

Optimal strategies for operating energy storage in an arbitrage market*

Lisa Flatley[†] Robert S MacKay[†] Michael Waterson[‡]

Abstract

We characterize profit-maximizing operating strategies, over some time horizon $[0, T]$, for an energy store which is trading in an arbitrage market. Our theory allows for leakage, operating inefficiencies, operating constraints and general cost functions. In the special case where the operating cost of a store depends only on its instantaneous power output (or input), we present an algorithm to determine the optimal strategies. A key feature is that this algorithm is localized in time, in the sense that the action of the store at a time $t \in [0, T]$ only requires information about electricity prices over some subinterval of time $[t, t_k] \subset [t, T]$. To introduce more complex storage models, we discuss methods for an example which includes minimum switching times between modes of operation.

Keywords. Dynamic optimization, energy storage, optimal localized algorithm.

AMS subject classifications. 49K30, 93C10

1 Introduction

Over the coming decades, the UK energy market faces significant challenges as it strives to meet its climate change targets. Low carbon and renewable generation will need to play a more dominant role in our future electricity supply market. However, renewable energy is intermittent and its availability is driven by uncontrollable elements, such as wind speed and solar intensity. Other options, such as nuclear power and Carbon Capture and Storage, are generally considered to be less flexible than traditional thermal plants. The problems are clear. The unreliability of renewable supply means that there will always be a need for a quick-reacting back-up, in order to ensure that demand is met. On the other hand, at off-peak demand times, supply may be curtailed (often at a high cost) if the system is not flexible enough to respond.

On the demand side, despite improved energy efficiency measures, we can expect to see an overall increase in consumption as the electrification of transport and heating become more commonplace [1]. If all environmental targets are met on time, then National Grid's Gone Green scenario predicts a 12.6% increase in Great Britain's peak demand by 2035, from 2013 levels [2]. In addition, future demand profiles are likely to look very different from today. Electric heat pumps, for example, will create a marked seasonal peaking of demand during the colder, winter months. Electric vehicle charging and new smart technologies, on the other hand, could distort our daily demand patterns.

Electricity storage is one potential option for improving the flexibility and reliability of our electricity system. It could offer services such as peak-shaving, frequency response, reactive power regulation and the provision of reserve. In the UK, large-scale electricity storage currently comes only in the

*This work was supported by EPSRC (grant EP/K002228/1) and the Alfred P. Sloan Foundation New York.

[†]Mathematics Institute, University of Warwick, Gibbet Hill Road, Coventry, CV4 7AL. Corresponding email address: L.Flatley@warwick.ac.uk

[‡]Economics Department, University of Warwick, Gibbet Hill Road, Coventry, CV4 7AL.

form of pumped hydro power plants, which are implemented largely to meet early evening winter peak demands between 4pm and 8pm [3]. Their main sources of revenue are through providing the balancing services, Short-Term Operating Reserve (STOR) and Fast Reserve, which are funded by National Grid. The stores then replenish their supply during the night, when electricity prices are at their lowest [4].

Much research has been dedicated to analyzing storage viability. Some papers focus on the benefits which a store could bring to the electricity system if operated in a socially optimal way (see, for example, [5, 6, 7, 8]). Typically, the approach of these papers is to solve the unit dispatch problem i.e. to select a suitable configuration of generators (including storage) to run at each time, in order to ensure that demand is met (or that demand is met with a certain high probability). An advantage of this method is that it can incorporate the interactions between all assets of the energy system into a single optimization problem, allowing a comparison between the actions of storage against its competing options. A drawback, however, is that the system optimum does not necessarily coincide with each individual firm's optimal strategy. In many cases, it is not even clear that each firm would make a profit under these solutions. Therefore, the whole-system approach is generally better suited to questions where the store is not privately owned, but instead owned and operated by a central controller, such as the system operator.

Other papers focus on the profits available to a store which faces stochastic or probabilistic prices. Some of these papers restrict the behaviour of the store to a pre-defined set of operating strategies (e.g. [9, 10]). Another approach is to implement Dynamic Programming techniques (e.g. [11, 12]). A drawback of this latter approach, however, is that it tends to be computationally heavy and requires information about prices over the entire time horizon over which we wish to optimize. Additionally, such methods do not give much scope for mathematical insight into the dynamics of the solutions.

In this paper, we present a method to determine the maximum profit available to an electricity store which is operating in the wholesale electricity market. We assume that the store can predict an electricity price function $p : [0, T] \rightarrow \mathbb{R}$ (or alternatively, a cost function) where $[0, T]$ is the period of time over which we wish to optimize. However, the method only ever uses price information over a smaller time interval, thus reducing the required prediction horizon and the amount of computation needed. It is worth commenting here that the prediction of future prices in this way is not an unreasonable assumption, since most electricity today is traded through "over-the-counter" bilateral contracts, which are agreed ahead of operating time.

