



Heriot-Watt University  
Research Gateway

# The Optimal Control of Storage for Arbitrage and Buffering, with Energy Applications

## Citation for published version:

Cruise, J & Zachary, S 2018, The Optimal Control of Storage for Arbitrage and Buffering, with Energy Applications. in M Mougeot, D Picard, P Tankov, R Plougonven & P Drobinski (eds), *Renewable Energy: Forecasting and Risk Management 2017*. Springer Proceedings in Mathematics & Statistics, Springer, pp. 209-227, Workshop on Forecasting and Risk Management for Renewable Energy 2017, Paris, France, 7/06/17. [https://doi.org/10.1007/978-3-319-99052-1\\_11](https://doi.org/10.1007/978-3-319-99052-1_11)

## Digital Object Identifier (DOI):

[10.1007/978-3-319-99052-1\\_11](https://doi.org/10.1007/978-3-319-99052-1_11)

## Link:

[Link to publication record in Heriot-Watt Research Portal](#)

## Document Version:

Peer reviewed version

## Published In:

Renewable Energy

## Publisher Rights Statement:

Copyright © 2018, Springer Nature Switzerland AG

## General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

## Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [open.access@hw.ac.uk](mailto:open.access@hw.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# The optimal control of storage for arbitrage and buffering, with energy applications

James Cruise\* and Stan Zachary\*

May 24, 2018

## Abstract

We study the optimal control of storage which is used for both arbitrage and buffering against unexpected events (shocks), with particular applications to the control of energy systems in a stochastic and typically time-heterogeneous environment. Our philosophy is that of viewing the problem as being formally one of stochastic dynamic programming (SDP), but of recasting the SDP recursion in terms of functions which, if known, would reduce the associated optimisation problem to one which is deterministic, except that it must be re-solved at times when shocks occur. In the case of a perfectly efficient store facing linear buying and selling costs the functions required for this approach may be determined exactly; otherwise they may typically be estimated to good approximation. We provide characterisations of optimal control policies.

We consider also the associated deterministic optimisation problem, outlining an approach to its solution which is both computationally tractable and—through the identification of a running forecast horizon—suitable for the management of systems over indefinitely extended periods of time.

We give examples based on Great Britain electricity price data.

## 1 Introduction

How should one optimally control storage which is used simultaneously for a number of different purposes? We study this problem in the case of a single store which is used for both price arbitrage, i.e. for buying and selling over time, and for buffering against unexpected events, or *shocks*. Here an optimal control must balance the sometimes conflicting controls which would apply to these two uses of storage considered individually. Of particular interest is the control of an energy store in a stochastic and typically time-heterogeneous environment, where at any time a full stochastic description of that environment may not be available over more than a relatively short future time horizon. The shocks correspond, for example, to the loss of a generator or transmission line, or a sudden surge in demand. Our philosophy is that of viewing the problem as being formally one of stochastic dynamic programming (SDP), but of recasting the SDP recursion in terms of functions which may be determined in advance, either exactly or approximately, and which reduce the associated optimisation problem to one which is deterministic, except that it must be re-solved at those times at which shocks occur.

There is considerable literature on the control of storage for each of the above two purposes considered on its own. There have been numerous studies of the use of storage for buffering

---

\*Heriot-Watt University. Research supported by EPSRC grant EP/I017054/1

against both the increased variability and the increased uncertainty in electrical power systems due to the higher penetration of renewable generation—the former due to the natural variability of such resources as wind power, and the latter due to the inherent uncertainty of forecasting. These studies have considered many different more detailed objectives; these range from the sizing and control of storage facilities co-located with the renewable generation so as to provide a smoother supply and so offset the need for network reinforcement [1, 2, 3], to studies on storage embedded within transmission networks so as to increase wind power utilisation and so reduce overall generation costs [4, 5, 6]. In addition there have been a number of studies into the more general use of storage for buffering, for example, so as to provide fast frequency response to power networks [7, 8, 9], or to provide quality of service as part of a microgrid [10, 11].

In the case of the use of storage for arbitrage, and with linear cost functions for buying and selling at each instant in time, the problem of optimal control is the classical *warehouse problem* (see [12, 13, 14] and also [15] for a more recent example). Cruise et al [16] consider the optimal control of storage in the case where the store is a price-maker (i.e. the size of the store is sufficiently large that its activities influence prices in the market in which it operates) and is subject to both capacity and rate constraints; they develop the associated Lagrangian theory, and further show that the optimal control at any point in time usually depends only on the cost functions associated with a short future time horizon. Recent alternative approaches for studying the value and use of storage for arbitrage can be found in the papers [17, 18, 19, 20, 21]—see also the text [22], and the further references given in [16]. For an assessment of the potential value of energy storage in the UK electricity system see [9].

In general the problem of using a store for buffering is necessarily stochastic. The natural mathematical approach is via stochastic dynamic programming. This, however, is liable to be computationally intractable, especially in the case of long time horizons and the likely time heterogeneity of the stochastic processes involved. Therefore much of the literature considers necessarily somewhat heuristic but nevertheless plausible control policies—again often adapted to meeting a wide variety of objectives. For example, for large stores operating within transmission networks, the buffering policies studied have included that of a fixed target level policy [23], a dynamic target level policy [24], and a two-stage process with day ahead generation scheduling and an online procedure to adapt load levels [25].

Control policies have been studied via a range of analytic and simulation-based methods. Examples of an analytic approach can be found in [26], where partial differential equations are utilised to model the behaviour and control of a store, and in [27, 28], where spectral analysis of wind and load data is used with models which also incorporate turbine behaviour. Simulation-based studies include [23, 24], which use a bootstrap approach based on real wind forecast error data, and [25], which uses Monte Carlo simulation of the network state.

In the present paper we use an economic framework to study the optimal control of a store which, as previously stated, is used both for price arbitrage and for buffering against occasional and unpredictable *shocks* whose occurrence is described by some stochastic process. The store seeks to operate in such a way as to minimise over time the expected total cost of its operation. We believe such an economic framework to be natural when the store operates as part of some larger and perhaps very complex system, provided the price signals under which the store operates are correctly chosen—see, for example, [29]. The store may be sufficiently large as to have market impact, leading to nonlinear cost

functions for buying and selling, may be subject to rate (as well as capacity) constraints, and, as will typically be the case, may suffer from round-trip inefficiencies. We formulate a stochastic model which is realistic in many circumstances and characterise some of the properties of an optimal control, relating the results to the existing experimental literature. Our approach is that of re-expressing the traditional SDP recursion so as to reduce the associated optimisation problem to one which is deterministic, except only that it must be dynamically re-solved whenever shocks occur. The specification of the associated optimisation problem requires that the cost functions  $C_t$ , which give the costs of buying or selling at each successive time  $t$ , are supplemented by further functions  $A_t$  associated with the expected costs of possible shocks. The cost functions  $C_t$  are formally assumed to be deterministic. However, when prices are stochastic, deterministic approximations may be used and updated at successive time steps; this deterministic re-optimisation approach is common when storage is used for price arbitrage alone—see, for example, [30, 15, 31, 32] and, for the case where storage is sufficiently large as to have market impact, see [16]. The functions  $A_t$  (which, although defined in terms of the stochastic process of shocks, are also deterministic) are formally introduced in Section 2. We show that in the case of a perfectly efficient store facing linear buying and selling costs the functions  $A_t$  may be determined exactly, and that otherwise they may typically be estimated to a good approximation.

