Energy Conversion and Management 76 (2013) 472–482

Contents lists available at ScienceDirect

Energy Conversion and Management

journal homepage: www.elsevier.com/locate/enconman

On the value and price-responsiveness of ramp-constrained storage

Ali Faghih [a,b,*], Mardavij Roozbehani [b,1], Munther A. Dahleh [a,b,2]

[a] Electrical Engineering and Computer Science (EECS) Department, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA
[b] Laboratory for Information and Decision Systems (LIDS), MIT, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Received 3 April 2013
Accepted 22 July 2013
Available online 29 August 2013

Keywords:
Ramp-constrained energy storage
Stochastic dynamic programming
Control systems

ABSTRACT

The primary concerns of this paper are twofold: understanding the value of storage in the presence of ramp constraints and exogenous energy prices, and understanding the implications of the associated optimal storage management policy for qualitative and quantitative characteristics of storage response to real-time prices. The optimal policy, along with the associated finite-horizon time-averaged value of storage, are analytically characterized in this paper. An analytical upper bound on the infinite-horizon time-averaged value of storage is also derived. This bound is valid for any achievable realization of prices when the support of the distribution is fixed, and highlights the dependence of the value of storage on ramp constraints and storage capacity. It is shown that while the value of storage is a non-decreasing function of price volatility, due to the finite ramp rate, the value of storage saturates quickly as the capacity increases, regardless of volatility. To study the implications of the optimal policy, computational experiments are presented that suggest optimal utilization of storage can, in expectation, induce a considerable amount of price elasticity near the average price. Then, a computational framework is presented for characterization of the behavior of storage as a function of price and the state of charge, which illustrates a steep buy/sell phase transition in the price-state plane.

1. Introduction

The growing demand for electricity and the urge to reduce greenhouse emissions promote large-scale integration of renewable energy sources, as well as storage and other demand-side management techniques (see [1] for a detailed discussion) to improve energy efficiency in the future grid. However, renewable energy sources are highly uncertain and intermittent. While energy storage technologies can help mitigate the intermittency and narrow the gap between generation from renewable resources and consumption, they may also add to the uncertainty in the system since the optimal response of storage to prices is a complicated function of both price and the amount of stored energy. Modeling and understanding the behavior of storage in response to real-time market prices is therefore critical for reliable operation of power systems with large amounts of storage.

Energy storage has a clear environmental value; it helps mitigate the intermittency of the renewable resources and thereby maximize their utilization, while reducing the risk associated with making advance commitments for a renewable generation owner. Nevertheless, despite all the potential advantages of storage, if the value of storage as an arbitrage mechanism is not attractive, the markets may not invest sufficiently in storage. Hence, unless proper incentives and pricing policies are in place, the environmental and reliability values of storage might not materialize due to underinvestment. Therefore, there is a need for development of econometric models and characterization of the associated optimal policies that can be used for assessing the value of storage. This paper seeks to provide such characterization by presenting a model for optimal utilization of ramp-constrained storage in response to stochastically varying energy prices and studying the corresponding value of storage.

Availability of econometric models of storage and characterization of the effects of storage on the price elasticity of demand (PED) is also important for system operators who need to maintain stability and guarantee reliability. This is particularly important in the context of electricity markets since the aggregate PED can affect price volatility and sensitivity to disturbances [2]. According to [2], in power grids with information asymmetry between consumers, producers, and system operators, robustness of the system to disturbances is greatly affected by consumers' real-time valuation of electricity and response to real-time prices.

The existing literature covering various dimensions of storage and its applications is extensive, both in the area of trading/warehousing in operations management and in the area of energy storage. The problem of optimizing purchase, i.e. injecting into storage,

* Corresponding author. Address: 77 Mass Ave, 32D-740, Cambridge, MA 02139, USA. Tel.: +1 617 324 1544.

E-mail addresses: afaghih@mit.edu (A. Faghih), mardavij@mit.edu (M. Roozbehani), dahleh@mit.edu (M.A. Dahleh).

1 Address: 77 Mass Ave, 32D-732, Cambridge, MA 02139, USA.

2 Address: 77 Mass Ave, 38-435, Cambridge, MA 02139, USA.

0196-8904/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.enconman.2013.07.072

and sale, i.e., withdrawing from storage, for the case of a warehouse with fixed size and an initial stock of a certain commodity, often referred to as the warehouse problem, is a classical problem in the trading and commercial management of commodities, and has been studied in the literature extensively. Some works in this line of research that are most relevant to the work herein include [3,4], which study the discrete-time case of the warehouse problem by imposing a limit on the amount that can be injected into or withdrawn from the warehouse at each time step, and also [5] which solves the same extension of the warehouse problem in continuous time. Several other works have also considered this extension of the warehouse problem, each for a different application (for instance, see [6]). Although the above-mentioned references and the work herein share similarities in the structure of the optimal policy and the associated value function, the differences in the assumptions on the stochastic price process make the analytical results of these papers different from one another. Unlike all the previous works in this area, it is assumed in this paper that the price at each time period is independent of previous prices. Under this assumption, explicit formulas (recursive and/or closed-form) are derived for the thresholds of the optimal policy and for the value of storage, using the principles of stochastic dynamic programming. The assumption on prices is justified by testing the performance of the optimal policy (in the Appendix) against price data from real-time markets of the Pennsylvania–Jersey–Maryland (PJM) Interconnection and the Independent System Operator (ISO) of New England (ISONE); the findings of these tests suggest that the sensitivity of the optimal policy and value of storage to the assumption of independent prices is low (please see the Appendix for details). Another justification for this assumption is that in practice, empirical estimation of conditional distributions (needed for a Markovian price model) requires significant amounts of data. Back of the envelop calculations show that collecting this much data would require going too far back in the price history. However, due to non-stationarity, doing so would make the data irrelevant. Although one can resort to calibrated models for estimation of correlation in data, or try to learn the thresholds directly, these directions are not pursued in this paper.

The relevant literature that focus particularly on energy storage is also extensive. Various types of electrical energy storage systems have been invented and developed for energy networks, and their characteristics and appropriate applications have been studied [7]. In addition, the feasibility and/or value of integrating particular storage technologies, from battery storage systems [8] to compressed air energy storage systems [9], etc., have been studied extensively. However, there is also a large body of literature on the prospects and economic viability of large-scale integration of storage (for instance, see [10]), without focusing on a particular storage technology. Many of these works study various aspects of the role and value of storage in the management and integration of fluctuating renewable energy sources. The quantitative aspect of the work herein is related to those works that study optimal management strategies for the core problem of controlling energy storage to maximize profits. For instance, [11] addresses this problem by developing a multistage looping algorithm to maximize the profit obtained from a pumped-storage plant using forecasted hourly prices. Some others use dynamic programming; for instance, [12] takes a numerical approach and deterministically solves the dynamic programming problem for particular finite sample paths and then averages the results of these paths. On the other hand, some works such as [13,14] use dynamic programming to address the storage management problem in an analytical framework and, like the work herein, study an analytically tractable model of storage and derive explicit formulas that can give insight into learning the behavior of storage and its implications in the particular setup of interest. In contrast to [13], which obtains its core results under the assumption of uniformly distributed generation from the wind farm, herein storage is studied purely as an arbitrage mechanism that interacts only with the main grid, without any particular assumption on price distribution other than the assumption of independent prices. In contrast to [14], ramp constraints are explicitly included in the model herein, highlighting the effects of ramp constraints on the optimal policy. Nevertheless, some studies have focused on managing storage without any renewable sources connected to it. For instance, [15] considers industrial consumers with time-of-use rates and uses dynamic programming to determine optimal contracts and optimal sizes of battery storage for such consumers. However, they use a deterministic approach and relax ramp constraints.

The model set forth in this paper and some preliminary results were reported in [16]. This paper provides a comprehensive exposition which adds several new ideas, core analytical results, and systematic computational experiments.

The contributions of this paper are summarized as follows: First, a dynamic model for optimal utilization of storage in the presence of ramp constraints under the assumption of independent and exogenous prices is proposed. This model assumes limited storage capacity and allows sell-back of energy to the grid. Using the principles of stochastic dynamic programming, the optimal policy and the corresponding value function are analytically characterized for the finite-horizon case. In particular, recursive equations are provided for computation of the exact value of storage for the finite-horizon storage problem. Then, a closed-form upper bound on the infinite-horizon optimal average value per stage of storage is obtained over all possible realizations of prices within a bounded support. This result highlights how the capacity limit and the ramp constraint bound the value of storage. Next, it is shown that while the value of storage is a non-decreasing function of price volatility, the value of storage saturates quickly due to finite ramping rates as the capacity increases, regardless of price volatility.

