The value function $V_1(\cdot)$ is the maximum of the linear functions $V_1^S(\cdot)$ and $V_1^B(\cdot)$, and is thus convex, as displayed in the left panel of Figure 3. Because $y_1^S(1) = 0$ and $y_1^B(0) = 1$, it follows from (10) and (11) in Lemma 3 that $X_1^S = 1$ and $X_1^B = 0$, which corresponds to Case 2(iii) in Proposition 1 with $Z_1^{(1)} = 0.5$, since the functions $V_1^S(\cdot)$ and $V_1^B(\cdot)$ cross once at 0.5. Hence, the optimal period 1 decision rule is to buy to fill up the storage facility when the storage facility is less than half full, and sell to empty the storage facility otherwise. The right panel of Figure 3 illustrates the period 1 inventory level after performing the action corresponding to this decision rule, which exemplifies the structure in panel (b) of Figure 1. \square

The intuition for Example 1 is as follows: since the price trajectory is -4 , -3 , and 0 in periods 1, 2, and 3, buying to fill up the storage facility in period 1 is always better than buying to fill it up in period 2, and in period 2 selling to empty it is never optimal. The only issue is whether it may be optimal to sell to empty the storage facility in period 1 and, once empty, to buy to fill it up in period 2, rather than buying to fill it up in period 1 and, once full, doing nothing in period 2. The payoff of the first policy is $3 - 2x_1$; the payoff of the second policy is $4 - 4x_1$. The first policy is better than the second policy if and only if $3 - 2x_1 > 4 - 4x_1$, i.e., $x_1 > 0.5$, which means that in period 1 it is optimal to sell to empty the storage facility when the inventory level exceeds 0.5, and it is optimal to buy to fill it up when the inventory level is less than or equal to 0.5. This example illustrates how an efficiency loss (in this example 50% for discharging) combined with negative prices can induce a nontrivial relationship between the inventory level and an optimal decision rule; i.e., selling to empty the storage facility at high inventory levels (above 0.5) and buying to fill it up at low inventory levels (at or below 0.5) in period 1.

Expressing the threshold functions in period 1 as depending on the period 1 price P_1 rather than the price-component vector in this period, which we do not model here, we can change this price in Example 1 to obtain different values for the inventory threshold function $Z_1^{(1)}(P_1)$ while keeping the values of the inventory threshold functions $X_1^S(P_1)$ and $X_1^B(P_1)$ fixed at 1 and 0, respectively. For instance, if the period 1 price P_1 varies in the interval $(-6, -3)$ then $Z_1^{(1)}(P_1)$ is the function $2 + 6/P_1$, which strictly decreases on this interval, approaching 1 from below and 0 from above

when P_1 tends to -6 from above and to -3 from below, respectively; $Z_1^{(1)}(-4) = 0.5$ as shown in Example 1.

As stated at the end of §3.2.1, negative prices and a round-trip efficiency strictly less than 1 are necessary for the optimal policy structure in Proposition 1 to differ from the one of Charnes et al. (1966). To illustrate, suppose that in Example 1 everything else is the same but there is no discharging loss ($\beta = 1$). Then the optimal policy changes to buy to fill up the storage facility at all inventory levels in both the first and second periods, and to do nothing at all inventory levels in the third period, which is consistent with the optimal policy structure of Charnes et al. (1966), because the value function is linear in inventory in each period. (In particular, when the round-trip efficiency is equal to 1, it can be shown that it is optimal to sell to empty the storage facility when the spot price in the current stage is more than or equal to the discounted expected next-stage spot price divided by the storing efficiency, and buy to fill up the storage facility otherwise.) On the other hand, everything else being the same as in Example 1, if we increase the prices of all the periods by 4, i.e., the price path is 0, 1, and 4 in periods 1 through 3, then the optimal policy structure of Charnes et al. (1966) ensues, again, because the value function is linear in inventory in each period: specifically, the optimal policy changes to buy to fill up the storage facility at all inventory levels in both the first and second periods, and sell to empty the storage facility at all inventory levels in the third period.

EXAMPLE 2 (CASE 1). This example has the same values for the charging, discharging, and storing efficiencies and the discount factor as Example 1, but it has four periods. The price in the first period is 4; the prices of the the last three periods are the following equally likely paths: $(-12, -10.8, 0)$, $(-12, -7.2, 0)$, and $(54, 0, 0)$. Proceeding as in Example 1, we can analytically determine—for each of the three price paths—the value function for period 2. Averaging these value functions yields the continuation function for period 1, from which we obtain the period 1 optimal decision rule: sell to empty the storage facility at low inventory levels, i.e., those smaller than 0.4, buy to fill it up at high inventory levels, i.e., those larger than 0.6, and do nothing for moderate inventory levels, i.e., those between 0.4 and 0.6. The left panel of Figure 4 presents the period 1 inventory level after performing the optimal action versus the feasible inventory level. This chart illustrates the optimal policy structure in panel (a) of Figure 1 and corresponds to Case 1 in Proposition 1 with the values of the threshold functions X_1^S and X_1^B equal to 0.4 and 0.6, respectively. \square

EXAMPLE 3 (CASE 3(III)). This example modifies Example 2 by taking the period 1 price to be 3.2. The optimal decision rule in period 1 yields only two types of actions: sell to empty the storage




Figure 4 The stage 1 inventory level after performing the optimal action for Example 2 (left panel) and Example 3 (right panel)

facility if the inventory level is below $1/8$, and otherwise buy to fill it up. The right panel of Figure 4 plots the inventory level after performing the optimal action in period 1 as a function of the feasible inventory level. This plot exemplifies the optimal policy structure in panel (c) of Figure 1 and corresponds to Case 3(iii), more specifically Case 3(iiiic), in Proposition 1 with the values of the threshold functions X_1^S , X_1^B , and $Z_1^{(2)}$ equal to $2/7$, $1/11$, and $1/8$, respectively ($1/11 < 1/8 < 2/7$).

In Example 3 varying the price in period 1 yields different values for the inventory threshold functions $X_1^S(P_1)$, $X_1^B(P_1)$, and $Z_1^{(2)}(P_1)$, again expressed as functions of the period 1 price P_1 rather than the, unmodeled here, period 1 price-component vector. For instance, if the period 1 price P_1 changes in the interval $(3, 3.6]$ then we have $X_1^S(P_1) = 0.8/(6 - P_1)$, $X_1^B(P_1) = [(3 - P_1)/(1 - P_1)] \cdot 1\{P_1 \in (3, 3.5]\} + [(3.4 - P_1)/(3 - P_1)] \cdot 1\{P_1 \in (3.5, 3.6]\}$, and $Z_1^{(2)}(P_1) = 2(P_1 - 3)/P_1$. These functions increase without crossing on the interval $(3, 3.6]$ for P_1 , coincide and are equal to $1/3$ when $P_1 = 3.6$, and $X_1^S(P_1)$ tends to $0.8/3$ and both $X_1^B(P_1)$ and $Z_1^{(2)}(P_1)$ tend to 0 (all from above) when P_1 approaches 3 from above.

4. Model for the disposal strategy

In this section we model the merchant disposal strategy for managing electricity surpluses in a wholesale electricity market. Consistent with the storage strategy, the merchant does not bid in a forward market and trades electricity in each period t from set \mathcal{T} . The price evolution satisfies the

same assumptions as in §3.1. The merchant carries out the disposal strategy using a device—e.g., a load bank—that in each period can consume an amount of electricity equal to C , which for consistency we take to be the same as the gross charging power capacity of the storage facility. Because the disposal strategy can either do nothing or buy up to C per period, the optimal disposal strategy is trivial: buy C when the price is negative and do nothing otherwise. Therefore, the period 1 market value of this strategy is

$$\sum_{t \in \mathcal{T}} \delta^{t-1} \mathbb{E}[-P_t(\vec{p}_t) \cdot C \cdot 1\{P_t(\vec{p}_t) < 0\} | \vec{p}_1]. \quad (12)$$

5. Numerical analysis

In this section we discuss our numerical results, which we obtain using the calibrated electricity price model of Zhou et al. (2014). We present this price model in §5.1, introduce the setup of our numerical analysis in §5.2, examine the values of the storage and disposal strategies in §5.3, and investigate the importance of considering negative prices when devising a storage policy in §5.4.