Our method is derived from standard Calculus of Variations techniques and is intended as an extension to [13]. Our work differs from [13] in several respects. Firstly, we present our model in a continuous time setting, rather than the discrete setting employed in [13]. Even if prices are declared at discrete time intervals, in accordance with our current market system, our approach allows for the input of piecewise constant prices but with continuous variation in the operation of the store. Secondly, [13] allows only for convex cost functions, whereas here we allow for much more general operating costs and, in particular, we remove the convexity assumption. This is realistic: in general, there is little reason to believe that the cost of running a motor, for example, is a convex function of the output power. Moreover, there may be cost jumps involved in switching on a new motor if additional power is required. Although these more general conditions do not guarantee the existence of solutions, we prove that optimal operating strategies of the form given in Proposition 2.1 exist if and only if the associated algorithm does not terminate early. Finally, we allow the inclusion of additional operating constraints. If these constraints can be expressed as bounds on the instantaneous power output (e.g. in the case of enforced periods of closure for maintenance), then they can easily be incorporated into the main result (Proposition 2.1) and the optimal strategy can be determined using the algorithm of Section 3.2. This method reduces the optimization problem to a set of smaller problems which are localized in time. For other constraints, however, this localization in time is not guaranteed. In Section 4.1, we demonstrate this issue and propose a method for handling an additional condition if the store must wait for a fixed amount of time (called the minimum switching time) whenever it wishes to change its mode of operation between the charging and the discharging mode.

The structure of the paper is as follows: Over the remainder of this section, we introduce our storage model and its associated costs. In Section 2, we present the main result, which characterizes the optimal strategies via a reference price function. In Section 3, we present an algorithm which determines both the optimal strategy and the reference price for the basic case where the store is only constrained by its power ratings and its capacity constraints. We prove that an optimal strategy exists in this case only if the algorithm does not terminate early. Section 4 contains a discussion of a simple storage model and in Section 4.1 we investigate the inclusion of a minimum switching time constraint. In Section 5, we present and discuss some illustrative results. Finally, the conclusions of the paper are summarized in Section 6.

1.1 The storage model

We use a simple technology-agnostic model for storage, which incorporates a number of key features, both physical and economic. To allow for multiple applications, we define all quantities in a general setting. Towards the end of the paper, in Section 4, we present a concrete (and technically relevant) example.

Let $[0, T]$ be the interval of time over which we want to optimize the actions of the store, for some $T > 0$. At each time $t \in [0, T]$, we denote by $U(t) \subset \mathbb{R}$ the set of admissible power outputs associated with the store at that time. The operator of the store may then choose an *operating strategy* $q : [0, T] \rightarrow \mathbb{R}$ which allocates, at each time t , the amount of power $q(t) \in U(t)$ that is to enter the store (using the convention that if $q(t) < 0$, then the store discharges $-q(t)$ units of power at time t). Allowing the power constraint set $U(t)$ to vary with time in this way means that we may incorporate constraints such as planned periods of closure or transmission congestion into the model. A planned closure over a time period $I \subset [0, T]$, for example, corresponds to the power constraint set $U(t) = \{0\}$ for all $t \in I$.

The amount of stored energy, or the *level of the store*, which results from the strategy q at time t is denoted $\ell[q](t)$ and evolves according to the differential equation

$$\frac{d}{dt}\ell[q](t) = -\alpha\ell[q](t) + q(t). \quad (1)$$

Here, $\alpha \in [0, 1)$ is the leakage rate, which is introduced to reflect the self-discharge which occurs due to imperfect sealing or porous casing. Given an initial condition $\ell[q](0) = \ell_0$, the level of the store at each time $t \in [0, T]$ is thus given by

$$\ell[q](t) = e^{-\alpha t}\ell_0 + e^{-\alpha t} \int_0^t e^{\alpha s} q(s) ds. \quad (2)$$

We will only consider strategies $q : [0, T] \rightarrow \mathbb{R}$ which are piecewise continuous functions of time. This is a physically reasonable assumption, since any strategy which has an accumulation point of discontinuities or a non-jump discontinuity would be hard to achieve by any piece of equipment. Throughout the remainder of this paper, we denote by $\mathcal{C}_{pw}(I)$ the space of piecewise continuous functions $u : I \rightarrow \mathbb{R}$, for any subset $I \subset \mathbb{R}$.