The average PED is then studied in a simulated electricity market, where the term "average" reflects the fact that the dependence of storage response on the stage and the internal state of storage has been averaged out. It is shown that optimal utilization of storage may, in expectation, induce a considerable amount of price elasticity near the average price, but little or no elasticity elsewhere. While the demand for electricity has often been considered to be highly inelastic, the existing literature on price elasticity are mostly based on empirical evidence and qualitative reasoning (see, for instance, [17,18]). In this paper, price elasticity is studied in a quantitative framework, characterizing the PED induced by optimally controlled, ramp-constrained storage through an input–output model of response to prices. Next, the study of the PED induced by storage is extended by addressing the PED as a function of the internal state of storage. The results highlight the interplay between state-dependence and price-dependence of the storage response in a computational framework. In particular, the buy/sell phase transition region of the storage response in the price-state plane is illustrated.

In addition to this paper's findings regarding the value and price-responsiveness of storage that have been presented in the main part of the paper, the results on the validity of the assumption of independent prices presented in the Appendix can be used as a benchmark for comparison with other results that assume more complicated models of the stochastic price process.

The remainder of this paper is organized as follows: The dynamic model of storage management is introduced in Section 2. The optimal policies for the storage management problem and the corresponding value function are presented in Section 3. In Section 4, an analytical upper bound on the optimal average value per stage of storage is derived and computational findings on the value of storage are reported. The implications of the optimal policy are then discussed in Section 5, by studying the average PED of a storage system in Subsection 5.1, and the price-responsiveness of a storage system as a function of the storage state in Subsection 5.2. The conclusions are presented in Section 6. The Appendix gives an evaluation of the optimal policy's performance.

2. A dynamic model of storage

2.1. Notation

The set of positive real numbers (integers) is denoted by $\mathbb{R}_+$ ($\mathbb{Z}_+$), and non-negative real numbers (integers) by $\bar{\mathbb{R}}_+$ ($\bar{\mathbb{Z}}_+$). The probability mass function (PMF) of a random variable Λ is denoted by P_Λ, and the cumulative distribution function (CDF) is denoted by F_Λ. P and F are simply used when there is no ambiguity.

2.2. The model

In this section, a dynamic model for optimal management of ramp-constrained storage in the presence of stochastically-varying prices is developed. The first step is to formulate the storage management problem as a stochastic dynamic programming problem over a finite horizon.

2.2.1. The decisions

The decision set of the storage owner at each discrete instant of time $k \in \bar{\mathbb{Z}}_+$ is characterized by a pair

$$(v_k^{\text{in}}, v_k^{\text{out}}) \in [0, \bar{v}^{\text{in}}] \times [0, \bar{v}^{\text{out}}] \tag{1}$$

where, v_k^{in} and v_k^{out} are, respectively, the amount of power that the consumer injects in, or withdraws from the storage. The corresponding upper bounds ($\bar{v}^{\text{in}}$ and $\bar{v}^{\text{out}}$) represent the physical ramp constraints on storage. Also, $v_k = v_k^{\text{in}} - v_k^{\text{out}} \in [-\bar{v}^{\text{out}}, \bar{v}^{\text{in}}]$ denotes the net storage response. With a slight abuse of terminology, the storage response v_k may be referred to as demand, with the understanding that $v_k \leqslant 0$ implies a negative demand.

2.2.2. The price

The price per unit of energy at each stage λ_k is sampled from an exogenous stochastic process that is independently distributed across time, with mean $\bar{\lambda}_k$ and standard deviation σ_k, and support over $[\lambda_k^{\min}, \lambda_k^{\max}] \subset [0, \infty)$. It is assumed that at the beginning of each time interval $[k, k+1]$, the random variable λ_k, is materialized and revealed to the consumer. Note that the distributions of λ_k can be different at different stages; however, it is assumed that the price distribution at each stage is known a priori. It is also assumed that the feed-in and usage tariffs are the same, i.e., λ_k is the price per unit for both purchase (corresponding to $v_k \geqslant 0$) and sell-back (corresponding to $v_k \leqslant 0$), and there are no transaction costs.

2.2.3. The states

The storage state is characterized by a variable

$$s_k \in [0, \bar{s}] \tag{2}$$

where s_k is the amount of energy stored, and $\bar{s}$ is the upper bound on storage capacity. The state s_k evolves according to:

$$s_{k+1} = \beta s_k + \eta^{\text{in}} v_k^{\text{in}} - \eta^{\text{out}} v_k^{\text{out}} \tag{3}$$

where $\beta \leqslant 1$ is the decay factor, $\eta^{\text{in}} \leqslant 1$ and $\eta^{\text{out}} \geqslant 1$ are charging and discharging efficiency factors. Note that efficiency factors and ramp rates might in general be complicated functions of the operating point, i.e., the storage level, but this paper focuses on an ideal case. The idealized model of the dynamics of storage can be written as:

$$s_{k+1} = s_k + v_k, \qquad v_k \in [-\bar{v}_{\text{out}}, \bar{v}_{\text{in}}] \tag{4}$$

which corresponds to $\beta = 1$, $\eta^{\text{in}} = 1$, and $\eta^{\text{out}} = 1$.

2.2.4. Penalty and salvage value

There is a penalty $h_k(s_k)$ associated with storage, where the sequence of functions $h_k : \bar{\mathbb{R}}_+ \mapsto \bar{\mathbb{R}}_+$ are assumed to be monotonic. Also, a salvage value of $\hat{t} \in [\lambda_N^{\min}, \lambda_N^{\max}]$ is assigned to each unit of energy left in storage by the end of the time-horizon.

2.2.5. The optimization-based model of ideal storage

Since the goal in this paper is to develop a tractable model that effectively highlights the important structural features of the optimal control law and the associated value of storage with an emphasis on the ramp constraint and storage capacity, the idealized model of storage will be adopted; this will allow for focusing on the ramp constraint and storage capacity as the parameters of interest in the model when analyzing the value and price-responsiveness of storage. Nevertheless, in terms of methodology, it would be straightforward to include a "price-adjustment" factor in the formulation to account for injection-withdrawal losses in a similar manner done in [4]. In addition, the piecewise-linear penalty function that has been embedded into the model herein can be used as a surrogate cost to model the dissipation losses associated with keeping energy in storage. Note also that according to [19], the academic state of the art is around 95% efficiency for a battery pack and 93% for the overall system with converters. The industrial state of the art is around 90% efficiency for batteries [20]. It is also assumed that the ramp constraint is symmetric, i.e. $\bar{v}^{\text{in}} = \bar{v}^{\text{out}} = \bar{v}$. The idealized storage management problem can be formulated as a finite-horizon dynamic programming problem as follows:

$$\begin{aligned}
&\min_{v_0,\ldots,v_{N-1}} \mathbf{E}\left[\sum_{k=0}^{N-1}\{h_k(s_k)+\lambda_k v_k\} - \hat{t}s_N\right] \\
&\text{s.t. } s_{k+1} = s_k + v_k \\
&\quad\;\; s_k \in [0, \bar{s}] \\
&\quad\;\; v_k \in [-\bar{v}, \bar{v}] \\
&\quad\;\; \lambda_k \text{ exogenous, and independently distributed according to a PMF} P_k
\end{aligned} \tag{5}$$

Remark 1. The storage problem is first formulated and solved for the finite-horizon case. Later in Section 4.1, the infinite-horizon case is considered and an upper bound on the optimal average value per stage of storage is obtained.

3. Characterization of the optimal policy

In this section, the optimal policy and the value function for problem (5) are characterized based on the principles of stochastic dynamic programming.