5.1. Calibrated price model

We summarize the price model and the calibration results from Zhou et al. (2014). As discussed in Zhou et al. (2014), and references therein, a model of electricity price evolution should capture four salient features: *negative prices*; *mean reversion*, a tendency to revert back to the mean price level; *spikes*, price jumps that quickly disappear; and *seasonality*, any repeated price pattern at any time scale. To capture these features, Zhou et al. (2014) combine an inverse hyperbolic transformation to accommodate negative prices as in Schneider (2011/12), a mean reverting model as in Lucia and Schwartz (2002), a spike component (a jump that lasts for only one period) as in Seifert and Uhrig-Homburg (2007), and a seasonality function similar to one of the functions in Lucia and Schwartz (2002). Specifically, the period t electricity price function $P_t(\xi_t, J_t)$ is the sum of the period t spike component J_t and the period t despiked-price function $P'_t(\xi_t)$, where ξ_t is the period t mean-reverting component:

$$P_t(\xi_t, J_t) = J_t + P'_t(\xi_t).$$

The spike component J_t is a compound Bernoulli process in which a spike occurs in period t with probability λ and with size distributed according to an empirical distribution. The despiked-price function $P'_t(\xi_t)$ satisfies

$$\sinh^{-1} \left(\frac{P'_t(\xi_t)}{\ell} \right) = \xi_t + f(t),$$

Table 1 Estimated parameters $\hat{\kappa}$ and $\hat{\sigma}$ of the AR(1) model and scale parameter $\hat{\ell}$ (Source: Zhou et al. 2014)

AR(1) model estimated parameters		Estimated scale parameter
$\hat{\kappa}$	$\hat{\sigma}$	$\hat{\ell}$
0.1176	0.1770	30

Table 2 Estimated parameters $\hat{\gamma}^0$, $\hat{\gamma}^{1i}$'s, $\hat{\gamma}^{2j}$'s, and $\hat{\gamma}^{3h}$'s of the seasonality function $f(\cdot)$ (Source: Zhou et al. 2014)

$\hat{\gamma}^0$	1.3778											
	i											
$\hat{\gamma}^{1i}(\cdot 10^{-3})$	1	2	3	4	5	6	7	8	9	10	11	
	9	25.9	40.1	57	-28.9	83.5	214.6	177.4	6.4	-5.3	-72.1	
	j											
$\hat{\gamma}^{2j}(\cdot 10^{-3})$	1	2	3	4	5	6						
	-40.2	-97.6	-5.6	14	20.9	33						
	h											
$\hat{\gamma}^{3h}(\cdot 10^{-3})$	1	2	3	4	5	6	7	8	9	10	11	12
	-73.5	-118.8	-177.6	-194	-153.8	-79.2	36.6	90.6	186.1	265.7	302.6	324.8
	h											
$\hat{\gamma}^{3h}(\cdot 10^{-3})$	13	14	15	16	17	18	19	20	21	22	23	
	320.7	317.1	302.7	295.5	310.1	356	337.7	339.9	313.1	243.7	145.7	

where $\sinh^{-1}(\cdot)$ is the inverse hyperbolic function; ℓ is a scaling parameter; and $f(t)$ is the seasonality function for period t . Zhou et al. (2014) assume that ξ_t is an auto-regressive model of order 1, AR(1):

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

where κ is the mean-reverting rate; σ is the volatility; and $\epsilon_t \sim N(0, 1)$ is an independent and identically distributed normal error term. The mean level of this AR(1) process is 0 because the mean of the despiked price process is captured by the seasonality component $f(t)$, which is modeled as follows:

$$f(t) = \gamma^0 + \sum_{i=1}^{11} \gamma^{1i} \cdot D_t^{1i} + \sum_{j=1}^6 \gamma^{2j} \cdot D_t^{2j} + \sum_{h=1}^{23} \gamma^{3h} \cdot D_t^{3h},$$

where γ^0 is the constant level; γ^{1i} , γ^{2j} , and γ^{3h} represent the seasonality coefficient of month $i \in \{1, \dots, 11\}$, week day $j \in \{1, \dots, 6\}$, and hour $h \in \{1, \dots, 23\}$, respectively; and D_t^{1i} , D_t^{2j} , and D_t^{3h} are dummy variables equal to 1 if period t is in month i , week day j , and hour h , respectively, and 0 otherwise.

Zhou et al. (2014) calibrate this model to prices observed between 2005 and 2008 in the New York City zone of the NYISO real-time market. The frequency of observations in this market is 5 minutes. We summarize in Tables 1-2 and in Figure 5 the calibration results of Zhou et al. (2014). The resulting mean absolute error and root mean square error are 7.6349 \$/MWh and 12.4023




Figure 5 Empirical spike distribution; the estimated spike probability $\hat{\lambda}$ is 0.0751 (Source: Zhou et al. 2014)

\$/MWh, respectively, which correspond to 9% and 15% of the average observed price, 85.1206 \$/MWh. The calibration error is thus relatively small. The frequency of negative prices estimated on 10,000 Monte Carlo simulated price paths sampled from the calibrated price model is close to that observed in the data: 0.44% in the simulation versus 0.45% in the data. Hence, this calibrated price model is adequate for our analysis purposes.

5.2. Numerical setup

We consider various industrial batteries with the same energy capacity of 10 MWh and the same power capacity of 1 MW (about 1% of the size of the NYISO New York City Zone real-time market; NYISO, 2014), but with different values of the round-trip efficiency parameter. This setting corresponds to the slow storage case. We focus on slow storage rather than fast storage because industrial batteries are currently more common in grid-level electricity storage applications than fast storage devices (such as flywheels) due to their larger energy capacity and longer duration of charge/discharge (Eyer and Corey 2010, Akhil et al. 2013). We use a period length equal to 5 minutes, which is consistent with the frequency of the observed prices used to calibrate the price model (see §5.1). Thus, the value of the parameter C is the power capacity rescaled to a five-minute period: $1\text{MW}/(12\text{ periods/hour}) = 1/12\text{ MWh/period}$.

The round-trip efficiency of an industrial battery depends on its technology and specific design: it varies between 0.4 and 0.5 for a metal air battery, 0.6 and 0.7 for a nickel-cadmium battery, 0.70 and 0.80 for a lead acid battery, 0.75 and 0.85 for a flow battery (such as zinc-bromine), 0.85 and 0.90 for a sodium-sulfur battery, and 0.9 and 0.95 for a lithium-ion battery (EPRI 2004, Eyer and Corey 2010, ESA 2013).

To consider a wide range of battery technologies, we choose values of the round-trip efficiency parameter r in the set $\{0.5, 0.6, 0.7, 0.8, 0.9\}$. For all the values of r in this set, we experimented with different possible feasible combinations of values for the charging efficiency α and the discharging

efficiency β (recall that $r = \alpha \cdot \beta$), and found that our results did not change substantially. We thus report the results for the most common combination: equal charging and discharging efficiencies (EPRI 2004, Eyer and Corey 2010, Hittinger et al. 2012), i.e., $\alpha = \beta = \sqrt{r}$. We considered values of the one-period storing efficiency parameter η in the set $\{1, 0.9999988, 0.9999965, 0.9999941\}$, which are derived from the one-month self-discharge rates of 0% for a sodium sulfur battery, 1% for a lead acid battery, 3% for a nickel-cadmium battery, and 5% for a nickel-metal hydride battery, respectively (EPRI 2004). Because of the low sensitivity of our findings to the value of this parameter, we report the results for the case of no self discharge ($\eta = 1$).

We consider a one year horizon, so the total number of (five-minute) periods in our MDP is 105,120 ($= 12 \cdot 24 \cdot 365$). We set the discount factor δ for each period to 0.9999999, corresponding to an annual risk-free interest rate of 1% with continuous compounding (recall that we use risk-neutral valuation).

For each given set of parameters we solve a discretized version of model (3) by using backward dynamic programming. As in Zhou et al. (2014), we use the method of Jaillet et al. (2004) to discretize the AR(1) process as a trinomial lattice with five-minute time increments that specifies attainable price levels and their transition probabilities for each stage. This constructed lattice converges in 6 stages to 11 levels. The spike process is, by definition, discrete. We discretize the feasible inventory set into 121 equally spaced levels (we tried using more inventory levels, and our results were unchanged), discretizing the feasible action space for each such inventory level accordingly.

We take the value of the storage strategy to be the value function in the initial stage with zero inventory level of the version of model (3) discretized as just discussed (the beginning of period 1 corresponds to 00:00 on January 1-st and we let $\vec{p}_1 = (0, 0)$). We let the value of the disposal strategy be the value of the version of expression (12) evaluated in a backward fashion again using the discretized AR(1) process and the discrete spike process.

5.3. Comparison of the values of the storage and disposal strategies

We compare the value of the storage strategy at different round-trip efficiencies and the value of the disposal strategy. Figure 6 plots these values for different frequencies of negative prices, which we achieve by introducing separate probabilities for positive and negative spikes, fixing the probability of a positive spike to the estimated spike probability, and varying the probability of a negative spike based on this estimated value. (We obtained qualitatively similar results when generating more frequent negative prices by increasing the absolute values of the negative calibrated values of the hourly seasonality parameters $\hat{\gamma}^{3h}$'s displayed in Table 2.) When negative prices occur more




Figure 6 Impact of different frequencies of negative prices on the values of the disposal and storage strategies

frequently, the values of all the strategies increase; for each given frequency of negative prices, the value of each considered storage strategy is always higher than the value of the disposal strategy. Additionally, while the value of the disposal strategy is close to zero at a frequency of negative prices equal to 0.5% (recall that 0.45% is the frequency of negative prices in the data used for calibration), the value of this strategy becomes quite substantial at a 10% negative price frequency (the frequency of negative prices in ERCOT in the past few years; Huntowski et al. 2012): about \$118,000 in our one year horizon given a power capacity of 1 MW (118,000 \$/MW-year = 118 \$/kW-year), but still less than the value of the storage strategy, e.g., about 391 \$/kW-year when the value of the round-trip efficiency parameter r equals 0.8.