The optimal control up to the time of the first shock is given by the solution, at the start of the control period, of an optimisation problem which can be regarded as that of minimising the costs associated with the store buying and selling added to those of notionally “insuring” for each future instant in time against the effects of the random fluctuations, i.e. the shocks, resulting from the provision of buffering services. The cost of such “insurance” depends on the absolute level of the store at the relevant time. Thus the deterministic problem is that of choosing the vector of successive levels of the store so as to minimise a total cost function  $\sum_t [C_t(x_t) + A_t(s_t)]$ , subject to rate and capacity constraints. Here  $C_t(x_t)$  is the cost of incrementing the level of the store (positively or negatively) at time  $t$  by  $x_t$ , and the function  $A_t$  is such that  $A_t(s_t)$  is the expected additional cost of dealing with any shock which may occur at the time  $t$  when the level of the store is then  $s_t$ . We define this optimisation problem more carefully in Sections 2 and discuss various possible approaches to its solution. In the stochastic environment in which the store operates, the solution of this problem determines the future control of the store until such time as its buffering services are actually required, following which the level of the store is perturbed and the optimisation problem must be re-solved starting at the new level. The continuation of this process provides what is in principle the exactly optimal stochastic control of the store on a potentially indefinite time scale.

In Section 2 we formulate the relevant stochastic model, discuss its applicability, and give various approaches to the determination of the optimal control. These approaches require the availability of good estimates of the above functions  $A_t$ , and in Section 4 we show how these may be obtained. In Section 3 we provide some characteristic properties of optimal solutions, which we relate to empirical work in the existing literature. In Section 5 gives examples.

## 2 Model and determination of optimal control

Consider the management of a store over a finite time interval which is divided into a succession of periods indexed by  $1, \dots, T$ . At the start of each time period  $t$  the store

makes a decision as to how much to buy or sell during that time period; however, the level of the store at the end of that time period may be different from that planned if, during the course of the period, the store is called upon to provide buffering services to deal with some unexpected event or random *shock*. Such a shock might be the need to supply additional energy during the time period  $t$  due to an unexpected failure—for example that of a generator—or might simply be the difference between forecast and actual renewable generation or demand.

We suppose the store has a capacity of  $E$  units of energy. Similarly we suppose that the total energy which may be input or output during any time period is subject to rate (i.e. power) constraints  $P_I$  and  $P_O$  respectively. This slotted-time model corresponds, for example, to real world energy markets where energy is typically traded at half-hourly or hourly intervals, with the actual delivery of that energy occurring in the intervening continuous time period. Detailed descriptions of the operation of the UK market can be found in [33, 34]. The theory developed here easily extends to the case where the above storage parameters are time dependent.

Define also the set  $X = [-P_O, P_I]$ . Both buying and selling prices associated with any time period  $t$  may be represented by a convex function  $C_t$  defined on  $X$  such that  $C_t(x)$  is the cost of a planned change of  $x$  to the level of the store during the time period  $t$ . Typically each function  $C_t$  is increasing and  $C_t(0) = 0$ ; then, for positive  $x$ ,  $C_t(x)$  is the cost of buying  $x$  units and, for negative  $x$ ,  $C_t(x)$  is the negative of the reward for selling  $-x$  units. Then the convexity assumption corresponds, for each time  $t$ , to an increasing cost of buying each additional unit, a decreasing reward obtained for selling each additional unit, and every unit buying price being at least as great as every unit selling price. When, as is usually the case, the store is not perfectly *efficient* in the sense that only a fraction  $\eta \leq 1$  of the energy input is available for output, then this may be captured in the cost function by reducing selling prices by the factor  $\eta$ ; under the assumption that the cost functions  $C_t$  are increasing it is easily verified that this adjustment preserves the above convexity of the functions  $C_t$ . We thus assume that the cost functions are so adjusted so as to capture any such round-trip inefficiency. The functions  $C_t$  are taken to be deterministic but, as discussed in the Introduction, in a stochastic environment a deterministic re-optimisation approach is possible.

A further form of possible inefficiency of a store is *leakage*, whereby a fraction of the contents of the store is lost in each unit of time. We do not explicitly model this here. However, only routine modifications are required to do so, and are entirely analogous to those described in [16].

Suppose that at the end of each time period  $t - 1$  the level of the store is given by the random variable  $S_{t-1}$ , where we take  $S_0$  to be given by the initial level  $s_0$  of the store. We assume that one may then choose a *planned* adjustment (contract to buy or sell)  $x_t \in X$  and such that  $S_{t-1} + x_t \in [0, E]$  to the level of the store during the time period  $t$ . The planned adjustment  $x_t$  is a (deterministic) function of the level  $S_{t-1}$  and the cost of this adjustment is  $C_t(x_t)$ . Subsequent to this, during the course of the time period  $t$ , the store may be subject to a *shock* or random disturbance, corresponding to the need to provide unexpected buffering services. This shock has an associated cost, typically due to the store not being able to provide the required services, and may further disturb the final level of the store at the end of the time period  $t$ . We assume that the cost of any shock occurring during the time period  $t$  and the resulting *actual* level of the store at the end of the time period  $t$  are given by random variables whose joint distribution is a function of the

*planned* final level  $S_{t-1} + x_t$  of the store for the end of that time period, but that, given this planned final level, these random variables are otherwise independent of all else. Thus we may assume that there are given  $T$  independent stochastic processes  $(D_t(s), S_t(s))_{s \in [0, E]}$  each taking values in  $\mathbb{R} \times \mathbb{R}$  (where each such “process” is indexed by the possible levels  $s$  of the store rather than by time), and that the shock cost and actual store level at the end of each time period  $t$  are then given respectively by  $D_t(S_{t-1} + x_t)$  and  $S_t(S_{t-1} + x_t)$ ; in the absence of any shock during the time period  $t$ , we have  $D_t(s) = 0$  and  $S_t(s) = s$  for all  $s$ . The assumption that the joint distribution of the shock cost and store level disturbance associated with any time period  $t$  depend only on the planned level  $S_{t-1} + x_t$  of the store at the end of the time period  $t$  is likely to be most accurate in applications where the store is able to adjust to its target level quickly within each time period, or where the level of the store does not change too much within a single time period; its relaxation—for example, by allowing a dependence of these random variables on a more general function of  $S_{t-1}$  and  $x_t$ —simply complicates without essentially changing the analysis below. Note that the model further assumes that disturbances do not persist beyond the end of the time periods in which they occur. Under any given control policy for the management of the store satisfying the above conditions (i.e. under any specification, for each time  $t$ , of the planned increment  $x_t$  as a function of the realised value of  $S_{t-1}$ ), the levels  $S_t$  of the store at the end of the successive time periods  $t$  form a Markov process.