Definition 1. Given a sequence of probability mass functions P_k, $k = 0, 1, \ldots, N$, let Θ_k and ψ_k be sequences of maps from the set of all subsets of $\bar{\mathbb{R}}_+$ to $\bar{\mathbb{R}}_+$, defined as follows:

$$\Theta_k : I \mapsto \sum_{\theta \in I} \theta P_k(\theta), \qquad \forall I \subset \bar{\mathbb{R}}_+ \tag{6}$$

$$\psi_k : I \mapsto \sum_{\theta \in I} P_k(\theta), \qquad \forall I \subset \bar{\mathbb{R}}_+. \tag{7}$$

Given $\bar{v} \in \mathbb{R}_+$, and maps Θ_k and ψ_k as defined in (6) and (7), $\Phi_k^{\bar{v}}$ is the map from the set of all subsets of $\bar{\mathbb{R}}_+$ to $\mathbb{R}$ defined according to

$$\Phi_k^{\bar{v}} : I \mapsto \bar{v}(\Theta_k - \rho\psi_k)I, \qquad \forall I \subset \bar{\mathbb{R}}_+, \tag{8}$$

where $\rho = \inf\, I$. For instance, Φ_k^1 maps an interval (a,b) to $(\Theta_k - a\psi_k)(a,b)$.

Theorem 1. *Consider the finite-horizon storage management problem (5) with $\bar{s} = n\,\bar{v}$ for some $n \in \mathbb{Z}_+$. Furthermore, assume that the penalty functions $h_k:[0,\infty) \to [0,\infty)$, $k = 0, \ldots, N$ are piecewise linear non-decreasing convex functions of the form:*

$$h_k(s) = h_k^i s + c_k^i,\ s \in [i\,\bar{v}, (i+1)\,\bar{v}), i \in \bar{\mathbb{Z}}_+ \tag{9}$$

Then, (i) the optimal policy is characterized as follows: if $s_k \in [i\,\bar{v}, (i+1)\,\bar{v})$, for some $i \in \bar{\mathbb{Z}}_+$, then

$$v_k^* = \begin{cases} \max(-s_k, -\overline{v}), & t_{k+1}^{\max(0,i-1)} < \lambda_k \quad \text{for all } i \\ i\overline{v} - s_k, & t_{k+1}^{i} < \lambda_k \leqslant t_{k+1}^{i-1} \quad \text{for } i \geqslant 1 \\ (i+1)\overline{v} - s_k, & t_{k+1}^{i+1} < \lambda_k \leqslant t_{k+1}^{i} \quad \text{for all } i \\ \overline{v}, & \lambda_k \leqslant t_{k+1}^{i+1} \quad \text{for all } i \end{cases} \tag{10}$$

where the thresholds are computed via the following recursive equations:

$$t_N^i = \hat{t}, \quad i \in \{0, 1, 2, \ldots, n-1\}$$
$$t_N^i = -h_N^i, \quad i \geqslant n$$
$$\text{for } k < N:$$
$$t_k^0 = t_{k+1}^1 + \Phi_k^1(t_{k+1}^1, \lambda_k^{\max}] - h_k^0$$
$$t_k^i = t_{k+1}^{i-1} - h_k^i + \Phi_k^1(t_{k+1}^{i+1}, t_{k+1}^{i-1}] + (t_{k+1}^{i+1} - t_{k+1}^{i-1})F_k(t_{k+1}^{i-1}), \quad i \geqslant 1$$

(ii) the value function is a piecewise linear convex function of the form:

$$V_k(s) = -t_k^i s + e_k^i, \quad s \in [i\,\bar{v}, (i+1)\,\bar{v}), \quad i \in \bar{\mathbb{Z}}_+ \tag{12}$$

where t_k^i are the thresholds given in (11) and the intercepts e_k^i are computed via the following recursive equations:

$$e_N^i = 0, \ \ i \in \{0, 1, 2, \ldots, n-1\}$$
$$e_N^i = \bar{s}(t_N^i - \hat{t}), \ \ i \geqslant n$$
$$\text{for } k < N:$$
$$e_k^0 = c_k^0 + e_{k+1}^0 + \left(\bar{v}\,\lambda_k^{\min} + e_{k+1}^1 - e_{k+1}^0 - \bar{v}\,t_{k+1}^1\right)F_k(t_{k+1}^0)$$
$$e_k^i = c_k^i + \bar{v}\,\bar{\lambda} + f(t_{k+1}^{i-1}, t_{k+1}^i, t_{k+1}^{i+1}, e_{k+1}^{i-1}, e_{k+1}^i, e_{k+1}^{i+1}) + g(t_{k+1}^{i-1}, t_{k+1}^i, t_{k+1}^{i+1}), i \geqslant 1 \tag{13}$$

where the functions f and g are given by

$$f(\cdot) = e_{k+1}^{i-1} - \overline{v}t_{k+1}^{i+1} + (e_{k+1}^i - e_{k+1}^{i-1})F_k(t_{k+1}^{i-1}) + (e_{k+1}^{i+1} - e_{k+1}^i)F_k(t_{k+1}^i),$$
$$g(\cdot) = (i+1)\Phi_k^{\overline{v}}[t_{k+1}^{i+1}, t_{k+1}^i] + i\Phi_k^{\overline{v}}[t_{k+1}^i, t_{k+1}^{i-1}] - \Phi_k^{\overline{v}}[t_{k+1}^{i-1}, \lambda_k^{\max}] - \Phi_k^{\overline{v}}[t_{k+1}^{i+1}, \lambda_k^{\max}].$$

Proof 1. Please see this paper's extended e-print on arXiv [21]. □

Fig. 1 shows how the thresholds vary with time and state for the case of a discretized truncated log-normal distribution with mean $\bar{\lambda}_k = 49$ and $\sigma_k = 9$ for all k, i.e. for independently and identically distributed (i.i.d) prices. For generating this plot, $\bar{v} = 1, n = 15, N = 24, \hat{t} = \bar{\lambda}_k$, and no storage penalties (i.e. $h_k^i = 0$ for $i < n$ and all k) are assumed.

Remark 2. The form of the optimal policy shows that if one starts with an empty storage ($s_0 = 0$), then the storage state s_k will only take integer multiples of $\bar{v}$ since $v_k^* \in \{-\bar{v}, 0, \bar{v}\}$ for all k. When $s_0 \neq 0$, the storage state will fall on the grid of integer multiples of $\bar{v}$ immediately after the first time that $v_k^* = i\,\bar{v} - s_k$ or $v_k^* = (i+1)\,\bar{v} - s_k$, and hence, $v_k^* \in \{-\bar{v}, 0, \bar{v}\}$ for the remainder of the time horizon. This conclusion holds for the infinite-horizon case as well and can help simplifying the policy and analysis by focusing on a finite state system.

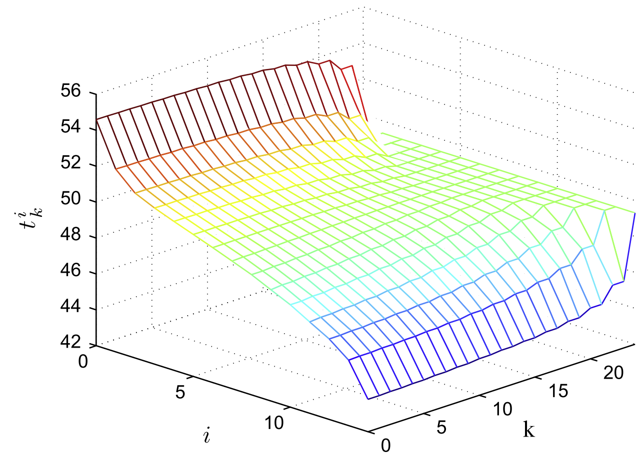

Fig. 1. Thresholds as a function of time and state for i.i.d prices with a discretized truncated log-normal distribution, with mean $\bar{\lambda}_k = 49$ and $\sigma_k = 9, \bar{v} = 1$, $n = 15, N = 24, \hat{t} = \bar{\lambda}_k$ and no storage penalties.

Remark 3. The upper bound $\bar{s}$ on the storage capacity is enforced by choosing h_k^i in (9) sufficiently large (i.e. $h_k^i > \lambda_k^{\max}$) for $i \geqslant n$, so that it would never be optimal to store energy beyond $\bar{s}$. It can be verified that the thresholds and consequently the optimal policy are invariant with respect to the choice of h_k^i for $i \geqslant n$ as long as $h_k^i > \lambda_k^{\max}$.

Definition 2. The economic value of storage, or simply the value of storage, is defined as the negative of the cost of the optimal value of problem (5) divided by the number of stages (N), and is denoted by $\mathcal{V}$ for the finite-horizon case and by $\mathcal{V}_\infty$ for the infinite-horizon case. Therefore, $\mathcal{V} = -V_0(s_0)/N$. For instance, if $s_0 = 0$ (i.e. the consumer starts with an empty storage), it then follows from (12) that for the finite-horizon case:

$$\mathcal{V} = -V_0(0) = -e_0^0/N. \tag{14}$$

Thus, $\mathcal{V}$ can be computed using the recursive equations in (13) for the finite-horizon case.