To better contrast the values of these strategies, Figure 7 breaks down both the value of the storage strategy corresponding to a value of 0.8 for the round-trip efficiency parameter r and the value of the disposal strategy for different price intervals. Each bar represents the stage 1 estimated market value of all the cash flows incurred during the entire horizon that are transacted at the corresponding price interval. The value of the considered storage strategy arises both from selling electricity at positive prices (most notably the [100 \$/MWh, 200 \$/MWh] interval) and from buying electricity at negative prices—this storage strategy also buys at low positive prices, mostly in the interval [0, 100 \$/MWh]. By definition, the disposal strategy can only buy at negative prices. For some (negative) price intervals, the value of the disposal strategy is slightly larger than that of the storage strategy, because the battery cannot buy at negative prices when it is full—this difference is not perceivable in Figure 7.

Despite being less valuable than the storage strategy, the disposal strategy might be appealing

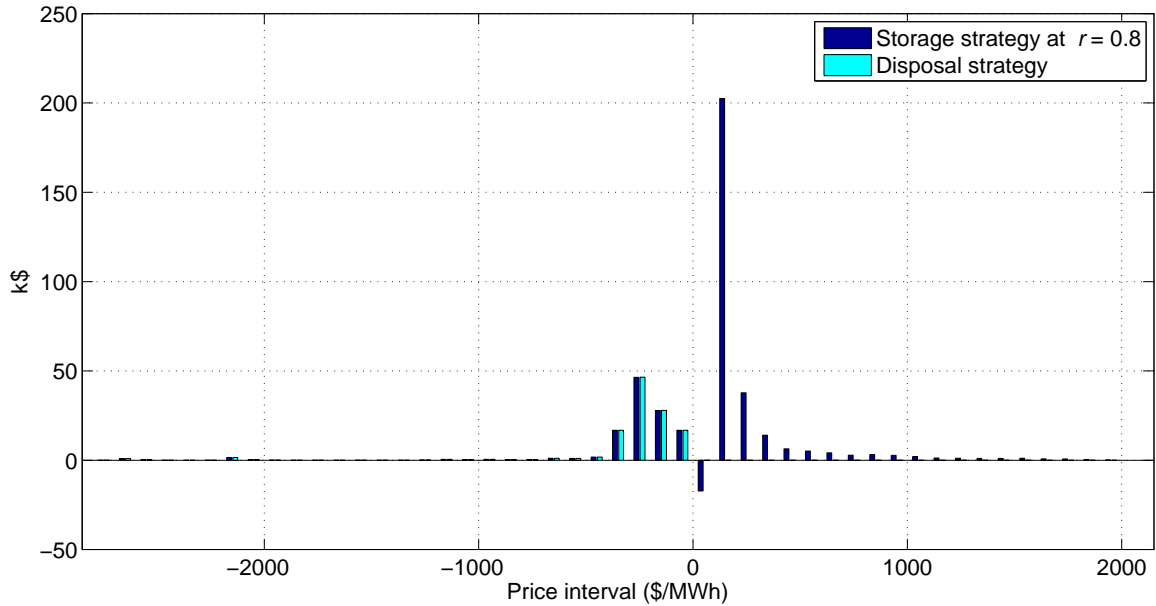


Figure 7 Value generated at different price intervals by the storage strategy when the round-trip efficiency parameter r is set equal to 0.8 and by the disposal strategy when negative prices occur 10% of the time

to merchants because devices for destroying electricity are much cheaper than those for storing it. For instance, a load bank costs about 60 \$/kW (Emerson 2014, Coffman Electrical Equipment Co 2014), but a battery with a round-trip efficiency parameter r equal to 0.8 costs between 1,000 \$/kW and 6,000 \$/kW (Akhil et al. 2013). Therefore, at high frequencies of negative prices the profitability of the disposal strategy (the estimated market value of its cash flows minus the capital cost) might be larger than that of the storage strategy. Merchants might thus be induced to invest in load banks in markets where negative prices occur often, such as the Electric Reliability Council of Texas (Huntowski et al. 2012) and the European Energy Exchange (Nicolosi 2010, Brandstätt et al. 2011).

5.4. Relevance of modeling negative prices when determining a storage policy

To examine the relevance of modeling the possibility of negative prices when devising a storage policy, we compare the values of the storage policies displayed in Figure 6 with their corresponding values when we ignore the possibility of negative prices in the optimization in our discretized version of model (3). Specifically, to compute the latter values, we set any negative price to zero when evaluating the payoff function $R(a_t, P_t(\vec{p}_t))$ in this optimization, obtain the resulting optimized actions for every stage and state, and finally evaluate these actions under the original price model that does admit negative prices.

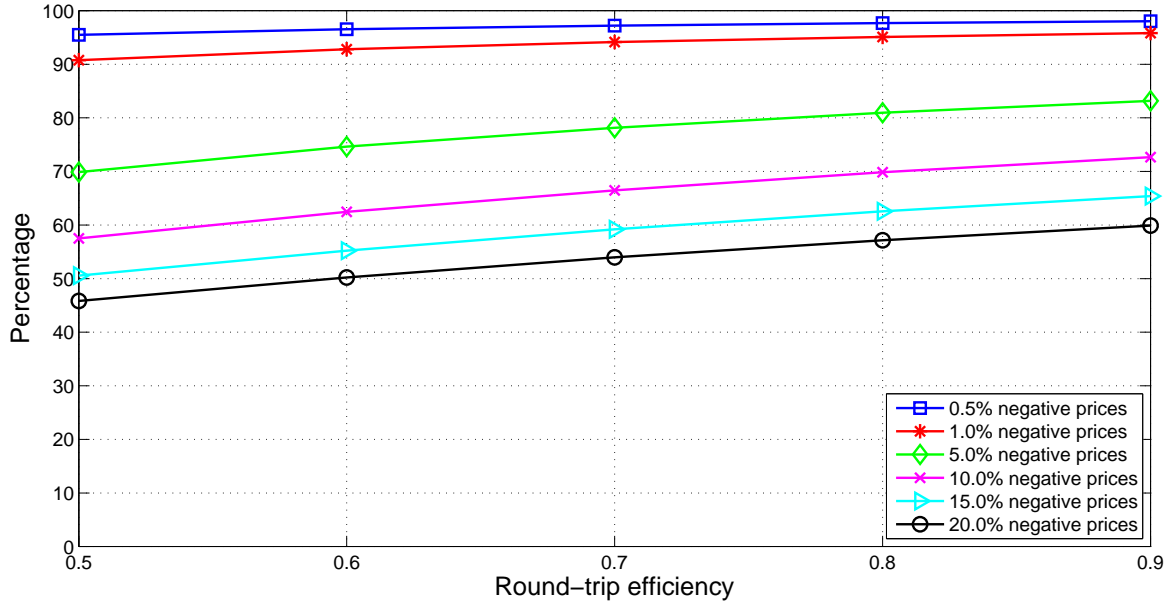


Figure 8 Percentage of the value of the storage strategy retained if negative prices are ignored when determining a storage policy

Figure 8 plots the percentage of the value of the storage strategy captured by storage policies determined when ignoring negative prices. As expected, if negative prices occur infrequently ignoring negative prices when determining a storage policy has little influence on the value of the storage strategy. However, if negative prices happen more than 5% of the time, neglecting negative prices in devising a storage policy results in a substantial loss of value, e.g., a loss of 30% for a battery with round-trip efficiency equal to 0.8 when negative prices occur 10% of the time. This loss ensues because this heuristic storage policy buys more than is optimal at low positive prices. Consequently, the battery is full more often than is optimal when negative prices do occur, and the ability to purchase at such attractive prices is thus lost.

Our results suggest that modeling negative prices when obtaining a storage policy is currently relevant in some markets, and is likely to become more important as negative prices become more frequent in practice.

6. Conclusions

Motivated by the empirical observation that electricity prices can be negative, we investigate how to manage electricity surpluses from the perspective of a merchant. We model as an MDP the problem of managing a storage facility in a wholesale market with prices that are potentially negative and derive the optimal policy structure of this MDP for the fast storage case. We demonstrate that this optimal policy structure generalizes a classic result of Charnes et al. (1966) and differs significantly

from typical threshold policies known to be optimal in the literature when prices are only non-negative. In the slow storage case, using an extant price model calibrated to NYISO real-time market data, we conclude that ignoring negative prices when devising a storage policy can incur a significant loss of value when negative prices occur more than 5% of the time. We also find that the value of the disposal strategy may be sizable, although smaller than that of the storage strategy. For instance, when negative prices occur 10% of the time, the value of the disposal strategy is 118 \$/kW-year compared to 391 \$/kW-year for the storage strategy corresponding to a battery with a 80% round-trip efficiency. However, as the capital costs of load banks are much cheaper than those of batteries, investing in such disposal devices might be more attractive to merchants.

While the storage strategy is common in most commodity markets, the disposal strategy is relatively rare. A type of disposal can be found in the practice of hydro spill: during rainy seasons, excess water in a dam can be spilled to satisfy environmental or irrigation requirements (Ikura and Gross 1984, BPA 2010, 2013). However, our analysis shows that there is potentially a hidden value in this practice, as pumping and spilling water could be used as a method to destroy negatively-priced electricity surpluses. As wind penetration increases and negative prices potentially become more frequent, the disposal strategy may become more common.