For each  $t$ , and conditional on each possible value  $s_{t-1}$  of the level  $S_{t-1}$  of the store at the end of the time period  $t-1$ , define  $V_{t-1}(s_{t-1})$  to be the expected future cost of subsequently managing the store under an optimal control—where, here and elsewhere, by an optimal control we mean a control defined as above under which the expected cost of managing the store is minimised. We then have the SDP recursion

$$V_{t-1}(s_{t-1}) = \min_{\substack{x_t \in X \\ s_{t-1} + x_t \in [0, E]}} [C_t(x_t) + \mathbf{E}[D_t(s_{t-1} + x_t) + V_t(S_t(s_{t-1} + x_t))]], \quad (1)$$

where  $\mathbf{E}$  denotes expectation and where, as above,  $s_{t-1} + x_t$  and  $S_t(s_{t-1} + x_t)$  ( $= S_t$ ) are respectively the planned and actual levels of the store at the end of the time period  $t$ . (The assumed independence of the “processes” of paired random variables  $\{(D_t(s), S_t(s))\}_{s \in [0, E]}$  defining shock costs and disturbances in successive time periods ensures that it is sufficient to consider unconditional expectations in (1).) We further have the terminal condition

$$V_T(s_T) = 0 \quad (2)$$

for all possible levels  $s_T$  of the store at the end of the time period  $T$ . The recursion (1) and the terminal condition (2) may in principle be used to determine an optimal control. In particular, given the level  $s_{t-1}$  of the store at the end of any time period  $t-1$ , the optimal planned increment to the level of the store for the time period  $t$  is given by  $\hat{x}_t(s_{t-1})$  where this is defined to be the value of  $x_t$  which achieves the minimisation in the recursion (1).

However, as discussed in the Introduction, an SDP approach may frequently be computationally intractable and is further not suitable for the management of a store over indefinite time horizons. Thus, for each  $t$ , let the (deterministic) function  $A_t$  on  $[0, E]$  be such that, for any *planned* level  $s_t = s_{t-1} + x_t$  of the store for the end of the time period  $t$ ,

$$A_t(s_t) = \mathbf{E}[D_t(s_t) + V_t(S_t(s_t))] - V_t(s_t), \quad (3)$$

where again the random variable  $S_t(s_t)$  is the *actual* level of the store at the end of the time period  $t$ . Given the planned level  $s_t$  of the store for the end of the time period  $t$ ,

the quantity  $A_t(s_t)$  is the difference between the expected cost  $\mathbf{E}(D_t(s_t) + V_t(S_t(s_t)))$  of optimally managing the store during and subsequent to the time period  $t$  and the corresponding expected cost  $V_t(s_t)$  which would be incurred in the guaranteed absence of any shock during that time period. We shall show in Section 4 that in many cases the functions  $A_t$  may be efficiently determined either exactly or to a very good approximation even in the absence of any knowledge of the functions  $V_t$ .

It now follows from (3) that the recursion (1) may be rewritten as

$$V_{t-1}(s_{t-1}) = \min_{\substack{x_t \in X \\ s_{t-1} + x_t \in [0, E]}} [C_t(x_t) + A_t(s_{t-1} + x_t) + V_t(s_{t-1} + x_t)], \quad (4)$$

where we again require the terminal condition (2). Further, given the level  $s_{t-1}$  of the store at the end of any time period  $t-1$ , the optimal planned increment  $\hat{x}_t(s_{t-1})$  to the level of the store for the time period  $t$  is given by the value of  $x_t$  which achieves the minimisation in the recursion (4).

The (backwards) recursion (4) is entirely deterministic. *Given a knowledge of the functions  $A_t$*  (see Section 4), a complete solution of the recursion (4) would determine, for all  $t = 1, \dots, T$  and for all possible levels  $s_{t-1}$  of the store at the end of the time period  $t-1$ , both the minimised expected future cost  $V_{t-1}(s_{t-1})$  and the optimal planned increment  $\hat{x}_t(s_{t-1})$  to the level of the store for the time period  $t$ . Now let the (Markov) process  $(\hat{S}_0, \dots, \hat{S}_T)$ , with  $\hat{S}_0 = s_0$ , correspond to the sequence of levels of the optimally controlled store. Then, since this process is random, the optimal planned increment  $\hat{x}_t(\hat{S}_{t-1})$  for each time period  $t$  is not known until the end of the time period  $t-1$ .

However, the solution of the recursion (4) as above, typically require the determination of each of the functions  $V_t$  for all possible values of its argument. We therefore define a deterministic optimisation problem whose solution  $s^* = (s_0^*, \dots, s_T^*)$ , with  $s_0^* = s_0$ , coincides with the optimal control of the store up to the time of the first shock. As we discuss below, the solution of this optimisation problem is typically computationally much simpler than the complete solution of recursion (4). However, it is necessary to re-solve this optimisation problem at the end of each time period in which a shock occurs.

For any vector  $s = (s_0, \dots, s_T)$  of possible store levels, where  $s_0$  is constrained to be the initial level of the store, and for each  $t = 1, \dots, T$ , define

$$x_t(s) = s_t - s_{t-1}. \quad (5)$$

Define also the optimisation problem:

$\mathbb{P}$ : choose  $s = (s_0, \dots, s_T)$ , where again  $s_0$  is the initial level of the store, so as to minimise

$$\sum_{t=1}^T [C_t(x_t(s)) + A_t(s_t)] \quad (6)$$

subject to the capacity constraints

$$0 \leq s_t \leq E, \quad 1 \leq t \leq T, \quad (7)$$

and the rate constraints

$$x_t(s) \in X, \quad 1 \leq t \leq T. \quad (8)$$

Let  $s^* = (s_0^*, \dots, s_T^*)$ , with  $s_0^* = s_0$ , denote the solution to the above problem  $\mathbb{P}$ . The recursion (4) is the dynamic programming recursion for the solution of the problem  $\mathbb{P}$  and

it follows straightforwardly from iteration of (4), using also the terminal condition (2), that  $x_1(s^*)$  achieves the minimisation in (4) for  $t = 1$ , i.e. that  $x_1(s^*) = \hat{x}_1(s_0)$  is planned first increment in the optimal control of the store. Thus, from (5), provided no shock occurs during the time period 1 so that  $\hat{S}_1 = s_0 + \hat{x}_1(s_0)$ , we have also that  $\hat{S}_1 = s_1^*$ . More generally, let the random variable  $T'$  index the first time period during which a shock does occur. Then repeated application of the above argument gives immediately the following result.