4. Characterization of the value of storage

4.1. Analytical upper bound on the value of storage

In this subsection, a bound on the long-term value of ramp-constrained storage will be derived. Herein, for the purpose of obtaining the bound, it is further assumed that $h_k^i = 0$, for all $i \leqslant n-1$ and $k \leqslant N$. This assumption is consistent with the objective of finding an upper bound on the value of storage.

Definition 3. Given a control policy $\pi_k : [0,\infty)^2 \mapsto [-\bar{v}, \bar{v}]$, and starting from an arbitrary initial state s, the infinite-horizon average cost per stage associated with problem (5) is defined as

$$\gamma_\pi \overset{\text{def}}{=} \lim_{N\to\infty} \frac{1}{N} \mathbf{E}\left[\sum_{k=0}^{N-1} \lambda_k v_k | s_0 = s\right], \tag{15}$$

where $v_k = \pi_k(x_k, \lambda_k)$. The problem of optimization of γ_π over all feasible stationary policies will be referred to as the infinite-horizon storage management problem. The associated optimal cost will be denoted by γ^*, and $\mathcal{V}_\infty \overset{\text{def}}{=} -\gamma^*$ will be referred to as the long-term expected value of storage.

Remark 4. It is standard to show that the optimal average cost is independent of the initial state s_0. Moreover, if the relative value iteration for the infinite-horizon storage problem converges to some differential cost function $H^*(s)$, then it is necessary for $H^*(s)$ and the optimal average cost per stage γ^* to satisfy the Bellman equation (see, for instance, [22]):

$$H^*(s) = \mathbf{E}\left[\min_{v\in[\max(-s,-\bar{v}),\min(\bar{v},\bar{s}-s)]} \lambda v + H^*(s+v)\right] - \gamma^*. \tag{16}$$

Theorem 2. *Consider the infinite-horizon storage management problem. Suppose that the support of the price distribution function at all stages lies within an interval $[\lambda_{min}, \lambda_{max}] \subseteq [0,\infty)$. All else held constant, the maximum over all possible distributions, of the long-term value of storage is given by*

$$\mathcal{V}_\infty = -\gamma^* = \frac{\bar{v}(\lambda_{\max}-\lambda_{\min})}{2}\frac{n}{n+1} = \frac{(\lambda_{\max}-\lambda_{\min})}{2}\frac{\bar{s}\ \bar{v}}{\bar{s}+\bar{v}}, \tag{17}$$

and is attained when the prices are sampled from a two-point symmetric distribution with nonzero probability masses placed at the endpoints of the fixed support:

$$P_\Lambda(\lambda) = \begin{cases} 1/2 & \text{if} \quad \lambda = \lambda_{\min} \\ 1/2 & \text{if} \quad \lambda = \lambda_{\max} \\ 0 & \text{otherwise} \end{cases}$$

And for any two-point distribution with PMF

$$P_\Lambda(\lambda) = \begin{cases} a & \text{if} \quad \lambda = \lambda_{\max} \\ 1-a & \text{if} \quad \lambda = \lambda_{\min} \\ 0 & \text{otherwise} \end{cases}$$

the following holds:

$$\mathcal{V}_\infty = -\gamma^* = \bar{v}(\lambda_{\max}-\lambda_{min})\frac{b(1+b+\cdots+b^{n-1})}{(b+1)(1+b+\cdots+b^n)} \tag{20}$$

where $b = (1-a)/a$.

Proof 2. Please see this paper's extended e-print on arXiv [21]. □

If in addition to the support of the price distribution, the mean of the distribution is also fixed, a tighter bound can be obtained as stated in Corollary 1:

Corollary 1. *Suppose that in addition to fixing the support of the price distribution function, the mean of the price distribution is also fixed to $\mu \in (\lambda_{min}, \lambda_{max})$. All else held constant, the maximum over all possible distributions, of the long-term expected value of storage is attained when the prices are sampled from a two-point distribution with the following PMF:*

$$P_\Lambda(\lambda) = \begin{cases} \frac{\mu-\lambda_{\min}}{\lambda_{\max}-\lambda_{\min}} & \text{if} \quad \lambda = \lambda_{\max} \\ \frac{\lambda_{\max}-\mu}{\lambda_{\max}-\lambda_{\min}} & \text{if} \quad \lambda = {}_{\min} \\ 0 & \text{otherwise} \end{cases}$$

The corresponding long-term value of storage is obtained by plugging $b = (\lambda_{max}-\mu)/(\mu-\lambda_{min})$ into (20).

Proof 3. Please see this paper's extended e-print on arXiv [21]. □

Remark 5. The $n/(n+1)$ scaling in the optimal average cost per stage implies that 90% of the maximum possible value of storage is achieved when the storage capacity is only 9 times the ramp constraint. Note that this was obtained for an extreme distribution. As it will be seen in the remainder of this section, for less extreme distributions with smaller variance over the support, the value of storage saturates even more quickly. This includes empirical distributions obtained from electricity market data. Furthermore, aging, dissipation, and non-ideal charging and discharging factors further reduce the value of storage.

4.2. Computational experiments for characterization of the value of storage

In this subsection, numerical computations are employed to characterize the value of the proposed model of storage over a finite time-horizon and highlight the effects of ramp constraints and price volatility on the value of storage. Herein, the following classes of distributions are considered:

- Discretized truncated log-normal distribution, with mean $\bar{\lambda} = 50$,
- Discretized uniform distribution, with mean $\bar{\lambda} = 50$.

The reason for choosing the log-normal distribution is that the empirical distributions from ISONE and PJM qualitatively resemble a log-normal distribution, at least for the cases tested in this paper. The choice of the mean ($\bar{\lambda} = 50$) is also a realistic choice for average hourly energy prices in markets such as PJM and ISONE. For the purpose of these computations, the same price distribution is used for all k (i.e. it is assumed that prices are independently and identically distributed). Throughout this section it is assumed that $s_0 = 0$, which means that the consumer starts with an empty storage. For each of these distributions, all quantities are fixed in the model other than σ, the standard deviation of price distribution, and n, the ratio of storage capacity $\bar{s}$ to physical ramp constraint of storage $\bar{v}$. The quantity $n = \bar{s}\ /\ \bar{v}$ is varied by fixing $\bar{v}$ and changing $\bar{s}$. Using the fixed quantities $N = 24$, $\bar{v} = 10$, and $\bar{\lambda} = 50$, how $\mathcal{V}$ varies as a function of σ and n is examined. For the purpose of these simulations, $\hat{t}$ is set equal to the mean of the price distribution. Herein, $h_k^i = 0$ for all $i \leqslant n-1$ and $k \leqslant N$, so that there is no penalty on storing energy up to capacity. Then, for a fixed time horizon, how $\mathcal{V}$ varies with σ and n is examined for each of the following price distributions:

4.2.1. Discretized truncated log-normal distribution

Fig. 2 illustrates how $\mathcal{V}$ changes with σ and n, for the discretized truncated log-normal distribution. The plots show that the value of storage increases linearly with σ. As one would expect, the value of storage also increases as the storage capacity increases. However, it is interesting to note that for a fixed standard deviation, the value of storage saturates fairly quickly as a function of n. Hence, for a given time horizon, a fixed ramp constraint, and a fixed σ, there exists a certain range for capacity beyond which the value of storage will no longer change noticeably. Also, the optimal storage capacity increases with price volatility.

4.2.2. Discretized uniform distribution

As can be seen in Fig. 3, saturation of the value of storage occurs at about the same value as in the log-normal case. However, for certain extreme distributions, such as an asymmetric 2-point distribution, saturation can occur more quickly. Note also that the value of storage is a linear function of the standard deviation, just like the log-normal case.