In addition to providing guidance on the management of electricity surpluses at the firm level, our results also raise awareness of potential issues at the societal level. Assuming that the disposal strategy is less socially desirable than the storage strategy, our results have three policy implications. First, government policies may need to be enacted to promote the storage strategy, as our results suggest that merchants may find the disposal strategy more profitable (i.e., once capital costs are accounted for) than the storage strategy. Second, the Production Tax Credit (as mentioned in §1) may ultimately need to be modified (e.g., made hour-dependent or price-dependent) in order to promote generating wind energy *only* when it is needed. Currently, this subsidy is constant within a year (DSIRE 2014), and may thus promote generation of wind energy that is subsequently destroyed. Third, curtailing wind energy may need to be used to mitigate the imbalance between electricity supply and demand (Brandstätt et al. 2011, Wu and Kapuscinski 2013). At present, operators in many regions curtail wind energy *only* for safety reasons (Genoese et al. 2010, Nicolosi 2010, Brandstätt et al. 2011), which may lead to negative prices and hence potential destruction of green energy. However, the assumption that the disposal strategy is less socially desirable than the storage strategy may not always hold. Sioshansi (2014), for example, identifies conditions when electricity storage reduces social welfare, and thus it may be possible that the disposal strategy, or a mix of the storage and disposal strategies, is socially beneficial at times. Investigating these issues requires an equilibrium model, which is an opportunity for future research.

Appendix A: Proofs

Proof of Lemma 1

We first prove finiteness. For each stage t , given any \vec{p}_t , it holds that

$$|V_t(x_t, \vec{p}_t)| \leq (1/\alpha) \cdot |P_t(\vec{p}_t)| + \sum_{\tau=t+1}^T |\mathbb{E}_t[P_\tau(\vec{p}_\tau)]| \cdot (1/\alpha) < \infty,$$

where the first inequality holds because in each stage the quantity bought cannot exceed $1/\alpha$, the quantity sold cannot exceed 1, and $1 \leq 1/\alpha$, and the second inequality is due to Assumption 1.

We prove convexity by induction considering the set $\mathcal{T} \cup \{T+1\}$. The function $V_{T+1}(\cdot, \vec{p}_{T+1}) \equiv 0$ is trivially convex on \mathcal{X} . We make the induction hypothesis that $V_\tau(\cdot, \vec{p}_\tau)$ is convex on \mathcal{X} for every stage τ from stage $t+1$ through T . For stage t , $\mathbb{E}_t[V_{t+1}(\cdot, \vec{p}_{t+1})]$ is convex since expectation preserves the convexity of $V_{t+1}(\cdot, \vec{p}_{t+1})$, which holds by the induction hypothesis. It follows that the objective function in (6) is convex in y_t on $[0, \eta x_t]$ given \vec{p}_t . Thus, an optimal solution to the optimization in (6) can be either 0 or ηx_t . By similar arguments, an optimal solution to the optimization in (7) can be either ηx_t or η . Hence, we can limit the search for an optimal solution to the optimization in (4) to the three candidate points 0, ηx_t , and η , which we include in the set \mathcal{Y}_t . We denote the objective function for the optimization in (4) by $v_t(x_t, y_t, \vec{p}_t)$. We can thus write

$$V_t(x_t, \vec{p}_t) = \max_{y_t \in \mathcal{Y}_t} v_t(x_t, y_t, \vec{p}_t).$$

We next show that the function $v_t(\cdot, y_t, \vec{p}_t)$ is convex on \mathcal{X} for each of the possible values of y_t in set \mathcal{Y}_t :

- If $y_t = 0$ then it follows from (5) and (6) that $v_t(x_t, 0, \vec{p}_t) = \delta \mathbb{E}_t[V_{t+1}(0, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t) \cdot \beta$, which is linear, and hence, convex, in x_t on \mathcal{X} .
- If $y_t = \eta x_t$ then, by (5) and either one of (6) and (7), we have $v_t(x_t, \eta x_t, \vec{p}_t) = \delta \mathbb{E}_t[V_{t+1}(\eta x_t, \vec{p}_{t+1})]$, which is easily shown to be convex in x_t on \mathcal{X} based on the induction hypothesis.
- If $y_t = \eta$ then (5) and (7) imply that $v_t(x_t, \eta, \vec{p}_t) = -P_t(\vec{p}_t)/\alpha + \delta \mathbb{E}_t[V_{t+1}(\eta, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t)/\alpha$, which is linear, and hence convex, in x_t on \mathcal{X} .

By Proposition A-3 in Porteus (2002), the function $V_t(\cdot, \vec{p}_t)$ is convex on \mathcal{X} . It follows from the principle of mathematical induction that $V_t(\cdot, \vec{p}_t)$ is convex on \mathcal{X} for all $t \in \mathcal{T} \cup \{T+1\}$, and hence is convex on \mathcal{X} when restricting attention to $t \in \mathcal{T}$. \square

Proof of Lemma 2

It follows from Lemma 1 that both the objective functions for the optimizations in (8) and (9) are finite and convex in y_t given any t and \vec{p}_t . Therefore, a maximizer for each of these optimizations must be one of the two end points of their corresponding feasible sets, i.e., $y_t^S(x_t, \vec{p}_t)$ is either 0 or ηx_t and $y_t^B(x_t, \vec{p}_t)$ is either ηx_t or η . \square

Proof of Lemma 3

Fix t and \vec{p}_t . Consider the optimization in (8). According to Lemma 2, the value of $y_t^S(x_t, \vec{p}_t)$ can be either 0 or ηx_t for each x_t on \mathcal{X} . The convexity of the function $w_t^S(\cdot, \vec{p}_t)$ on $[0, \eta]$, which follows from Lemma 1, implies that there exists a critical inventory level (the largest one if there are multiple such levels) denoted

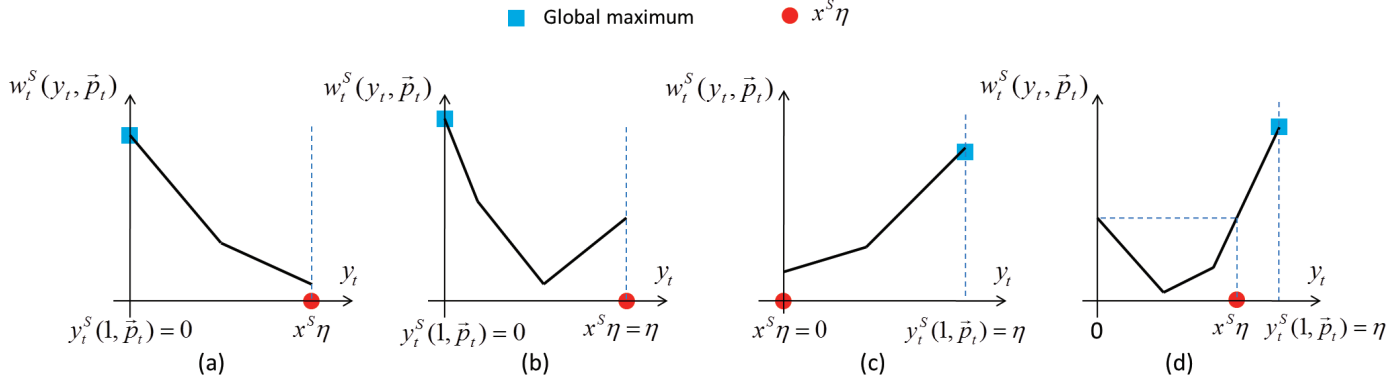


Figure 9 The value of $x^S \eta$ for the possible values of $y_t^S(1, \vec{p}_t)$

by x^S such that $y_t^S(x_t, \vec{p}_t) = 0$ for all x_t in $[0, x^S]$, and $y_t^S(x_t, \vec{p}_t) = \eta x_t$ for all x_t in $(x^S, 1]$. We can characterize x^S by considering the two possible values of $y_t^S(1, \vec{p}_t)$, i.e., 0 and η :

- Suppose that $y_t^S(1, \vec{p}_t) = 0$, which occurs when the function $w_t^S(\cdot, \vec{p}_t)$ is either non-increasing on $[0, \eta]$, as in panel (a) of Figure 9, or non-monotonic on $[0, \eta]$, as in panel (b) of Figure 9. In this case we have $y_t^S(x_t, \vec{p}_t) = 0$ for all x_t in \mathcal{X} , so that $x^S = 1$.

- Suppose that $y_t^S(1, \vec{p}_t) = \eta$, which occurs when the function $w_t^S(\cdot, \vec{p}_t)$ is either non-decreasing on $[0, \eta]$, as in panel (c) of Figure 9, or non-monotonic on $[0, \eta]$, as in panel (d) of Figure 9. In this case we have $y_t^S(x_t, \vec{p}_t) = 0$ for all x_t smaller than the maximum value of y_t/η such that $w_t^S(y_t, \vec{p}_t) = w_t^S(0, \vec{p}_t)$, and $y_t^S(x_t, \vec{p}_t) = \eta x_t$ otherwise. Therefore, x^S is this maximum value of y_t/η .

Consequently, we let $X_t^S(\vec{p}_t)$ be the function that returns x^S at the given t and \vec{p}_t .