**Theorem 1.** *For all  $t < T'$ , we have  $\hat{S}_t = s_t^*$ .*

The solution to the problem  $\mathbb{P}$  therefore defines the optimal control of the store up to the end of the time period  $T'$  defined above. At that time it is necessary to reformulate the problem  $\mathbb{P}$ , starting at the end of the time period  $T'$ , instead of at time 0, and replacing the initial level of the store  $s_0$  by the perturbed level  $\hat{S}_{T'}$  at that time. Iterative application of this process at the times of successive shocks leads to the dynamically determined stochastic optimal control—which is exact to the extent that the functions  $A_t$  are known exactly.

Given that the functions  $A_t$  are known, either exactly or to a sufficiently good approximation (again see Section 4), the deterministic optimisation problem  $\mathbb{P}$  may be solved by using strong Lagrangian techniques to derive a *forward* algorithm which is computationally much simpler than the use of a dynamic programming approach, and which further identifies a running *planning* or *forecast horizon*. The latter is such that, for each time  $t$  there exists a time  $t' > t$  such that the optimal decision at time  $t$  does not depend on the functions  $C_u$  and  $A_u$  for  $u > t'$ . This is proved in [16] for the case in which the functions  $A_t$  are zero, but the more general result and algorithm may be derived along the same lines. The existence of such a running forecast horizon further reduces the computation required in the solution of the problem  $\mathbb{P}$  and makes the present approach particularly suitable for the management of storage over a very long or indefinite time period. It further means that, in an environment in which prices—and so the cost functions  $C_t$ —are uncertain, in order to make the optimal decision at any time  $t$  as above it is only necessary to estimate the cost functions  $C_u$  for values of  $u$  up to the associated forecast horizon  $t'$ . In the case where, as in the fairly realistic examples of Section 5, the store fills and partially empties on an approximate daily cycle, the length of this forecast horizon is typically of the order of a day or two. In practice electricity prices in particular may often be estimated accurately on such time scales, and a deterministic re-optimisation approach, as discussed in the Introduction, is likely to suffice for the optimal control of the store.

### 3 Characterisation of optimal solutions

In this section we establish some properties of the functions  $\hat{x}_t(\cdot)$  defined in the previous section (as achieving the minimisation in the recursion (4)) and determining the optimal control of the store.

One case of particular interest is that where the store is a price-taker (i.e. the store is not so large as to impact itself on market prices), so that, for each  $t$ , the cost function  $C_t$  is given by

$$C_t(x) = \begin{cases} c_t^{(b)} x, & \text{if } x \geq 0 \\ c_t^{(s)} x, & \text{if } x < 0, \end{cases} \quad (9)$$



where the unit “buying” price  $c_t^{(b)}$  and the unit “selling” price  $c_t^{(s)}$  are such that  $c_t^{(s)} \leq c_t^{(b)}$  (possible inequality resulting, for example, from the round-trip inefficiency of the store—see the discussion of Section 2.)

Theorem 2 below is a simple result which shows that in the case where buying and selling prices are equal, and provided rate constraints are nonbinding, the optimal policy is a “target” policy. That is, for each time period  $t$  there exists a target level  $\hat{s}_t$  such that, given that the level of the store at the end of the immediately preceding time period is  $s_{t-1}$ , the optimal planned level  $s_{t-1} + x_t$  of the store to be achieved during the time period  $t$  is set equal to  $\hat{s}_t$ , independently of  $s_{t-1}$ .

**Theorem 2.** *Suppose that, for each  $t$ , we have  $c_t^{(b)} = c_t^{(s)} = c_t$ ; define*

$$\hat{s}_t = \arg \min_{s \in [0, E]} [c_t s + A_t(s) + V_t(s)], \quad (10)$$

where the functions  $A_t$  and  $V_t$  are as introduced in Section 2. Then, for each  $t$  and for each  $s_{t-1}$ , we have  $\hat{x}_t(s_{t-1}) = \hat{s}_t - s_{t-1}$  provided only that this quantity belongs to the set  $X$ .

*Proof.* The recursion (4) here becomes, for each  $t$ ,

$$V_{t-1}(s_{t-1}) = \min_{\substack{x_t \in X \\ s_{t-1} + x_t \in [0, E]}} [c_t x_t + A_t(s_{t-1} + x_t) + V_t(s_{t-1} + x_t)], \quad (11)$$

and the above minimisation is achieved by  $x_t$  such that  $s_{t-1} + x_t = \hat{s}_t$ , provided only that  $x_t \in X$ .  $\square$

In order to deal with the possibility of rate constraint violation, or the more general price-taker case where  $c_t^{(s)} < c_t^{(b)}$ , or the general case where the cost functions  $C_t$  are merely required to be convex, we require the additional assumption of convexity of the functions  $A_t$ . This condition, while not automatic, is reasonably natural in many applications—see the examples of Section 5.

**Theorem 3.** *Suppose that, in addition to convexity of the functions  $C_t$ , each of the functions  $A_t$  is convex. Then, for each  $t$ :*

- (i) *the function  $V_{t-1}$  is convex;*
- (ii)  *$\hat{x}_t(s_{t-1})$  is a decreasing function of  $s_{t-1}$ ;*
- (iii)  *$s_{t-1} + \hat{x}_t(s_{t-1})$  is an increasing function of  $s_{t-1}$ .*

*Proof.* To show (i) we use backwards induction in time. The function  $V_T$  is convex. Suppose that, for any given  $t \leq T$ , the function  $V_t$  is convex; we show that the function  $V_{t-1}$  is convex. For any given values  $s_{t-1}^{(i)}$ ,  $i = 1, \dots, n$ , of  $s_{t-1}$  and for any convex combination  $\bar{s}_{t-1} = \sum_{i=1}^n \kappa_i s_{t-1}^{(i)}$ , where each  $\kappa_i \geq 0$  and where  $\sum_{i=1}^n \kappa_i = 1$ , define also  $\bar{x}_t = \sum_{i=1}^n \kappa_i \hat{x}_t(s_{t-1}^{(i)})$ . Note that  $\bar{x}_t \in X$  and that  $\bar{s}_{t-1} + \bar{x}_t \in [0, E]$ . Then, from (4),

$$\begin{aligned} V_{t-1}(\bar{s}_{t-1}) &\leq C_t(\bar{x}_t) + A_t(\bar{s}_{t-1} + \bar{x}_t) + V_t(\bar{s}_{t-1} + \bar{x}_t) \\ &\leq \sum_{i=1}^n \kappa_i \left( C_t(\hat{x}_t(s_{t-1}^{(i)})) + A_t(s_{t-1}^{(i)} + \hat{x}_t(s_{t-1}^{(i)})) + V_t(s_{t-1}^{(i)} + \hat{x}_t(s_{t-1}^{(i)})) \right) \\ &= \sum_{i=1}^n \kappa_i V_{t-1}(s_{t-1}^{(i)}), \end{aligned}$$

where the second inequality above follows from the convexity of the functions  $C_t$ ,  $A_t$  and  $V_t$  (the latter by the inductive hypothesis). Thus  $V_{t-1}$  is convex as required.