One interesting observation in these results is that in the presence of ramp constraints, several distributed storage systems would be more profitable than one large storage system of equal ramp constraint and aggregate capacity, due to the quick saturation of $\mathcal{V}$ as n increases. Although this observation is based on the assumption that the ramp constraint and capacity are independent, this assumption might actually be valid for the case of

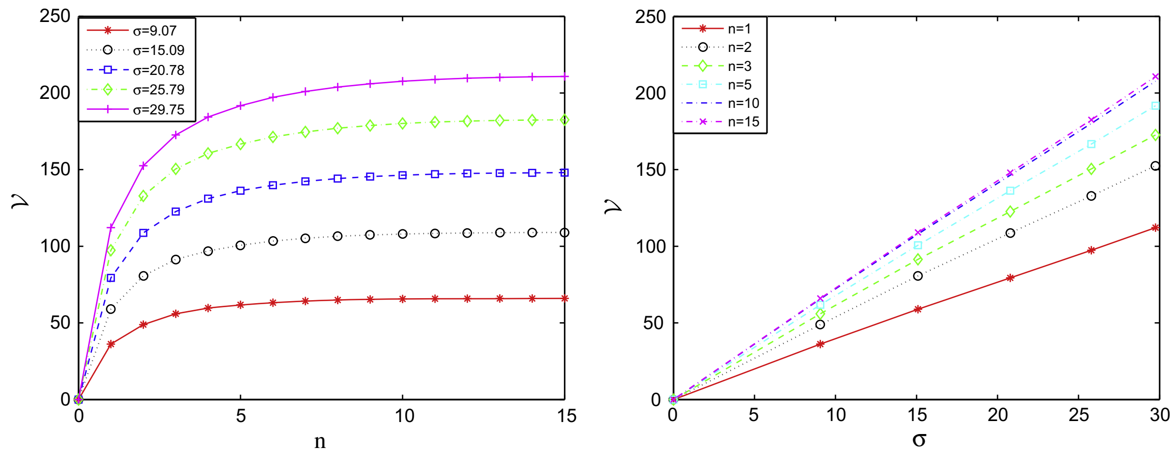

Fig. 2. $\mathcal{V}$ vs. n (left) and σ (right) for a few samples, using a discretized truncated log-normal distribution.

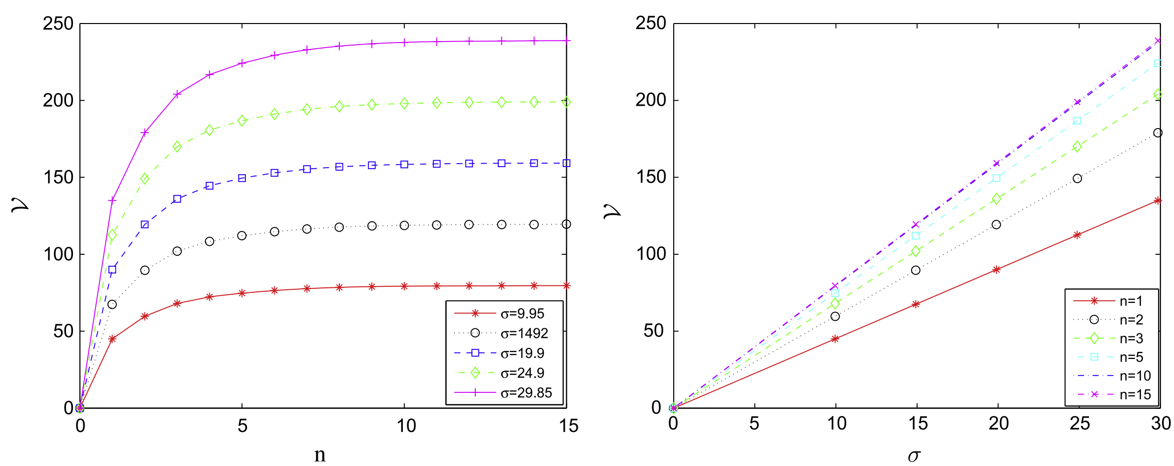

Fig. 3. $\mathcal{V}$ vs. n (left) and σ (right) for a few samples, using a discretized uniform distribution.

distribution grids in which the ramp constraint is imposed by the power lines. Another interesting observation is that although the shape of the plots look quite similar for both distributions, the value of the uniform distribution is higher than that of the log-normal. This makes intuitive sense because with the uniform distribution, on average, the consumer has as many opportunities for buying at a low price as there are opportunities for selling at a high price, while for the log-normal case, most of the probability mass is centered around the mode, which creates fewer opportunities for arbitrage in expectation.

5. Implications of the optimal policy for price-responsiveness of storage

Herein, the focus is on understanding the price-dependence and state-dependence of the optimal policy, and the impact of the state-price interplay on the storage response. The state-dependence of storage response is first averaged out in order to focus on the "average" price-responsiveness only, which allows for addressing the expected price elasticity of demand induced by storage in the first subsection. Then, in the next subsection, the interplay between state-dependence and price-dependence of storage response is explored.

5.1. Average price elasticity of demand induced by storage

In this subsection, in order to study the implications of the optimal policy on price elasticity of demand (PED) in an energy market, a computational framework for studying the average PED in a simulated energy market is presented. In the dynamic model studied in this paper, the storage response depends on price, stage, state, time-horizon, storage capacity, and ramp constraint. The term "average PED" reflects the fact that the dependence of storage response on the stage and the internal state of storage have been averaged out. It is assumed that there is a fixed time horizon N, and state-dependence is eliminated by taking expectations. In particular:

$$v(k,\lambda) = \mathbf{E}_{s_0,s_k}[v_k^*|\lambda_k = \lambda].$$

In order to eliminate stage-dependence, one can think of the storage response-measuring observer as sampling a random time τ uniformly over $\{0,\cdots,N\}$. By averaging over this randomness, dependence is maintained on the prices only:

$$v_{avg}(\lambda) = \mathbf{E}_\tau[v(\tau,\lambda)],$$

which is captured in the simulations by clustering real-time prices, and averaging over each cluster.

In these numerical simulations, the average is taken over random instances of prices and storage initial states. It is assumed that $N = 288$, which corresponds to a period of 24 h, where real-time prices are updated once in every 5 min. The storage system implements the optimal policy given in Theorem 1. For the purpose of these computations, the same price distribution is used for all k (i.e. it is assumed that prices are independently and identically distributed). For generating random price_sequences, a discretized truncated log-normal distribution with $\bar{\lambda} = 52$ and $\sigma = 22$ is used, because the log-normal distribution qualitatively resembles the empirical distribution of prices from ISONE and PJM, as mentioned in the previous section. Also, a mean of 52 and standard deviation of 22 are realistic choices for real-time energy prices in markets such as PJM and ISONE on a day with moderate volatility. Based on the results in Section 4.2, for these model parameters, a storage capacity of $\bar{s} = 5\,\bar{v}$ is a reasonable choice for all consumers. Just like all previous simulations, $\hat{t}$ is set equal to $\bar{\lambda}$. Also, $\bar{v}$ is set equal to 10, and $h_k^i = 0$ for all $k \leqslant N$ and $i < n$.

Fig. 4 illustrates how the average storage response changes as a function of price.

As the plot suggests, the expected demand seems to be considerably more responsive to changes in the prices that fall in the mid-portion of the price range. This portion serves as a steep transition region, in which the policy quickly switches from the "buy" policy to the "sell" policy.

To characterize price elasticity, first recall the standard definition of PED:

$$PED = \frac{\Delta d/d}{\Delta\lambda/\lambda} \tag{21}$$

where d denotes demand. To characterize PED more accurately, one needs to bear in mind that the overall PED should have the firm component of demand in it. The firm component of demand is the amount of energy purchased by the consumers from the grid to satisfy any demand other than the energy they store in their storage devices; note that the amount of the firm demand does not depend on the prices. Hence, in this setup, $d = d^f + v_{avg}(\lambda)$, where d^f denotes the firm component of demand. It can be observed in Fig. 4 that the average PED is almost zero for prices that are considerably larger or smaller than the mean price, and only in the mid-portion of the plot (i.e. around the mean price), a substantial average PED can be noticed. One can verify using Eq. (21) with $d = d^f + v_{avg}(\lambda)$ that the average PED (i.e. the PED computed using the average storage response) depends on how the average storage response compares with the firm demand. Table 1 shows the average PED around the mean price, using different values for d^f.