Consider the optimization in (9). As stated in Lemma 2, the value of $y_t^B(x_t, \vec{p}_t)$ can be either ηx_t or η for each x_t on \mathcal{X} . By the convexity of the function $w_t^B(\cdot, \vec{p}_t)$ on $[0, \eta]$, there exists a critical inventory level (the smallest one if there are multiple such levels) denoted by x^B such that $y_t^B(x_t, \vec{p}_t) = \eta x_t$ for all x_t in $[0, x^B)$, and $y_t^B(x_t, \vec{p}_t) = \eta$ for all x_t in $[x^B, 1]$. Considering the two possible values of $y_t^B(0, \vec{p}_t)$, i.e., 0 and η , allows us to characterize x^B :

- If $y_t^B(0, \vec{p}_t) = 0$ then we have $y_t^B(x_t, \vec{p}_t) = \eta x_t$ for all x_t smaller than the minimum value of y_t/η such that $w_t^B(y_t, \vec{p}_t) = w_t^B(\eta, \vec{p}_t)$, and $y_t^B(x_t, \vec{p}_t) = \eta$ otherwise. Hence, x^B is this minimum value of y_t/η .

- If $y_t^B(0, \vec{p}_t) = \eta$ then we have $y_t^B(x_t, \vec{p}_t) = \eta$ for all x_t in \mathcal{X} , so that $x^B = 0$.

We thus let $X_t^B(\vec{p}_t)$ be the function that returns x^B at the given t and \vec{p}_t . \square

Proof of Proposition 1

We fix both t and \vec{p}_t throughout this proof.

Case 1: $0 \leq X_t^S(\vec{p}_t) < X_t^B(\vec{p}_t) \leq 1$. We consider the following exhaustive and mutually exclusive cases illustrated in panel (a) of Figure 10:

- For x_t in $[0, X_t^S(\vec{p}_t)]$, according to Lemma 3 the optimal action for the optimization in (6) is to sell to empty the storage facility, thus dominating the feasible do-nothing action for this optimization. Moreover, the do-nothing action is optimal for the optimization in (7) by Lemma 3 and the assumption $X_t^S(\vec{p}_t) < X_t^B(\vec{p}_t)$.

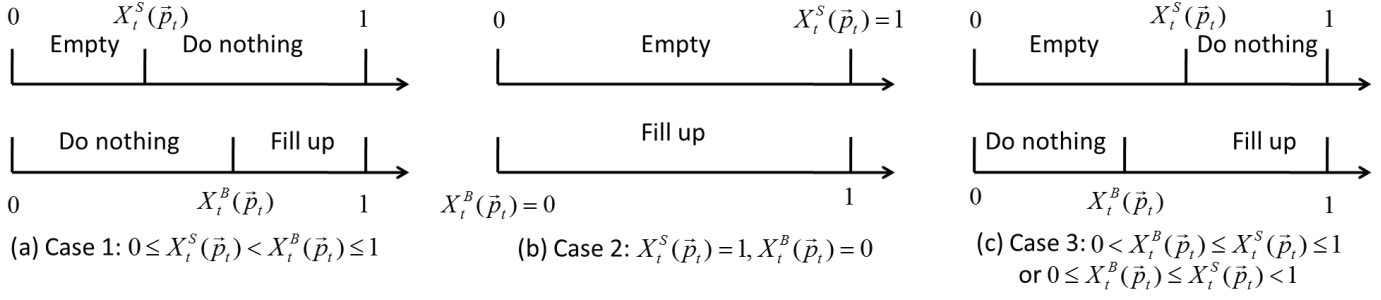


Figure 10 Optimal solutions to the optimizations in (8) and (9) in different regions for Cases 1, 2, and 3 in Proposition 1

It follows from (5) that the optimal action for the optimization in (4) is to sell to empty the storage facility, i.e., $a_t^*(x_t, \vec{p}_t) = -x_t$.

- For x_t in $[X_t^B(\vec{p}_t), 1]$, buying to fill up the storage facility is the optimal action for the optimization in (7) by Lemma 3, and hence this action dominates the do-nothing action for this optimization. Furthermore, Lemma 3 and the assumption $X_t^S(\vec{p}_t) < X_t^B(\vec{p}_t)$ imply that the do-nothing action is optimal for the optimization in (6). By (5), the optimal action for the optimization in (4) is to buy to fill up the storage facility, i.e., $a_t^*(x_t, \vec{p}_t) = 1 - x_t$.

- For x_t in $(X_t^S(\vec{p}_t), X_t^B(\vec{p}_t))$, doing nothing is optimal for both the optimization in (6) and the optimization in (7) according to Lemma 3. By (5), doing nothing is also the optimal action for the optimization in (4), i.e., $a_t^*(x_t, \vec{p}_t) = 0$.

Case 2: $X_t^S(\vec{p}_t) = 1$ and $X_t^B(\vec{p}_t) = 0$ (as illustrated in panel (b) of Figure 10). By the assumption $X_t^S(\vec{p}_t) = 1$ and Lemma 3 we have $w_t^S(0, \vec{p}_t) \geq w_t^S(\eta, \vec{p}_t)$, or, equivalently,

$$\delta \mathbb{E}_t [V_{t+1}(0, \vec{p}_{t+1})] \geq -P_t(\vec{p}_t) \cdot \beta + \delta \mathbb{E}_t [V_{t+1}(\eta, \vec{p}_{t+1})]. \quad (13)$$

Moreover, the assumption $X_t^B(\vec{p}_t) = 0$ and Lemma 3 imply that $w_t^B(0, \vec{p}_t) \leq w_t^B(\eta, \vec{p}_t)$, or, equivalently,

$$\delta \mathbb{E}_t [V_{t+1}(0, \vec{p}_{t+1})] \leq -P_t(\vec{p}_t)/\alpha + \delta \mathbb{E}_t [V_{t+1}(\eta, \vec{p}_{t+1})]. \quad (14)$$

Combining inequalities (13) and (14) yields

$$\left(\frac{1}{\alpha} - \beta \right) \cdot P_t(\vec{p}_t) \leq 0. \quad (15)$$

Inequality (15) and the assumption $r < 1$, i.e., $\alpha \neq 1$ or $\beta \neq 1$, yield

$$P_t(\vec{p}_t) \leq 0. \quad (16)$$

The assumption $X_t^S(\vec{p}_t) = 1$ and Lemma 3 imply that $y_t^S(x_t, \vec{p}_t) = 0$ for all x_t in \mathcal{X} , and it then follows from (6) that $V_t^S(x_t, \vec{p}_t) = \delta \mathbb{E}_t [V_{t+1}(0, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t) \cdot \beta$ for all x_t in \mathcal{X} , which is a linear function of x_t on \mathcal{X} . By the assumption $X_t^B(\vec{p}_t) = 0$ and Lemma 3, we have $y_t^B(x_t, \vec{p}_t) = \eta$ for all x_t in \mathcal{X} , so from (7) we obtain $V_t^B(x_t, \vec{p}_t) = -P_t(\vec{p}_t)/\alpha + \delta \mathbb{E}_t [V_{t+1}(\eta, \vec{p}_{t+1})] + x_t \cdot P_t(\vec{p}_t)/\alpha$ for all x_t in \mathcal{X} , which is a linear function

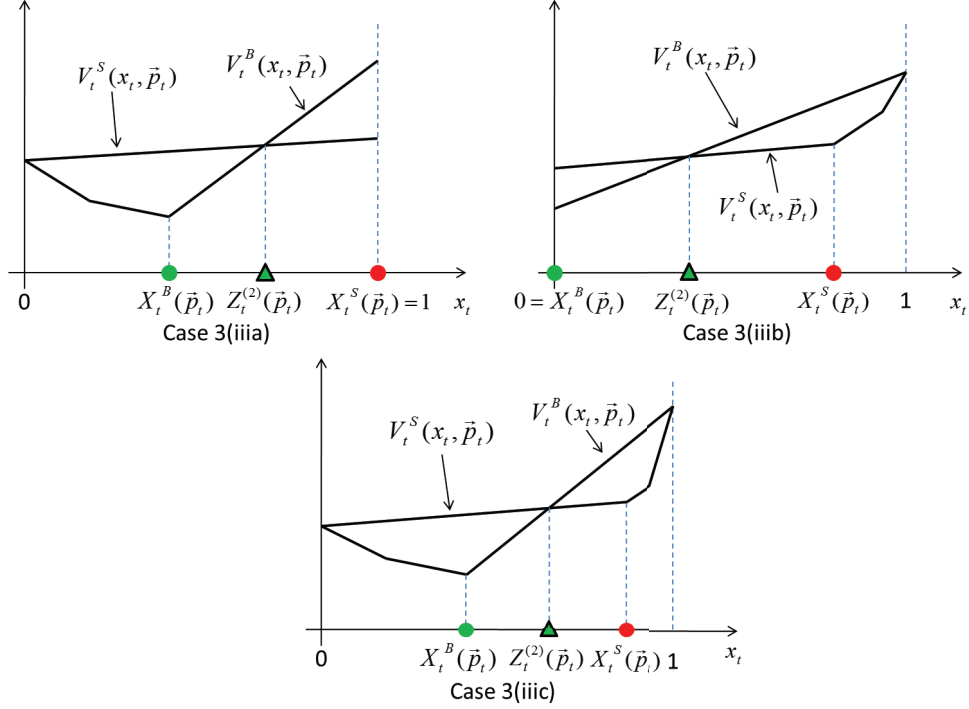


Figure 11 Cases 3(iia)-3(iic)

of x_t on \mathcal{X} . The linearity of the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ on \mathcal{X} implies that they can cross at most once on the interior of \mathcal{X} . If no such crossing occurs then there are two cases to consider:

2(i) If $V_t^S(x_t, \vec{p}_t) \geq V_t^B(x_t, \vec{p}_t)$ for all x_t in \mathcal{X} then $a_t^*(\vec{p}_t) = -x_t$ for all x_t in \mathcal{X} .