To show (ii) and (iii) we make use of the following result: let  $f$  and  $g$  be functions defined on the real line  $\mathbb{R}$  such that  $g$  is convex, and suppose that, for each fixed  $s$ , the function of  $x \in \mathbb{R}$  given by  $f(x) + g(s+x)$  is minimised by  $\hat{x}(s)$ ; then  $\hat{x}(s)$  is a decreasing function of  $s$ . To see this, suppose that  $s_1 < s_2$  and note that, under the given assumptions,

$$f(\hat{x}(s_1)) + g(s_1 + \hat{x}(s_1)) \leq f(x) + g(s_1 + x), \quad x \in \mathbb{R}. \quad (12)$$

The convexity of  $g$  implies straightforwardly that

$$g(s_2 + \hat{x}(s_1)) - g(s_2 + x) \leq g(s_1 + \hat{x}(s_1)) - g(s_1 + x), \quad \text{for all } x > \hat{x}(s_1). \quad (13)$$

It follows from (12) and (13) that

$$f(\hat{x}(s_1)) + g(s_2 + \hat{x}(s_1)) \leq f(x) + g(s_2 + x), \quad \text{for all } x > \hat{x}(s_1).$$

It now follows from the above that  $\hat{x}(s_2)$  is (or, in the absence of uniqueness, may be taken to be) less than or equal to  $\hat{x}(s_1)$ .

The result (ii) of the theorem now follows by applying the above result with the function  $f$  given by  $C_t$  and the function  $g$  given by  $A_t + V_t$ , since  $A_t$  is assumed convex and, from (i),  $V_t$  is also convex. (That the minimisation in (4) is taken over those  $x$  within a closed interval of the real line causes no problems: for example, this restriction may be formally dropped by extending the domains of definition of  $C_t$ ,  $A_t$  and  $V_t$  to the entire real line, taking them to be infinite outside the intervals on which they are naturally defined.)

The result (iii) of the theorem similarly follows by applying the above general result with the function  $f$  given by  $A_t + V_t$  and the function  $g$  given by the convex function  $C_t$  (in the recursion (4) writing  $C_t(x_t) = C_t(-s_{t-1} + (s_{t-1} + x_t))$  and, for each fixed value of  $s_{t-1}$ , regarding the minimisation in (4) as being over the variable  $s_{t-1} + x_t$ ).  $\square$

*Remark 1.* Given initial levels  $s_0^{(1)}$  and  $s_0^{(2)}$  of the store, let  $\{S_t^{(1)}\}$  and  $\{S_t^{(2)}\}$  (with  $S_0^{(1)} = s_0^{(1)}$  and  $S_0^{(2)} = s_0^{(2)}$ ) be the respective optimally controlled stochastic processes of levels of the store—coupled with respect to the underlying stochastic process of shocks. Suppose we additionally assume that the level of the store immediately following any shock is an increasing function of the level immediately prior to that shock. It then follows from (iii) of Theorem 3, that under the conditions of the theorem, if  $s_0^{(1)} \leq s_0^{(2)}$  then  $S_t^{(1)} \leq S_t^{(2)}$  for all subsequent  $t$ . This monotonicity property proves useful in Section 4.

We now return to the price-taker case, in which the cost functions are as defined by (9), and which corresponds to a store which is not sufficiently large as to have market impact. Here we prove a strengthened version of Theorem 3. For each  $t$ , given that the function  $A_t$  is convex, define

$$s_t^{(b)} = \arg \min_{s \in [0, E]} [c_t^{(b)} s + A_t(s) + V_t(s)] \quad (14)$$

and similarly define

$$s_t^{(s)} = \arg \min_{s \in [0, E]} [c_t^{(s)} s + A_t(s) + V_t(s)]. \quad (15)$$

Note that the above convexity assumption and the condition that, for each  $t$ , we have  $c_t^{(s)} \leq c_t^{(b)}$  imply that  $s_t^{(b)} \leq s_t^{(s)}$ . We now have the following result.

**Theorem 4.** *Suppose that the cost functions  $C_t$  are as given by (9) and that the functions  $A_t$  are convex. Then the optimal policy is given by: for each  $t$  and given  $s_{t-1}$ ,*

$$\hat{x}_t(s_{t-1}) = \begin{cases} \min(s_t^{(b)} - s_{t-1}, P_I) & \text{if } s_{t-1} < s_t^{(b)}, \\ 0 & \text{if } s_t^{(b)} \leq s_{t-1} \leq s_t^{(s)}, \\ \max(s_t^{(s)} - s_{t-1}, -P_O) & \text{if } s_{t-1} > s_t^{(s)}. \end{cases} \quad (16)$$

*Proof.* For each  $t$ , it follows from the convexity of the functions  $C_t$ ,  $A_t$  and  $V_t$  (the latter by the first part of Theorem 3) that, for  $s_{t-1} < s_t^{(b)}$  the function  $C_t(x_t) + A_t(s_{t-1} + x_t) + V_t(s_{t-1} + x_t)$  is minimised by  $x_t = s_t^{(b)} - s_{t-1}$ , for  $s_t^{(b)} \leq s_{t-1} \leq s_t^{(s)}$  it is minimised by  $x_t = 0$ , while for  $s_{t-1} > s_t^{(s)}$ , it is minimised by  $x_t = s_t^{(s)} - s_{t-1}$ . The required result now follows from the recursion (4) (on again using the convexity of the functions  $C_t$ ,  $A_t$  and  $V_t$  to account for the rate constraint in that recursion).  $\square$

Thus in general in the price-taker case there exists, for each time period  $t$ , a “target interval”  $[s_t^{(b)}, s_t^{(s)}]$  such that, if the level of the store at the end of the previous time period is  $s_{t-1}$ , the optimal policy is to chose  $\hat{x}_t(s_{t-1})$  so that  $s_{t-1} + \hat{x}_t(s_{t-1})$  is the nearest point (in absolute distance) to  $s_{t-1}$  lying within, or as close as possible to, the above interval. In the case where  $c_t^{(b)} = c_t^{(s)} = c_t$ , the above interval shrinks to the single point  $\hat{s}_t$  defined by (10).

These results shed some light on earlier, more applied, papers of Bejan et al [23] and Gast et al [24], in which the uncertainties in the operation of a energy store result from errors in wind power forecasts. The model considered in those papers is close to that of the present paper, as we now describe. The costs of operating the store result (a) from round-trip inefficiency, which in the formulation of the present paper would be captured by the cost functions  $C_t$  as defined by (9) with  $c_t^{(s)} < c_t^{(b)}$  and with  $C_t$  the same for all  $t$ , and (b) from buffering events, i.e. from failures to meet demand through insufficient energy available to be supplied from the store when it is needed, and from energy losses through store overflows. In the formulation of the present paper these costs would be captured by the functions  $A_t$ . In contrast to the present paper decisions affecting the level of the store (the amount of conventional generation to schedule for a particular time) are made  $n$  time steps—rather than a single time step—in advance, when wind power is forecast and conventional generation scheduled. The underlying arguments leading to Theorems 2–4 continue to apply, at least to a good approximation. In particular sample path arguments suggest that the reduction of round-trip efficiency slows the rate at which the store-level trajectories—started from different initial levels but with the same stochastic description of future shock processes—converge over subsequent time. In particular Gast et al [24] confirm these results empirically, considering round-trip efficiencies less than 1 and noting that in this case simple “target” policies such as that described by Theorem 2 (which is applicable in the case of round-trip efficiencies equal to 1) are here suboptimal.