5.2. State-dependent price-responsiveness of a storage system

In this subsection, a computational framework is presented for understanding the behavior of storage as a function of price and the amount of stored energy, and for characterization of the buy/sell phase transition region in the price-state plane. In order to eliminate stage dependence, the infinite-horizon version of the storage problem (5) will be considered and policy iteration (see, e.g., [22]) will be performed to numerically obtain a stationary

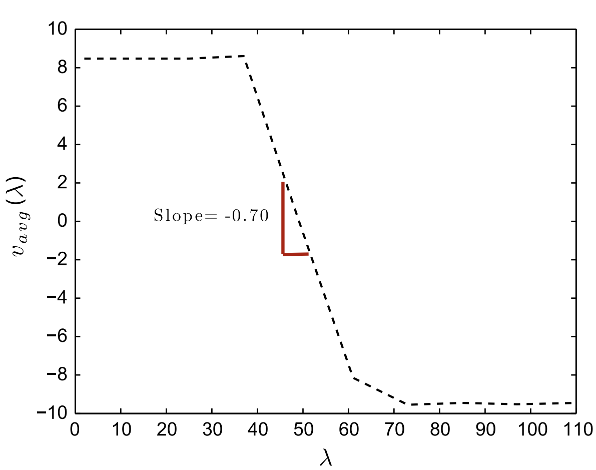

Fig. 4. $v_{avg}(\lambda)$ vs. λ, using the discretized log-normal price distribution.

(stage-independent) policy for purchase/sale as a function of both state and price. This section provides a qualitative picture of the structural characteristics of the behavior of storage, and a framework for estimating the PED as a function of the state. Herein, it is assumed that the consumer starts with an empty storage, implying that the states would only take on integer multiples of the ramp constraint. The same price distribution is also used for all stages (i.e. the prices are assumed to be i.i.d). In these computations, a discretized truncated log-normal price distribution is first used, and then the results are compared against the case of a discretized uniform distribution. For both distributions, a mean of about $\bar{\lambda} = 50$ and a standard deviation of about $\sigma = 30$ are used. Both distributions also have the same support. Furthermore, $\bar{v} = 1$ and $n = 10$. Fig. 5 illustrates how the storage response varies with price for three cases of the state (when the storage is empty ($i = 0$), when the storage is half full ($i = n/2$), and when the storage is nearly full ($i = n - 1$)).

Note in Fig. 5 the considerable effect of the storage state on storage response, compared to the small effect of the price distribution. More specifically, for both distributions, when the storage is empty, the optimal policy recommends purchasing from the grid even when the prices are somewhat above the mean price. Note that for the log-normal case, this policy change occurs at a slightly lower price because of the left skewness of the log-normal distribution. Though, when the storage is half full, the optimal decision switches from the "buy-it-all" policy to the "sell-it-all" policy right at the mean price for the uniform distribution; for the log-normal distribution, this policy change occurs slightly before the mean price, which is again due to the left skewness of the log-normal distribution. Finally, when the storage is nearly full, for both distributions the optimal policy is to sell as much energy as possible for most prices, and to do nothing for the low prices.

The transition points in the infinite-horizon policy from sell-it-all to buy-it-all on the $s - \lambda$ plane are shown in Fig. 6. Any point to the left of and/or below the transition points is a buying policy, which corresponds to $v^*(s, \lambda) = \bar{v}$, and any point to the right of and/or above the transition points is a selling policy, which corresponds to $v^*(s, \lambda) = -\bar{v}$. The plus signs show a direct transition from buying to selling when moving along the vertical axis, i.e., as storage state varies, unless they are immediately followed by a star on their right. The stars denote a transition through a "Do Nothing" policy when moving along the horizontal axis, i.e., as price varies. Therefore, at the prices denoted by $*$ the optimal policy is $v^* = 0$.

Fig. 6 clearly illustrates the interplay between the state-dependence and the price-dependence of the storage response. Note also in Fig. 6 the small effect of the price distribution on the storage response. The optimal policy for the log-normal case is slightly more shifted to left compared to that of the uniform case, which is again caused by the left skewness of the log-normal distribution.

The above results can now be used to characterize the PED as a function of the storage state. In order to compute the overall PED using (21), once again the demand (d) is quantified such that it includes the firm component of the demand as well: $d = d^f + v^*(s, \lambda)$. Note that PED $(s) = 0$ for all the points in the "Buy Region" and in the "Sell Region" because the storage response is constant in those regions. However, around the transition curve, the PED is non-zero

Table 1
Average price elasticity of demand around the mean price using different values for d^f.

d^f	Average price elasticity
$\bar{v}$	−3.6
$3\,\bar{v}$	−1.2
$8\,\bar{v}$	−0.45

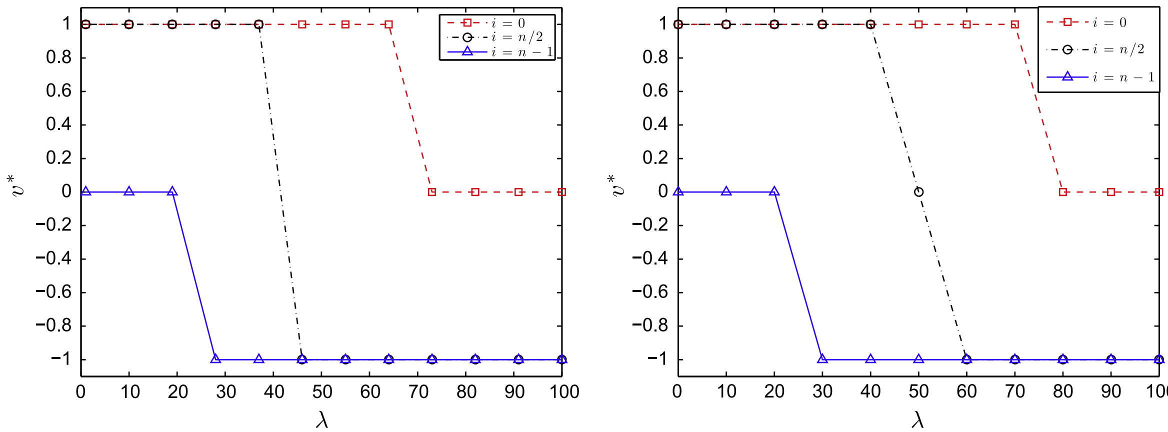

Fig. 5. Storage response vs. price for three sample states for discretized truncated log-normal (left) and discretized uniform (right) price distributions, both with mean 50.

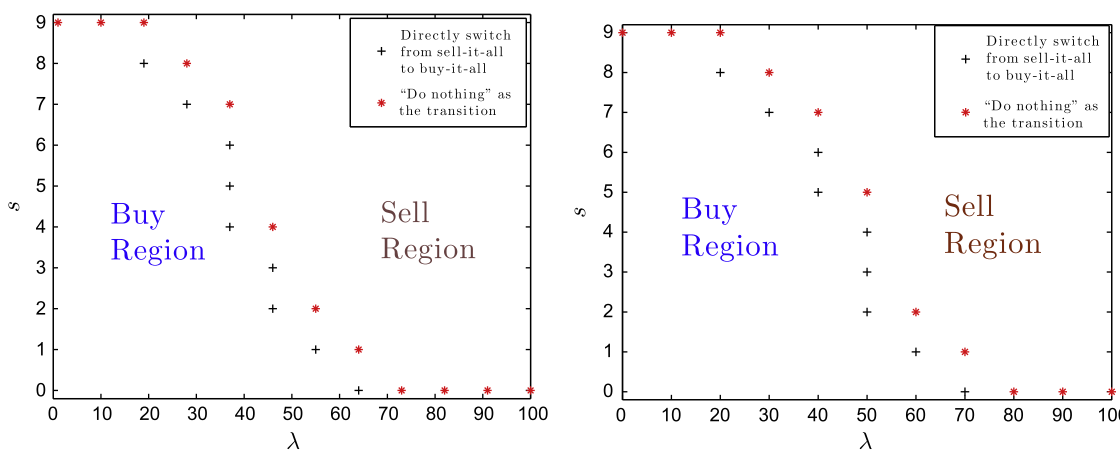

Fig. 6. Occurrence points of policy change from sell-it-all to buy-it-all in the $s - \lambda$ plane, for discretized truncated log-normal (left) and discretized uniform (right) price distributions.

because $\Delta d < 0$ in that region. But Δd can only take two values: $\Delta d = -2\,\bar{v}$ when there is a direct transition from the "Buy Region" to the "Sell Region", and $\Delta d = -\,\bar{v}$ when the transition is through a "Do Nothing" policy. Hence,

$$PED(s) = -\frac{2\,\bar{v}\,\lambda}{(d^f + v^*(s,\lambda))\Delta\lambda} \quad \text{and} \quad PED(s) = -\frac{\bar{v}\,\lambda}{(d^f + v^*(s,\lambda))\Delta\lambda}$$

around the points denoted by + and $*$, respectively, where, depending on the point at which the PED shall be computed, v^* takes on one of the values in $\{-\bar{v}, 0, \bar{v}\}$.