2(ii) If $V_t^S(x_t, \vec{p}_t) \leq V_t^B(x_t, \vec{p}_t)$ for all x_t in \mathcal{X} and $V_t^S(x_t, \vec{p}_t) < V_t^B(x_t, \vec{p}_t)$ for at least one x_t in \mathcal{X} , then $a_t^*(\vec{p}_t) = 1 - x_t$ for all x_t in \mathcal{X} .

If the stated crossing occurs we have the following case:

2(iii) We let $z^{(1)}$ be the point on the interior of \mathcal{X} where the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ cross. Inequality (16) and the fact that the stated crossing cannot occur when $P_t(\vec{p}_t) = 0$, because in this case the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ are constant on \mathcal{X} , imply that $P_t(\vec{p}_t) < 0$. Thus, also by the assumption $r < 1$, the slope of $V_t^S(\cdot, \vec{p}_t)$, which is $\beta \cdot P_t(\vec{p}_t)$, is strictly greater than the slope of $V_t^B(\cdot, \vec{p}_t)$, which is $P_t(\vec{p}_t)/\alpha$. Hence, we conclude that $V_t^S(x_t, \vec{p}_t) < V_t^B(x_t, \vec{p}_t)$ for all x_t in $[0, z^{(1)})$ and $V_t^S(x_t, \vec{p}_t) > V_t^B(x_t, \vec{p}_t)$ for all x_t in $(z^{(1)}, 1]$. In other words, we have $a_t^*(x_t, \vec{p}_t) = 1 - x_t$ for all x_t in $[0, z^{(1)})$, and $a_t^*(x_t, \vec{p}_t) = -x_t$ for all x_t in $(z^{(1)}, 1]$, where we arbitrarily break the tie between the optimal actions $1 - x_t$ and $-x_t$ at $x_t = z^{(1)}$ in favor of the action $-x_t$. We let $Z_t^{(1)}(\vec{p}_t)$ be the function that takes the value $z^{(1)}$ at the given t and \vec{p}_t .

Case 3: $0 < X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t) \leq 1$ or $0 \leq X_t^B \leq X_t^S(\vec{p}_t) < 1$ (as illustrated in panel (c) of Figure 10). We consider the following exhaustive and mutually exclusive cases for the potential values taken by the functions $X_t^S(\vec{p}_t)$ and $X_t^B(\vec{p}_t)$.

- $0 < X_t^B(\vec{p}_t) \leq 1 = X_t^S(\vec{p}_t)$. As in Case 2, if $X_t^S(\vec{p}_t) = 1$ then $V_t^S(\cdot, \vec{p}_t)$ is linear on \mathcal{X} . Lemma 3 implies that $V_t^B(\cdot, \vec{p}_t)$ is linear on $[X_t^B(\vec{p}_t), 1]$. By Lemma 3 and the assumption $X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t)$, it holds that on

$[0, X_t^B(\vec{p}_t)]$ doing nothing is optimal for the optimization in (7) and suboptimal for the optimization in (6), so that $V_t^B(\cdot, \vec{p}_t) \leq V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^B(\vec{p}_t)]$. Thus, the functions $V_t^B(\cdot, \vec{p}_t)$ and $V_t^S(\cdot, \vec{p}_t)$ can intersect (meet or cross) at most once on $[X_t^B(\vec{p}_t), 1]$. If $V_t^S(x_t, \vec{p}_t) \geq V_t^B(x_t, \vec{p}_t)$ for all x_t in \mathcal{X} then we have Case 3(i) for which, given that $X_t^S(\vec{p}_t) = 1$, $a_t^*(x_t, \vec{p}_t) = -x_t$ for all x_t in \mathcal{X} . Otherwise we have Case 3(iia): the functions $V_t^B(\cdot, \vec{p}_t)$ and $V_t^S(\cdot, \vec{p}_t)$ either cross on $(X_t^B(\vec{p}_t), 1)$ or they coincide on $[0, X_t^B(\vec{p}_t)]$ and $V_t^B(\cdot, \vec{p}_t)$ is strictly larger than $V_t^S(\cdot, \vec{p}_t)$ on $(X_t^B(\vec{p}_t), 1]$; the case when crossing occurs is illustrated in panel 3(iia) of Figure 11, which also relies on the equality $V_t^B(0, \vec{p}_t) = V_t^S(0, \vec{p}_t)$, as implied by Lemma 3 and the assumptions $X_t^B(\vec{p}_t) > 0$ and $X_t^S(\vec{p}_t) = 1$. We provide the optimal action for Case 3(iia), together with the optimal actions for Cases 3(iib) and 3(iic), later in Case 3(iii) of this proof.

- $X_t^B(\vec{p}_t) = 0 \leq X_t^S(\vec{p}_t) < 1$. As in Case 2, the function $V_t^B(\cdot, \vec{p}_t)$ is linear on \mathcal{X} . The function $V_t^S(\cdot, \vec{p}_t)$ is linear on $[0, X_t^S(\vec{p}_t)]$ by Lemma 3. Lemma 3 and the assumption $X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t)$ yield that on $(X_t^S(\vec{p}_t), 1]$ doing nothing is optimal for the optimization in (6) and suboptimal for the optimization in (7), and hence $V_t^S(\cdot, \vec{p}_t) \leq V_t^B(\cdot, \vec{p}_t)$ on $(X_t^S(\vec{p}_t), 1]$. Thus, the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ can intersect (meet or cross) at most once on $[0, X_t^S(\vec{p}_t)]$. If $V_t^S(x_t, \vec{p}_t) \leq V_t^B(x_t, \vec{p}_t)$ for all x_t in \mathcal{X} then we have Case 3(ii) for which, because $X_t^B(\vec{p}_t) = 0$, $a_t^*(x_t, \vec{p}_t) = 1 - x_t$ for all x_t in \mathcal{X} . Otherwise we have Case 3(iib): the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ either cross on $(0, X_t^S(\vec{p}_t))$ or they coincide on $[X_t^S(\vec{p}_t), 1]$ and $V_t^B(\cdot, \vec{p}_t)$ is strictly smaller than $V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^S(\vec{p}_t)]$; the case with crossing is illustrated in panel 3(iib) of Figure 11, which also uses the equality $V_t^S(1, \vec{p}_t) = V_t^B(1, \vec{p}_t)$, as implied by Lemma 3 and the assumptions $X_t^B(\vec{p}_t) = 0$ and $X_t^S(\vec{p}_t) < 1$.

- $0 < X_t^B(\vec{p}_t) \leq X_t^S(\vec{p}_t) < 1$. This is Case 3(iic). As in Case 3(iia), the function $V_t^B(\cdot, \vec{p}_t)$ is linear on $[X_t^B(\vec{p}_t), 1]$ and we have $V_t^B(\cdot, \vec{p}_t) \leq V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^B(\vec{p}_t)]$. As in Case 3(iib), the function $V_t^S(\cdot, \vec{p}_t)$ is linear on $[0, X_t^S(\vec{p}_t)]$ and it holds that $V_t^S(\cdot, \vec{p}_t) \leq V_t^B(\cdot, \vec{p}_t)$ on $(X_t^S(\vec{p}_t), 1]$. Thus, the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ can cross at most once on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$. The case when the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ cross on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$ is illustrated in panel 3(iic) of Figure 11, which is also based on the equalities $V_t^S(0, \vec{p}_t) = V_t^B(0, \vec{p}_t)$ and $V_t^S(1, \vec{p}_t) = V_t^B(1, \vec{p}_t)$, as implied by Lemma 3 and the assumptions $X_t^B(\vec{p}_t) \neq 0$ and $X_t^S(\vec{p}_t) \neq 1$. If the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ do not cross on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$, then they must coincide on \mathcal{X} . Moreover, we must have $X_t^B(\vec{p}_t) = X_t^S(\vec{p}_t)$: if $P_t(\vec{p}_t) \neq 0$ and $X_t^B(\vec{p}_t) \neq X_t^S(\vec{p}_t)$ then the respective slopes of the linear functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$ are $\beta P_t(\vec{p}_t)$ and $P_t(\vec{p}_t)/\alpha$, which cannot be equal if $r < 1$, as assumed; if $P_t(\vec{p}_t) = 0$ it is shown later in this proof that this case cannot subsist. Hence, the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ meet at $X_t^B(\vec{p}_t) = X_t^S(\vec{p}_t)$.