## 4 Determination of the functions $A_t$

We described in Section 2 how, given a knowledge of the functions  $A_t$  defined by (3), the optimal control of the store may be reduced to the solution of an optimisation problem which must be re-solved at those randomly occurring times at which shocks occur. In this section we consider conditions under which the functions  $A_t$  may be thus known,

either exactly or to good approximations—in all cases without the need for the prior determination of the functions  $V_t$ .

It is convenient to rewrite slightly the definition (3) of each of the functions  $A_t$  as

$$A_t(s_t) = \mathbf{E}D_t(s_t) + \mathbf{E}[V_t(S_t(s_t)) - V_t(s_t)], \quad (17)$$

and to regard  $A_t(s_t)$  as the sum of the two given expectations on the right side of (17). We shall argue below that in many applications it is the first of these two expectations, i.e.  $\mathbf{E}D_t(s_t)$ , that is likely to be much the dominant term on the right side of (17). Since, for each  $t$  and for each  $s_t$ , the distribution of  $D_t(s_t)$  is part of the model specification, the computation of  $\mathbf{E}D_t(s_t)$  is straightforward. We do, however, consider below how one might reasonably *obtain* this major part of the model specification, i.e. the cost of dealing with a random shock as a function of the level of the store at the time at which the shock occurs.

The second expectation on the right side of (17) is the difference between the expected cost  $\mathbf{E}V_t(S_t(s_t))$  of optimally managing the store subsequent to the time period  $t$  (when the actual level of the store at the end of that time period is then given by the random variable  $S_t(s_t)$ ) and the corresponding expected cost  $V_t(s_t)$  which would be incurred in the absence of any shock during the time period  $t$  (so that the level of the store at the end of that time period was then its planned value  $s_t$ ). This difference  $\mathbf{E}V_t(S_t(s_t)) - V_t(s_t)$  may also be understood in terms of a coupling of optimally controlled processes, started at the end of the time period  $t$  at the levels  $S_t(s_t)$  and  $s_t$ , and is the expectation of the difference of the costs of their optimal control up to the time at which the coupled processes first agree.

Now consider the somewhat idealised conditions of Theorem 2, where the store is a perfectly efficient price-taker, so that each cost function  $C_t$  is given by  $C_t(x) = c_t x_t$  for some market price  $c_t$ , and where each target level  $\hat{s}_t$  given by (10) is assumed to be always achievable. It follows from Theorem 2 that, regardless of any shock which may occur during any given time period  $t$ , the planned level of the store for the end of the time period  $t + 1$  is  $\hat{s}_{t+1}$ . Hence, from (4),

$$V_t(S_t(s_t)) - V_t(s_t) = C_{t+1}(s_t - S_t(s_t)), \quad (18)$$

so that, from (17),

$$A_t(s_t) = \begin{cases} \mathbf{E}D_t(s_t) + C_{t+1}(s_t - \mathbf{E}S_t(s_t)), & t < T, \\ \mathbf{E}D_T(s_T), & t = T. \end{cases} \quad (19)$$

Thus the functions  $A_t$  may here be determined—in terms of the given distributions of the random variables  $D_t(s)$  and  $S_t(s)$ —without the need to estimate the functions  $V_t$ .

More generally, the relations (19) correspond to the modified control in which, following any shock and hence store level disturbance during any time period  $t$ , the disturbed level  $S_t(s_t)$  of the store is immediately returned to the planned level  $s_t$  for the end of that period at a cost  $C_{t+1}(s_t - S_t(s_t))$ ; subject to this the store is otherwise optimally managed.

In the absence of this modification, the relations (19) may be viewed as providing a reasonable first approximation to the functions  $A_t$ —given the difficulties, in applications, of estimating the both the likelihood and the precise consequences of shocks, it is not clear that one could do significantly better. Better approximations, if required, might be made by allowing more time for the disturbed and undisturbed processes to couple as described

above, and by reasoning as before so as to obtain a more refined version of (19). For example, one might extend the coupling time until a known future time at which it is planned that the store will be full. (This is often realistic for electricity storage which may aim to be full at the end of each night so as to take advantage of much higher daytime prices—we give examples based on real price data and realistic store characteristics in Section 5.) Then, for any planned level  $s_t$  of the store at the end of any time period  $t$ , the quantity  $V_t(S_t(s_t)) - V_t(s_t)$  may be estimated analogously to (18) by considering optimal controls from the end of the time period  $t$  up to the first subsequent time at which the store is planned to be full.

Finally in the important special case in which shocks are rare but potentially expensive (as might be the case when the store is required to pay the costs of failing to have sufficient energy to deal with an emergency), then, for each  $t$  and  $s_t$ , the probability that  $S_t(s_t)$  is not equal to  $s_t$  is small, and the major contribution to  $A_t(s_t)$  as defined by (17) is likely to be  $\mathbf{E}D_t(s_t)$ . In this case either the simple approximation  $A_t(s_t) = \mathbf{E}D_t(s_t)$ , or the more refined approximation given by (19), may well suffice in applications.

In applications there is also a need, as part of the model specification, to realistically estimate—for each possible planned level  $s_t$  of the store at the end of each time period  $t$ —the joint distribution of the random vector  $(D_t(s_t), S_t(s_t))$  modelling the cost of any shock and the corresponding store level disturbance. This joint distribution is in general a function of the amount of energy  $Y_t$  required to deal with any shock during the time period  $t$ , where in practice the distribution of the random variable  $Y_t$  may need to be determined by observation. We consider two particular possibilities, both of which are natural in the context of modelling risk in power systems, where the focus may either be on *loss of load* or on *energy unserved* (see, for example, [35]):

- (i) the cost of a shock occurring during the time period  $t$  is simply a constant  $a_t > 0$  if there is insufficient energy within the store to meet it, and is 0 otherwise; we then have  $D_t(s_t) = a_t \mathbf{I}(Y_t > s_t)$ , where  $\mathbf{I}(\cdot)$  is the indicator function, and  $S_t(s_t) = \max(0, s_t - Y_t)$ ;
- (ii) the cost of a shock occurring during the time period  $t$  is proportional to the shortfall in the energy necessary to meet that shock; we then have  $D_t(s_t) = a'_t \max(0, Y_t - s_t)$ , where  $a'_t$  is the constant of proportionality, and again  $S_t(s_t) = \max(0, s_t - Y_t)$ .

Given the model of Section 2, the functions  $A_t$  may be determined (as described in this section) from the specification of the joint distributions of the random vectors  $(D_t(s), S_t(s))$ , together with the specification of the cost functions  $C_t$  and the store characteristics. In Section 5 we consider some plausible functional forms of the functions  $A_t$ .