Remark 6. The effects of inefficiencies such as conversion losses and aging on the optimal policy and price-responsiveness of storage are not directly considered in this paper. While such inefficiencies certainly limit the economic value of storage, an important question to ask is what is their effect on the response of storage to prices? For the case of battery storage, aging is proportional to the amount of current withdrawn or injected. Intuitively, this would make buying energy for, or selling energy from storage less profitable at moderate prices. Within the class of threshold policies, the corresponding optimal selling thresholds would be higher and buying thresholds would be lower than what has been derived in this paper and the "do nothing" range would be wider. Qualitatively, this would mean an overall choppier response from storage, with high elasticity over a narrow range of prices and low elasticity over a wide range, which would be undesirable from a system operation and reliability perspective.

6. Conclusion

In this paper, a dynamic model for optimal control of storage under ramp constraints and exogenous, stochastic prices was proposed. The associated optimal policy and value function were derived, and explicit formulas for their computation were given. Moreover, an analytical upper bound on the long-term average value of storage was derived, which can be useful in assessing viability of investment in storage. Also, computational experiments were presented for characterizing the value of storage, and important implications of the associated optimal policy for price-responsiveness of storage were studied. In particular, it was shown that regardless of price volatility, the value of storage saturates quickly as the storage capacity increases. Furthermore, the non-trivial interplay between state-dependence and price-dependence of the optimal storage response was highlighted, and it was shown that on average, a considerable amount of price elasticity is expected near the mean price. These results provide insight into learning the behavior of storage, particularly modeling and estimating the value and response of a ramp-constrained storage system when used as an arbitrage mechanism.

While this paper largely focused on the economic value of storage, it is important to recognize and quantify the environmental and the reliability value of storage. With proper control policies, storage can help matching stochastic supply with demand, improving system frequency and voltage profiles, and possibly mitigating large blackouts. The development of a systematic framework for quantifying the value of storage, the trade-offs between reliability, environmental, and economic value of storage, and design of real-time pricing and market mechanisms for optimally striking these trade-offs are important directions for future research.

Acknowledgements

Funding Sources: National Science Foundation Graduate Research Fellowship; MIT Energy Initiative Seed Fund; Draper Laboratories. The funding sources were not involved in the study design, in the collection, analysis and interpretation of data, in the writing of the report, and in the decision to submit the article for publication.

Appendix A. An evaluation of the optimal policy

Herein, to examine the results of this paper and verify the validity of its assumptions, the optimal policy is tested against actual market price data. The finite-horizon optimal policy (10) is applied to real-time wholesale market data taken from PJM [23] and ISONE [24] to examine the competitive ratio (CR) of the policy when applied to real-world price data. This would allow for assessing whether the assumption of independent prices is actually reasonable for the storage problem.

Definition A.1. The *empirical value* of storage is defined as the negative of the cost of the optimal value of problem (5) given a price sequence, divided by the number of stages (N), and is denoted by $\mathcal{V}$.

Definition A.2. The *competitive ratio* (CR) is defined as the ratio of the empirical value obtained from the optimal policy to the absolute maximum empirical value that would have been obtained deterministically, had the entire price sequence been known a priori.

A relatively high CR would not necessarily suggest that prices are independent; rather, it could mean that the sensitivity of the optimal policy to the correlations, if any, in real-world prices is low. If the CR obtained form the assumption of independent prices is relatively high, one can suggest that the additional information that the current price provides on future prices (as a conditional distribution) has little value.

The setup of the experiment is as follows. Actual data is taken for hourly energy prices (for 16 h of each day, from 8 a.m. to 12 a.m.) for two different months (December 2010 and July 2011) from PJM, and May and November 2011 from ISONE. The choice of these dates and times was arbitrary. For the purpose of these simulations, $\hat{t}$, the salvage value, is set equal to $\bar{\lambda}$, the empirical mean. To perform the simulations, the empirical distribution of the data for each case is first found. In these computations, the same price distribution is used for all k; i.e., all the price data for the entire month is taken, and this data is used to find the empirical distribution of prices in that month so that the thresholds for the optimal policy for all hours are computed from the same empirical distribution. In other words, it is assumed that prices are independently and identically distributed (i.i.d). Then, at the beginning of each stage the actual price of that stage is revealed to the optimal policy and the decision is recorded. The profit resulting from the optimal policy can then be computed using these recorded decisions. Next, for the purpose of comparison, it is assumed that the entire price sequence is perfectly known a priori, and the value that results from deterministically and omnisciently maximizing the profit against the materialized prices is computed. This deterministic value is the absolute best that an omniscient agent could have done, and its corresponding deterministic policy shall be referred to as the omniscient policy. Then, the CR is found by computing the ratio of the value obtained from the optimal policy to the value obtained from the omniscient policy. The CR gives a measure of how well the optimal policy has performed.

For each month, there is a number of days in which the average price is well above the average price of the entire month. These days are outliers in the sense that the empirical distribution is way off for modeling their price sequence. This issue is addressed by performing three sets of simulations. For the first set of simulations, only those days in which the average price is within one standard deviation of the average price of all the selected days in the month are chosen. In the second set of simulations, those days in which the average price is within about 1.5 times the standard deviation of the average price of all the selected days in the month are chosen. Then, in the third set of simulations, all days of the month are taken into account, even the outliers. The CR is recorded for each day, and the average CR for each month and each case is reported in Table A.1. Comparing the results in Table A.1 reveals how much these outlier days affect the CR. Note that a ramp constraint of $\bar{v} = 1$ and a storage capacity of $\bar{s} = 10$ is used in all the simulations, and it is assumed that there are no penalties on storage.

Fig. A.1 shows the plots of the empirical value of storage for each day of the month in the second set of simulations (i.e. for those days whose average price is within 1.5 times the standard deviation of the average price of all the selected days in the month). As it can be seen in Fig. A.1, nearly perfect matching is obtained in some days, while in some other days there is discrepancy. This discrepancy appears to be mainly due to two reasons. The first reason seems to be that multiple spikes exist in some days with only a few (or no) prices that are below the buying thresholds. So, even though there are ample opportunities for arbitrage, even the lowest of these spiky prices is not within the normal range. In Table A.1, this effect can be observed for the month of December in PJM, in which the removal of those days with high average prices improved the CR from 0.77 to 0.88. The second reason seems to be that in some days all the prices are almost in the same range compared to the thresholds (i.e. they are either mostly above the buying thresholds or mostly below the selling thresholds). In other words, the low CR is just an outcome of an undesirable sequence of prices (sample path). So, even though for the deterministic case with the price sequence known a priori it is possible to take advantage of these small price differentials, the thresholds are unable to capture these opportunities. This effect can be observed in the month of May in Table A.1, in which the removal of those days with high average prices did not really improve the CR; this is because the average price in the days with a relatively flat price sequence is not necessarily considerably higher than the month's average. So far, the thresholds have been computed for the optimal policy

Table A.1
Competitive Ratio for each set of simulations.

Source/month	Comp ratio	# Of days used	Range from mean
PJM/July	0.86	27	$1 \times \sigma$
	0.84	28	$1.5 \times \sigma$
	0.81	All	All
PJM/December	0.88	22	$1 \times \sigma$
	0.87	25	$1.5 \times \sigma$
	0.77	All	All
ISONE/May	0.72	25	$1 \times \sigma$
	0.71	All	$1.5 \times \sigma$
	0.71	All	All
ISONE/November	0.89	29	$1 \times \sigma$
	0.88	All	$1.5 \times \sigma$
	0.88	All	All