We group Cases 3(iia), 3(iib), and 3(iic) into Case 3(iii): we let $z^{(2)}$ be the point in $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)] \cap (0, 1)$, more specifically, $[X_t^B(\vec{p}_t), 1 = X_t^S(\vec{p}_t)]$ for 3(iia), $(0 = X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$ for 3(iib), and $[X_t^B(\vec{p}_t) \neq 0, X_t^S(\vec{p}_t) \neq 1]$ for 3(iic), where the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ intersect and $Z_t^{(2)}(\vec{p}_t)$ be the function that evaluates to $z^{(2)}$ at the given t and \vec{p}_t . Arbitrarily breaking the tie between the optimal actions $-x_t$ and $1 - x_t$ at $x_t = z^{(2)}$ in favor of the action $-x_t$, we thus have

$$a_t^*(x_t, \vec{p}_t) = \begin{cases} -x_t, & \forall x_t \in [0, Z_t^{(2)}(\vec{p}_t)], \\ 1 - x_t, & \forall x_t \in (Z_t^{(2)}(\vec{p}_t), 1]. \end{cases}$$

Impossibility of cases: $P_t(\vec{p}_t) < 0$.

Case 1. By Lemma 2, the only two possible values for $y_t^B(0, \vec{p}_t)$ are 0 and η . If $y_t^B(0, \vec{p}_t) = 0$ we have $y_t^S(1, \vec{p}_t) = 0$ because

$$\begin{aligned} w_t^S(0, \vec{p}_t) &= w_t^B(0, \vec{p}_t) \\ &> w_t^B(\eta, \vec{p}_t) \\ &= -\frac{P_t(\vec{p}_t)}{\alpha} + \delta \mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})] \\ &> -P_t(\vec{p}_t) \cdot \beta + \delta \mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})] \\ &= w_t^S(\eta, \vec{p}_t), \end{aligned}$$

where the first inequality is due to our convention that $y_t^B(\cdot, \vec{p}_t)$ is the largest maximizer for the optimization in (9) and the premise $y_t^B(0, \vec{p}_t) = 0$, and the second inequality to the assumption $P_t(\vec{p}_t) < 0$. Lemma 3 thus implies that $X_t^S(\vec{p}_t) = 1$, and Case 1 is impossible. If $y_t^B(0, \vec{p}_t) = \eta$ it follows from Lemma 3 that $X_t^B(\vec{p}_t) = 0$, and Case 1 cannot occur.

Case 3(iii). The slope of the linear function $V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^S(\vec{p}_t)]$ is $\beta P_t(\vec{p}_t)$ and the slope of the linear function $V_t^B(\cdot, \vec{p}_t)$ on $[X_t^B(\vec{p}_t), 1]$ is $P_t(\vec{p}_t)/\alpha$. The assumptions $P_t(\vec{p}_t) < 0$ and $r < 1$ yield

$$\beta P_t(\vec{p}_t) > \frac{P_t(\vec{p}_t)}{\alpha}. \quad (17)$$

Case 3(iii) when the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ meet at $X_t^B(\vec{p}_t) = X_t^S(\vec{p}_t)$ is impossible because otherwise it would follow from inequality (17) that the function $V_t(\cdot, \vec{p}_t)$ is strictly concave on \mathcal{X} , which contradicts its convexity established in Lemma 1. By construction, Case 3(iia), Case 3(iib), and Case 3(iic) with the qualification that the functions $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ cross on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$ cannot subsist when inequality (17) holds: for Case 3(iia), given that $V_t^B(\cdot, \vec{p}_t) \leq V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^B(\vec{p}_t)]$ it follows from inequality (17) that $V_t^B(\cdot, \vec{p}_t)$ and $V_t^S(\cdot, \vec{p}_t)$ cannot intersect on $[X_t^B(\vec{p}_t), 1 = X_t^S(\vec{p}_t)]$; for Case 3(iib), we have $V_t^B(\cdot, \vec{p}_t) \leq V_t^S(\cdot, \vec{p}_t)$ on $[X_t^S(\vec{p}_t), 1]$ and inequality (17) implies that $V_t^B(\cdot, \vec{p}_t)$ and $V_t^S(\cdot, \vec{p}_t)$ cannot intersect on $(0 = X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$; for Case 3(iic) with the stated qualification, by a similar logic $V_t^B(\cdot, \vec{p}_t)$ and $V_t^S(\cdot, \vec{p}_t)$ cannot cross on $[X_t^B(\vec{p}_t), X_t^S(\vec{p}_t)]$ if inequality (17) is true.

Impossibility of cases: $P_t(\vec{p}_t) = 0$. If $P_t(\vec{p}_t) = 0$ and $\mathbb{E}_t[V_{t+1}(\cdot, \vec{p}_{t+1})]$ is constant on $[0, \eta]$ then Case 2(i) occurs, and hence all the other cases are impossible, because our convention that $y_t^B(\cdot, \vec{p}_t)$ is the largest maximizer for the optimization in (9) and $y_t^S(\cdot, \vec{p}_t)$ is the smallest maximizer for the optimization in (8) yields $y_t^S(1, \vec{p}_t) = 0$ and $y_t^B(0, \vec{p}_t) = \eta$, and then Lemma 3 implies $X_t^S(\vec{p}_t) = 1$ and $X_t^B(\vec{p}_t) = 0$, and $V_t^S(\cdot, \vec{p}_t)$ and $V_t^B(\cdot, \vec{p}_t)$ are both equal to the same constant on \mathcal{X} .

Suppose that $P_t(\vec{p}_t) = 0$, $\mathbb{E}_t[V_{t+1}(\cdot, \vec{p}_{t+1})]$ is not constant on $[0, \eta]$, and $\mathbb{E}_t[V_{t+1}(0, \vec{p}_{t+1})] = \mathbb{E}_t[V_{t+1}(\eta, \vec{p}_{t+1})]$. Then by our convention that $y_t^S(\cdot, \vec{p}_t)$ is the smallest maximizer for the optimization in (8) and $y_t^B(\cdot, \vec{p}_t)$ is the largest maximizer for the optimization in (9) we have $y_t^B(0, \vec{p}_t) = \eta$ and $y_t^S(1, \vec{p}_t) = 0$, and it follows from Lemma 3 that $X_t^S(\vec{p}_t) = 1$ and $X_t^B(\vec{p}_t) = 0$, as well as $V_t^S(\cdot, \vec{p}_t) = \delta \mathbb{E}_t[V_{t+1}(0, \vec{p}_{t+1})] = \delta \mathbb{E}_t[V_{t+1}(\eta, \vec{p}_{t+1})] = V_t^B(\cdot, \vec{p}_t)$ on \mathcal{X} , which corresponds to Case 2(i), and thus all the other cases are impossible.

If $P_t(\vec{p}_t) = 0$, $\mathbb{E}[V_{t+1}(\cdot, \vec{p}_{t+1})]$ is not constant on $[0, \eta]$, and $\mathbb{E}[V_{t+1}(0, \vec{p}_{t+1})] > \mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})]$ then $y_t^S(1, \vec{p}_t) = y_t^B(0, \vec{p}_t) = 0$ and Lemma 3 thus implies that $X_t^S(\vec{p}_t) = 1$ and $X_t^B(\vec{p}_t) \neq 0$, which corresponds to Case 3(i) or Case 3(iia), and hence all the other cases are impossible. By Lemma 3 we have $V_t^S(\cdot, \vec{p}_t) = \delta\mathbb{E}[V_{t+1}(0, \vec{p}_{t+1})]$ on \mathcal{X} and $V_t^B(\cdot, \vec{p}_t) = \delta\mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})] < \delta\mathbb{E}[V_{t+1}(0, \vec{p}_{t+1})] = V_t^S(\cdot, \vec{p}_t)$ on $[X_t^B(\vec{p}_t), 1]$. As shown in the construction of Cases 3(i) and 3(iia), we have $V_t^B(\cdot, \vec{p}_t) \leq V_t^S(\cdot, \vec{p}_t)$ on $[0, X_t^B(\vec{p}_t)]$. We thus conclude that Case 3(i) occurs and Case 3(iia) is also impossible.

If $P_t(\vec{p}_t) = 0$, $\mathbb{E}[V_{t+1}(\cdot, \vec{p}_{t+1})]$ is not constant on $[0, \eta]$, and $\mathbb{E}[V_{t+1}(0, \vec{p}_{t+1})] < \mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})]$ it follows from Lemma 3 that $X_t^S(\vec{p}_t) \neq 1$ and $X_t^B(\vec{p}_t) = 0$, which conforms with Case 3(ii) or Case 3(iiib), and consequently all the other cases are impossible. Lemma 3 implies that $V_t^B(\cdot, \vec{p}_t) = \delta\mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})]$ on \mathcal{X} and $V_t^S(\cdot, \vec{p}_t) = \delta\mathbb{E}[V_{t+1}(0, \vec{p}_{t+1})] < \delta\mathbb{E}[V_{t+1}(\eta, \vec{p}_{t+1})] = V_t^B(\cdot, \vec{p}_t)$ on $[0, X_t^S(\vec{p}_t)]$. In the construction of Cases 3(ii) and 3(iiib) we showed that $V_t^S(\cdot, \vec{p}_t) \leq V_t^B(\cdot, \vec{p}_t)$ on $(X_t^S(\vec{p}_t), 1]$. Hence, Case 3(ii) occurs and Case 3(iiib) is also impossible.