## 5 Examples

We give some examples, in which we solve (exactly) the optimal control problem  $\mathbb{P}$  formally defined in Section 2. We investigate how the optimal solution depends on the cost functions  $C_t$  and on the functions  $A_t$  which reflect the costs of providing buffering services. The cost functions  $C_t$  are derived from half-hourly electricity prices in the Great Britain spot market over the entire year 2011, adjusted for a modest degree of market impact, as described in detail below. Thus we work in half-hour time units, with the time horizon  $T$  corresponding to the number of half-hour periods in the entire year. These spot market prices show a strong daily cyclical behaviour (corresponding to daily demand variation), being low at night and high during the day. This price variation can be seen in Figure 1




Figure 1: GB half-hourly spots prices (£/MWh) for March 2011.

which shows half-hourly GB spot prices (in pounds per megawatt-hour) throughout the month of March 2011. There is a similar pattern of variation throughout the rest of the year.

Without loss of generality, we choose energy units such that the rate (power) constraints are given by  $P_I = P_O = 1$  unit of energy per half-hour period. For illustration, we take the capacity of the store to be given by  $E = 10$  units of energy; thus the store can completely fill or empty over a 5-hour period, which is the case, for example, for the large Dinorwig pumped storage facility in Snowdonia [36].

We choose cost functions  $C_t$  of the form

$$C_t(x) = \begin{cases} c_t x(1 + \delta x), & \text{if } x \geq 0 \\ \eta c_t x(1 + \delta x), & \text{if } x < 0, \end{cases} \quad (20)$$

where the  $c_t$  are proportional to the half-hourly electricity spot prices referred to above, where  $\eta$  is an adjustment to selling prices representing in particular round-trip efficiency as described in Section 2, and where the factor  $\delta > 0$  is chosen so as to represent a degree of market impact (higher unit prices as the store buys more and lower unit prices as the store sells more). For our numerical examples we take  $\eta = 0.85$  which is a typical round-trip efficiency for a pumped-storage facility such as Dinorwig. We choose  $\delta = 0.05$ ; since the rate constraints for the store are  $P_I = P_O = 1$  this corresponds to a maximum market impact of 5%. While this is modest, our results are qualitatively little affected as  $\delta$  is varied over a wide range of values less than one, covering therefore the range of possible market impact likely to be seen for storage in practice.

Finally we need to choose the functions  $A_t$  reflecting the costs of providing buffering services. Our aim here is to give an understanding of how the optimal control of the

store varies according to the relative economic importance of cost arbitrage and buffering, i.e. according to the relative size of the functions  $C_t$  and  $A_t$ . We choose functions  $A_t$  which are constant over time  $t$  and of the form  $A_t(s) = ae^{-\kappa s}$  and  $A_t(s) = b/s$  for a small selection of the parameters  $a$ ,  $\kappa$  and  $b$ . The extent to which a store might provide buffering services in applications is extremely varied, and so the likely balance between arbitrage and buffering cannot be specified in advance. Rather we choose just sufficient values of the above parameters to show the effect of varying this balance. For a possible justification of the chosen forms of the functions  $A_t$ , see Section 4; in particular the form  $A_t(s) = ae^{-\kappa s}$  is plausible in the case of light-tailed shocks, while the form  $A_t(s) = b/s$  shows the effect of a slow rate of decay in  $s$ . (Note that in these examples we allow that the functions  $A_t$  should not necessarily be constant for values of their arguments greater than the rate constraint of 1: it is plausible that in practice greater quantities in store than can immediately be discharged to deal with a shock may nevertheless assist in dealing with its ongoing effects at subsequent times and, in the event of such a shock, may be considered as being notionally set aside for this purpose.)

In each of our examples, we determine the optimal control of the store over the entire year, with both the initial level  $S_0^*$  and the final level  $S_T^*$  given by  $S_0^* = S_T^* = 0$ . Figure 2 shows this optimal control (the sequence of successive levels of the store) for the time window corresponding to the month of March for each of the four cases  $A_t(s) = 0$ ,  $A_t(s) = e^{-s}$ ,  $A_t(s) = 10e^{-s}$ , and  $A_t(s) = 1/s$ . In each case the corresponding running forward horizon, as defined in Section 2, is generally of the order of a day or two. (Recall that the cost functions for March are determined by the prices illustrated in Figure 1. Although the optimal control is determined over the entire year, it may be verified empirically that in every case the restriction of this optimal control to any given time window is independent of the functions  $C_t$  and  $A_t$  for times  $t$  which are outside of a period which includes this time window and a few days on either side of it.)

The case  $A_t(s) = 0$  corresponds to the store incurring no penalty for failing to provide buffering services and optimising its control solely on the basis of arbitrage between energy prices at different times. The daily cycle of prices (again see Figure 1) is sufficiently pronounced that here the store fills and empties—or nearly so—on a daily basis, notwithstanding the facts that the round-trip efficiency of 0.85 is considerably less than 1 and that the minimum time for the store to fill or empty is 5 hours.

In the case  $A_t(s) = e^{-s}$  the store is just sufficiently incentivised by the need to reduce buffering costs that it rarely empties completely (though it does so very occasionally). Otherwise the behaviour of the store is very similar to that in the case  $A_t(s) = 0$ . In both the cases  $A_t(s) = 10e^{-s}$  and  $A_t(s) = 1/s$  the costs of failing to provide buffering services are much higher, and so the optimised level of the store rarely falls below 25% of its capacity. Note the very similar behaviour in these two cases despite the very different forms of the “penalty” functions  $A_t$ .

## Acknowledgements

The authors acknowledge the support of the Engineering and Physical Sciences Research Council for the programme (EPSRC grant EP/I017054/1) under which the present research is carried out. They also acknowledge the benefit of very helpful discussions with numerous colleagues, notably: Janusz Bialek, Chris Dent, Lisa Flatley, Richard Gibbens, Frank Kelly and Phil Taylor. They are further most grateful to the referees for many




Figure 2: Optimally controlled store level throughout March 2011 for each of the four cases  $A_t(s) = 0$ ,  $A_t(s) = e^{-s}$ ,  $A_t(s) = 10e^{-s}$ , and  $A_t(s) = 1/s$ .

insightful comments and suggestions which have considerably improved the exposition of the paper.