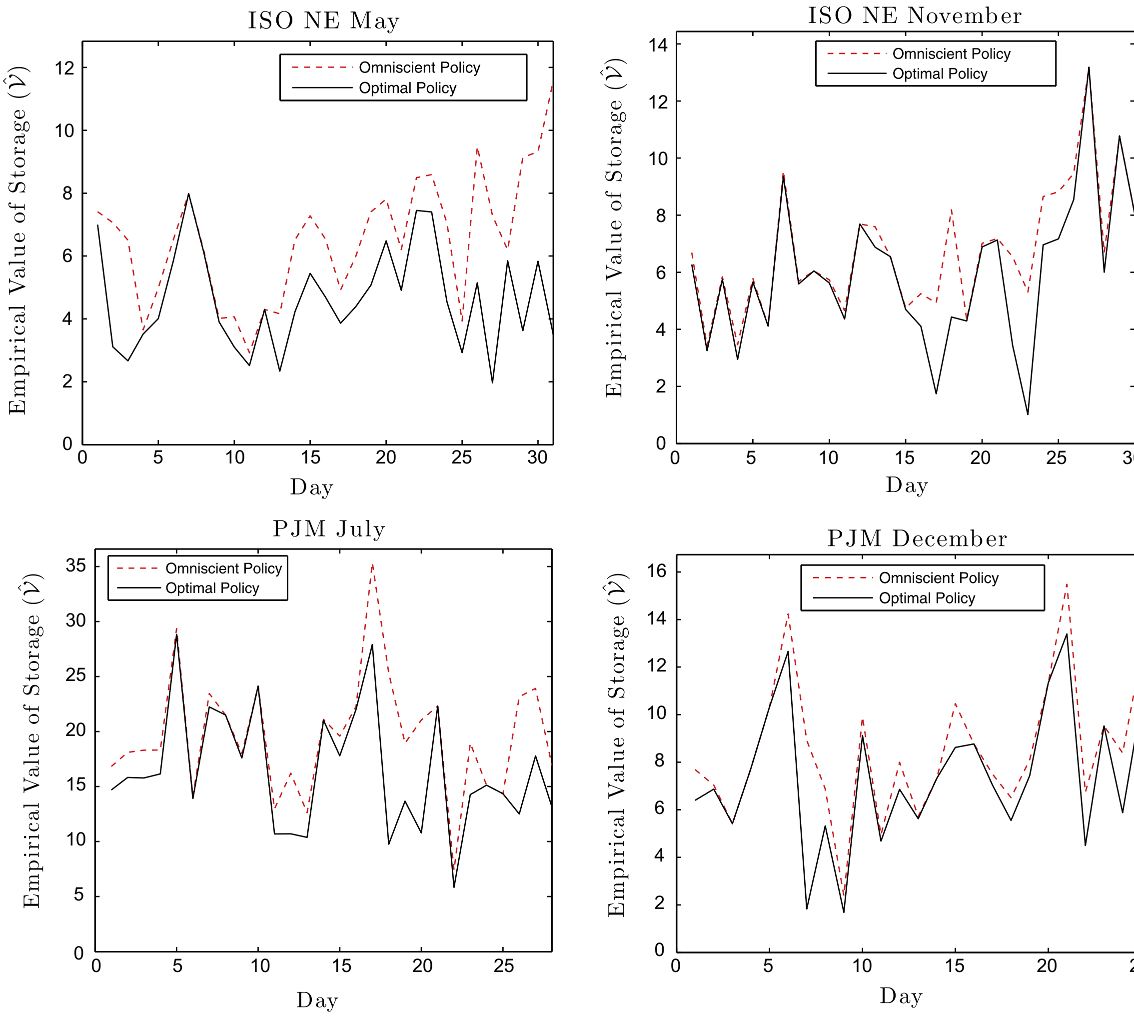

Fig. A.1. The empirical value of storage for each day of the month in the second set of simulations, obtained from applying the optimal policy (solid line) and the omniscient policy (dashed line).

using the empirical price distributions from the prices in that same month. This can be taken as a proxy for the sensitivity of the value of storage to correlations in the actual prices, given that the optimal policy assumes independent prices. Although there is no benchmark to compare with, the CR obtained from applying the optimal policy seems reasonably high, and it suggests that the sensitivity of the value of storage to the assumption of independent prices is low. Also, although the optimal policy does not do as well on those days with higher than normal, or very flat price profiles, it does not appear that a Markovian or Martingale assumption on prices could do any better, because these prices are outliers and do not seem to follow a structured stochastic pattern that can be learned from past data. However, this needs to be substantiated by further studies and systematic experiments.

In the next set of experiments, instead of computing the empirical distribution from the data of that same month, the empirical distributions are computed from the price data of the past 30 days, the past 20 days, and the past 10 days, respectively. The results are shown in Table A.2.

An interesting observation is that for both months of May and November in ISONE, using the empirical distribution from the past 20 days gives almost the same CR as using the empirical distribution from May and November themselves (as reported in Table A.1). However, comparing the competitive ratios shown in Table A.2 for both months in PJM with the results shown in Table A.1, it can be observed that for PJM, using the empirical distribution from historical data does not do as well as the empirical distribution from that month itself.

Table A.2
Competitive ratios using the empirical distribution from the past 30, 20, and 10 days.

Source/Month	Past 30 days	Past 20 days	Past 10 days
PJM/July	0.68	0.70	0.75
PJM/December	0.55	0.65	0.67
ISONE/May	0.69	0.70	0.68
ISONE/November	0.91	0.89	0.89

References

[1] Bellarmine GT, Turner MC. Energy conservation and management in the US. Energy Convers Manage 1994;35:363–73.
[2] Roozbehani M, Dahleh MA, Mitter S. Volatility of power grids under real-time pricing. IEEE Trans Power Syst 2012;27:1926–40.
[3] Rempala R. Optimal strategy in a trading problem with stochastic prices. Syst Model Optim Lect Notes Control Inf Sci 1994;197:560–6.
[4] Secomandi N. Optimal commodity trading with a capacitated storage asset. Manage Sci 2010;56:449–67.
[5] Kaminski V, Feng Y, Pang Z. Value, trading strategies and financial investment of natural gas storage assets. In: Northern Finance Association Conf. Kananaskis, Canada; 2008.
[6] Devalkar S, Anupindi R, Sinha A. Integrated optimization of procurement, processing and trade of commodities. Oper Res 2011;59:1369–81.

[7] Kondoh J, Ishii I, Yamaguchi H, Murata A, Otani K, Sakuta K. Electrical energy storage systems for energy networks. Energy Convers Manage 2000;41:1863–74.
[8] Jenkins DP, Fletcher J, Kane D. Model for evaluating impact of battery storage on microgeneration systems in dwellings. Energy Convers Manage 2008;49:2413–24.
[9] Lund H, Salgi G. The role of compressed air energy storage (CAES) in future sustainable energy systems. Energy Conversi Manage 2009;50:1172–9.
[10] Ekman CK, Jensen SH. Prospects for large scale electricity storage in Denmark. Energy Convers Manage 2010;51:1140–7.
[11] Kanakasabapathy P, Shanti Swarup K. Bidding strategy for pumped-storage plant in pool-based electricity market. Energy Convers Manage 2010;51:572–9.
[12] Korpaas M, Holen AT, Hildrum R. Operation and sizing of energy storage for wind power plants in a market system. Int J Electrical Power Energy Syst 2003;25:599–606.
[13] Kim JH, Powell WB. Optimal energy commitments with storage and intermittent supply. Oper Res 2011;59:1347–60.
[14] Harsha P, Dahleh MA. Optimal sizing of energy storage for efficient integration of renewable energy. In: Proc of 50th IEEE conf on decision and control. Orlando, Fl; 2011.
[15] Lee TY, Chen N. Determination of optimal contract capacities and optimal sizes of battery energy storage systems for time-of-use rates industrial customers. IEEE Trans Energy Convers 1995;10:562–8.
[16] Faghih A, Roozbehani M, Dahleh MA. Optimal utilization of storage and the induced price elasticity of demand in the presence of ramp constraints. In: Proceedings of the 50th IEEE conference on decision and control. Orlando, Fl; 2011.
[17] Kirschen DS, Strbac G, Cumperayot P, de Paiva Mendes D. Factoring the elasticity of demand in electricity prices. IEEE Trans Power Syst 2000;15:612–7.
[18] Kirschen DS. Demand-side view of electricity markets. IEEE Trans Power Syst 2003;18:520–7.
[19] Qian H, Zhang J, Lai JS, Yu W. A high-efficiency grid-tie battery energy storage system. IEEE Trans Power Electron 2011;26:886–96.
[20] A123 Systems. Smart grid stabilization and smart energy storage solutions. <www.a123systems.com> [retrieved April 2012.
[21] Faghih A, Roozbehani M, Dahleh MA. On the economic value and price-responsiveness of ramp-constrained storage. e-print; 2012. <http://arxiv.org/pdf/1211.1696v1.pdf>.
[22] Bertsekas DP. Dynamic programming and opt control, vols. I & II, 2nd ed. Belmont: Athena Scientific; 2000.
[23] PJM Interconnection website. Data Dictionary: Monthly real-time locational marginal pricing. <www.pjm.com> [retrieved December 2011].
[24] ISO New England website. Historical hourly price data. <http://www.iso-ne.com> [retrieved December 2011].