Impossibility of cases: $P_t(\vec{p}_t) > 0$. If $P_t(\vec{p}_t) > 0$ then Case 2 is impossible because inequality (16) is necessary for the existence of this case. \square

Acknowledgments

Nicola Secomandi was supported by NSF grant CMMI1129163. Yangfang Zhou thanks Jay Apt for comments on initial versions, and also thanks Liang LV, Eric Hittinger, Brandon Mauch, and Todd Ryan for stimulating discussions. The authors thank the department editor, Serguei Netessine, and the anonymous associate editor and three referees for their constructive feedback that significantly improved the paper.

References

- AEMO. 2009. High prices: Queensland and New South Wales; negative prices: Victoria, South Australia and Tasmania. Australian Energy Market Operator, Market Event Report, November 3. Available at http://www.aemo.com.au/Electricity/Resources/Reports-and-Documents/Market-Event-Reports/~media/Files/Other/reports/market_events/0180-0003%20pdf.ashx.
- Akhil, A. A., G. Huff, A. B. Currier, B. C. Kaun, D. M. Rastler, S. B. Chen, A. L. Cotter, D. T. Bradshaw, W. D. Gauntlett. 2013. DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA. Tech. rep., Sandia National Laboratories, Albuquerque, NM and Livermore, CA, USA. SANDIA REPORT SAND2013-5131. Available at <http://energy.gov/sites/prod/files/2013/08/f2/ElecStorageHndbk2013.pdf>.
- Boogert, A., C. de Jong. 2008. Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15**(3) 81–98.
- BPA. 2010. Northwest overgeneration: An assessment of potential magnitude and cost. Available at: https://www.bpa.gov/Projects/Initiatives/Oversupply/OversupplyDocuments/BPA_Overgeneration_Analysis.pdf.

-
- BPA. 2013. Seasonal power oversupply in 2012. Available at: <https://www.bpa.gov/Projects/Initiatives/Oversupply/OversupplyDocuments/2013/20130123-Oversupply-2012-Lookback.pdf>.
- Brandstätt, C., K. Brunekreeft, K. Jahnke. 2011. How to deal with negative power price spikes? Flexible voluntary curtailment agreements for large-scale integration of wind. *Energy Policy* **39**(6) 3732–3740.
- Brown, P. 2012. U.S. renewable electricity: How does wind generation impact competitive power markets? Tech. rep., CRS Report for Congress. Available at <http://www.fas.org/sgp/crs/misc/R42818.pdf>.
- Cahn, A. S. 1948. The warehouse problem. *Bulletin of the American Mathematical Society* **54**(11) 1073.
- 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.
- Coffman Electrical Equipment Co. 2014. Avtron load banks. Available at <http://www.steadypower.com/pages.php?pageid=60>.
- Devalkar, S.K., R. Anupindi, A. Sinha. 2011. Integrated optimization of procurement, processing, and trade of commodities. *Operations Research* **59**(6) 1369–1381.
- Dreyfus, S. E. 1957. An analytic solution of the warehouse problem. *Management Science* **4**(1) 99–104.
- DSIRE. 2014. Federal incentives/policies for renewables and efficiency–renewable electricity production tax credit (PTC). Available at: http://dsireusa.org/incentives/incentive.cfm?Incentive_Code=US13F.
- EIA. 2012. Annual energy review. Available at: <http://www.eia.gov/totalenergy/data/annual/showtext.cfm?t=ptb0811a>.
- Emerson. 2014. Load banks. Available at <http://www.emersonnetworkpower.com/en-US/Products/Load-Banks>.
- EPRI. 2004. EPRI-DOE handbook supplement of energy storage for grid connected wind generation applications. Tech. rep., Electric Power Research Institute, Palo Alto, CA, and the U.S. Department of Energy, Washington, DC. Available at <http://www.epri.com/search/Pages/results.aspx?k=EPRI-DOE%20Handbook%20Supplement%20of%20Energy%20Storage%20for%20Grid%20Connected%20Wind%20Generation%20Applications>.
- ESA. 2013. Electricity storage technology comparison. Available at: www.electricitystorage.org/technology/storage_technologies/technology_comparison.
- Eyer, J., G. Corey. 2010. Energy storage for the electricity grid: Benefits and market potential assessment guide—A study for the DOE energy storage systems program. Tech. rep., Sandia National Laboratories, Albuquerque, NM and Livermore, CA, USA. SANDIA REPORT SAND2010-0815. Available at <http://www.sandia.gov/ess/publications/SAND2010-0815.pdf>.

- Figueiredo, F. C., P. C. Flynn, E. A. Cabral. 2006. The economics of energy storage in 14 deregulated power markets. *Energy Studies Review* **14**(2) 131–152.
- Fink, S., C. Mudd, K. Porter, B. Morgenstern. 2009. Wind energy curtailment case studies. Tech. rep., National Renewable Energy Laboratory. Subcontract Report NREL/SR-550-46716.
- Geman, H. 2005. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley and Sons, Chichester, UK.
- 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. *Proceedings of the Seventh International Conference on the European Energy Market*. 1–6.
- Graves, F., T. Jenkin, D. Murphy. 1999. Opportunities for electricity storage in deregulating markets. *The Electricity Journal* **12**(8) 46–56.
- Gyuk, I., M. Johnson, J. Vetrano, K. Lynn, W. Parks, R. Handa, L. Kannberg, S. Hearne, K. Waldrip, R. Braccio. 2013. Grid energy storage. Tech. rep., U.S. Department of Energy. Available at: <http://energy.gov/sites/prod/files/2013/12/f5/Grid%20Energy%20Storage%20December%202013.pdf>.
- Hittinger, E., J.F. Whitacre, J. Apt. 2012. What properties of grid energy storage are most valuable? *Journal of Power Sources* **206**(15 May) 436–449.
- Huntowski, F., A. Patterson, M. Schnitzer. 2012. Negative electricity prices and the production tax credit—Why wind producers can pay us to take their power and why that is a bad thing. Tech. rep., The NorthBridge Group. Available at http://www.nbggroup.com/publications/Negative_Electricity_Prices_and_the_Production_Tax_Credit.pdf.
- Ikura, Y., G. Gross. 1984. Efficient large-scale hydro system scheduling with forced spill conditions. *IEEE Transactions on Power Apparatus and Systems* **PAS-103**(12) 3502–3520.
- Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50**(7) 909–921.
- Kaminski, V., Y. Feng, Z. Pang. 2008. Value, trading strategies and financial investment of natural gas storage assets. Presented at the 2008 Northern Finance Association Conference. Available at <http://www.northernfinance.org/2008/papers/75.pdf>.
- Knittel, K., M. Roberts. 2005. An empirical examination of restructured electricity prices. *Energy Economics* **27**(5) 791–817.
- Lai, G., F. Margot, N. Secomandi. 2010a. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58**(3) 564–582.
- Lai, G., M. X. Wang, N. Secomandi, S. Kekre, A. Scheller-Wolf. 2010b. Valuation of the real option to store liquefied natural gas at a regasification terminal. *Operations Research* **59**(3) 602–616.
- 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, Stanford, CA.
- Nicolosi, M. 2010. Wind power integration and power system flexibility: An empirical analysis of extreme events in Germany under the new negative price regime. *Energy Policy* **38**(11) 7257–7268.
- NYISO. 2011. Markets & Operations–Pricing Data. Available at http://www.nyiso.com/public/markets_operations/market_data/pricing_data/index.jsp.
- NYISO. 2014. The energy markets. Available at http://www.nyiso.com/public/about_nyiso/understanding_the_markets/energy_market/index.jsp.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA.
- 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. Springer, Berlin, 560–566.
- Schneider, S. 2011/12. Power spot price models with negative prices. *The Journal of Energy Markets* **4**(4) 77–110.
- Secomandi, N. 2010. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.
- 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., C. de Jong. 2003. Negative prices in electricity markets. *Commodities Now* **June**.
- Sioshansi, R. 2014. When energy storage reduces social welfare. *Energy Economics* **41**(January) 106–116.
- Sioshansi, R., P. Denholm, T. Jenkin, J. Weiss. 2009. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics* **31**(2) 269–277.
- Smith, J. E., K. F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research* **47**(1) 1–15.
- The Economist. 2012. Packing some power. Available at: <http://www.economist.com/node/21548495>.
- Thompson, M., M. Davison, H. Rasmussen. 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics* **56**(3) 226–238.
- Walawalkar, R., J. Apt, R. Mancini. 2007. Economics of electric energy storage for energy arbitrage and regulation in New York. *Energy Policy* **35**(4) 2558–2568.

- Williams, J. C., B. D. Wright. 1991. *Storage and Commodity Markets*. Cambridge University Press, Cambridge, UK.
- Wu, O., R. Kapuscinski. 2013. Curtailing intermittent generation in electrical systems. *Manufacturing and 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. Smith. 2014. Managing wind-based electricity generation in the presence of storage and transmission capacity. Tepper Working Paper 2011-E36, Carnegie Mellon University, Pittsburgh, PA.