## References

- [1] E. D. Castronuovo and J. A. P. Lopes, “Optimal operation and hydro storage sizing of a wind–hydro power plant,” *International Journal of Electrical Power & Energy Systems*, vol. 26, pp. 771–778, Dec. 2004.
- [2] P. Denholm and R. Sioshansi, “The value of compressed air energy storage with wind in transmission-constrained electric power systems,” *Energy Policy*, vol. 37, pp. 3149–3158, Aug. 2009.
- [3] M. Korpaas, A. T. Holen, and R. Hildrum, “Operation and sizing of energy storage for wind power plants in a market system,” *International Journal of Electrical Power & Energy Systems*, vol. 25, pp. 599–606, Oct. 2003.
- [4] P. Harsha and M. Dahleh, “Optimal management and sizing of energy storage under dynamic pricing for the efficient integration of renewable Energy,” *IEEE Transactions on Power Systems*, vol. 30, pp. 1164–1181, May 2015.
- [5] N. Richmond, P. Jacko, and A. M. Makowski, “Optimal planning of slow-ramping power production in energy systems with renewables forecasts and limited storage,” in *2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pp. 1–6, IEEE, July 2014.
- [6] A. Tuohy and M. O’Malley, “Impact of pumped storage on power systems with increasing wind penetration,” in *2009 IEEE Power & Energy Society General Meeting*, pp. 1–8, IEEE, July 2009.
- [7] O. Megel, J. L. Mathieu, and G. Andersson, “Maximizing the potential of energy storage to provide fast frequency control,” in *IEEE PES ISGT Europe 2013*, pp. 1–5, IEEE, Oct. 2013.
- [8] A. Oudalov, D. Chartouni, and C. Ohler, “Optimizing a battery energy storage system for primary frequency control,” *IEEE Transactions on Power Systems*, vol. 22, pp. 1259–1266, Aug. 2007.
- [9] F. Teng, J. Miles, A. Thomson, G. Strbac, N. Brandon, and D. Pudjianto, “Potential value of energy storage in the UK electricity system,” *Proceedings of the ICE - Energy*, vol. 168, pp. 107–117, May 2015.
- [10] A. Bernstein, L. Reyes-Chamorro, J.-Y. Le Boudec, and M. Paolone, “A composable method for real-time control of active distribution networks with explicit power set-points. Part I: Framework,” *Electric Power Systems Research*, vol. 125, pp. 254 – 264, 2015.
- [11] Y. Huang, S. Mao, and R. Nelms, “Adaptive electricity scheduling in microgrids,” in *Proc. IEEE INFOCOM*, Turin, Italy, 2013.
- [12] A. S. Cahn, “The warehouse problem,” *Bulletin of the American Mathematical Society*, vol. 54, no. 11, pp. 1073–1073, 1948.
- [13] R. Bellman, “On the theory of dynamic programming—a warehousing problem,” *Management Science*, vol. 2, no. 3, pp. 272–275, 1956.

- [14] S. E. Dreyfus, “An analytic solution of the warehouse problem,” *Management Science*, vol. 4, no. 1, pp. 99–104, 1957.
- [15] N. Secomandi, “Optimal commodity trading with a capacitated storage asset,” *Management Science*, vol. 56, no. 3, pp. 449–467, 2010.
- [16] J. R. Cruise, L. C. Flatley, R. J. Gibbens, and S. Zachary, “Control of energy storage with market impact: Lagrangian approach and horizons,” *Operations Research*, 2018. To appear.
- [17] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, “Optimal energy storage control policies for the smart power grid,” in *Proc. IEEE SmartGridComm*, pp. 475–480, 2011.
- [18] D. Pudjianto, M. Aunedi, P. Djapic, and G. Strbac, “Whole-systems assessment of the value of energy storage in low-carbon electricity systems,” *IEEE Transactions on Smart Grid*, vol. 5, pp. 1098–1109, 2014.
- [19] R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss, “Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects,” *Energy Economics*, vol. 31, pp. 269–277, Mar. 2009.
- [20] P. van de Ven, N. Hegde, L. Massoulié, and T. Salonidis, “Optimal control of end-user energy storage,” *IEEE Transactions on Smart Grid*, vol. 4, pp. 789–797, 2013.
- [21] R. Walawalkar, J. Apt, and R. Mancini, “Economics of electric energy storage for energy arbitrage and regulation in New York,” *Energy Policy*, vol. 35, pp. 2558–2568, Apr. 2007.
- [22] J. C. Williams and B. D. Wright, *Storage and Commodity Markets*. Cambridge University Press, 2005.
- [23] A. Bejan, R. Gibbens, and F. Kelly, “Statistical aspects of storage systems modelling in energy networks,” in *46th Annual Conference on Information Sciences and Systems (invited session on Optimization of Communication Networks)*, Princeton University, USA, 2012.
- [24] N. G. Gast, D.-C. Tomozei, and J.-Y. Le Boudec, “Optimal storage policies with wind forecast uncertainties,” in *Greenmetrics 2012*, Imperial College, London, UK, 2012.
- [25] M. Arnold and G. Andersson, “Model predictive control of energy storage including uncertain forecasts,” *Proceedings of the 17th Power Systems Computation Conference*, 2011.
- [26] S. D. Howell, P. W. Duck, A. Hazel, P. V. Johnson, H. Pinto, G. Strbac, N. Proudlove, and M. Black, “A partial differential equation system for modelling stochastic storage in physical systems with applications to wind power generation,” *IMA Journal of Management Mathematics*, vol. 22, pp. 231–252, 2011.
- [27] J. P. Barton and D. G. Infield, “Energy storage and its use with intermittent renewable energy,” *IEEE Transactions on Energy Conversion*, vol. 19, no. 2, pp. 441–448, 2004.
- [28] J. P. Barton and D. G. Infield, “A probabilistic method for calculating the usefulness of a store with finite energy capacity for smoothing electricity generation from wind and solar power,” *Journal of Power Sources*, vol. 162, pp. 943–948, 2006.

- [29] J. R. Cruise, L. C. Flatley, and S. Zachary, “Impact of storage competition on energy markets,” *European Journal of Operational Research*, 2018. To appear.
- [30] G. Lai, F. Margot, and N. Secomandi, “An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation,” *Operations Research*, vol. 58, no. 3, pp. 564–582, 2010.
- [31] O. Q. Wu, D. D. Wang, and Z. Qin, “Seasonal energy storage operations with limited flexibility: the price-adjusted rolling intrinsic policy,” *Manufacturing & Service Operations Management*, vol. 14, no. 3, pp. 455–471, 2012.
- [32] N. Secomandi, “Merchant commodity storage practice revisited,” *Operations Research*, vol. 63, no. 5, pp. 1131–1143, 2015.
- [33] “Electricity balancing services.” National Audit Office Briefing, May 2014.  
<http://www.nao.org.uk/wp-content/uploads/2014/05/Electricity-Balancing-Services.pdf>.
- [34] “The electricity trading arrangements.” ELEXON.  
[https://www.elexon.co.uk/wp-content/uploads/2013/11/beginners\\_guide\\_to\\_trading\\_arrangements\\_v4.0\\_cgi.pdf](https://www.elexon.co.uk/wp-content/uploads/2013/11/beginners_guide_to_trading_arrangements_v4.0_cgi.pdf).
- [35] R. Billinton and R. N. Allan, *Reliability Evaluation of Power Systems*. Springer, 2 ed., 1996.
- [36] “Dinorwig power station.”  
<http://www.electricmountain.co.uk/Dinorwig-Power-Station>.