The *capacity constraints* of the store are characterized by two functions $E^+, E^- : [0, T] \rightarrow \mathbb{R}$, so that any strategy $q : [0, T] \rightarrow U$ is constrained by the inequalities

$$E^-(t) \leq \ell[q](t) \leq E^+(t) \quad \forall t \in [0, T].$$

We assume that the initial and terminal levels of the store are pre-specified, so that

$$E^-(0) = E^+(0) \quad \text{and} \quad E^-(T) = E^+(T).$$

Note that, if not specified, then the terminal condition is in fact implicit since any optimal strategy q^* should satisfy that $\ell[q^*](T) = E^-(T)$. We also assume, without loss of generality, that $\{(d/dt)E^-(t) +$

$\alpha E^-(t), (d/dt)E^+(t) + \alpha E^+(t)\} \subset U(t)$ for all $t \in [0, T]$ (otherwise E could be adjusted so that this is true, without changing the set of admissible strategies).

The set of possible levels of the store, at each time, is represented by the *admissible energy domain* $E \subset [0, T] \times \mathbb{R}$, which is the set of all pairs $(t, m) \in [0, T] \times [E^-(t), E^+(t)]$. Often, one may choose to replace the capacity constraints E^+ and E^- with constants, so that $E^+ > 0$ is the physical size of the store and $E^- \geq 0$ is the minimum technically feasible level of the store (which may be strictly positive, for example in the case of compressed air energy storage, where a cushion of air needs to be maintained in the store at all times in order to provide sufficient pressure to instigate discharging). We have chosen here to represent E in this more general form, to allow the store to participate in multiple markets. For example, if a store decides to participate in a Balancing Mechanism reserve contract, it may be required by National Grid to make supply available over certain agreed periods. Thus, at these contracted times the store will have less capacity available for use in the wholesale market.

In addition to the capacity and power constraints outlined above, the operation of a store may be restricted by further factors, such as warm-up times, ramping constraints and minimum switching times between modes of operation. We denote the collection of additional constraints associated with the store by Γ . The only condition on these constraints is that they must be entirely expressible as functions of the power output function q . A maximum ramp-up rate $r > 0$ on the charge side, for example, may be written as inequalities

$$\dot{q}(t) \leq r \quad \forall t \in [0, T] \text{ such that } q(t) > 0.$$

We do not allow for constraints of this kind in Section 3. However, the impact of a non-empty set Γ is discussed in Section 4.

The total cost incurred by the store, given an operating strategy q , is denoted $C[q] \in \mathbb{R}$. Here, C is a functional which maps $\mathcal{C}_{pw}([0, T])$ into \mathbb{R} . This can incorporate factors such running costs, warming-up costs and costs of storing, as well as the cost of purchasing power and the payments (counted as negative cost) received for providing power. The aim of this paper is to identify those operating strategies q which minimize the total cost $C[q]$, whilst adhering to all of the physical constraints outlined above. If the store is profitable to run, then the minimal cost should, of course, be negative.

Definition 1.1 (Admissible strategies and optimal strategies). *We say that $q : [0, T] \rightarrow \mathbb{R}$ is an admissible strategy if q is piecewise continuous and satisfies the following properties:*

1. For each $t \in [0, T]$, we have $q(t) \in U(t)$ and $\ell[q](t) \in [E^-(t), E^+(t)]$.
2. If $\Gamma \neq \emptyset$, then the strategy q satisfies each of the constraints in Γ .

We denote by X the set of all admissible strategies and say that $q^* \in X$ is an optimal strategy if

$$C[q^*] \leq C[q] \quad \forall q \in X. \quad (3)$$

2 Characterization of optimal strategies

The following proposition provides a characterization of optimal strategies in terms of a ‘‘reference price’’ function μ^* . We introduce the *indicator functional* $S : \mathcal{C}_{pw} \rightarrow \{0, 1\}$, which is defined by the property that $S[q]$ takes the value 0 if and only if q satisfies each of the constraints in Γ .

Proposition 2.1 (Characterization of optimal strategies). *Suppose there exists a constant $\rho^* \in \mathbb{R}$, a piecewise differentiable function $\mu^* : [0, T] \rightarrow \mathbb{R}$ and a strategy $q^* \in X$ with the following properties:*

- (i) The strategy q^* is a minimizer of

$$C[q] - \int_0^T e^{\alpha t} \mu^*(t) q(t) dt + \rho^* S[q] \quad (4)$$

over all piecewise continuous functions $q : [0, T] \rightarrow \mathbb{R}$ such that $q(t) \in U(t)$ for all $t \in [0, T]$.

(ii) If μ^* is differentiable at $t \in [0, T]$, then the following complementary slackness conditions are satisfied:

$$\frac{d\mu^*}{dt}(t) = 0 \quad \text{if } E^-(t) < \ell[q^*](t) < E^+(t), \quad (5)$$

$$\frac{d\mu^*}{dt}(t) \geq 0 \quad \text{if } \ell[q^*](t) = E^+(t), \quad (6)$$

$$\frac{d\mu^*}{dt}(t) \leq 0 \quad \text{if } \ell[q^*](t) = E^-(t). \quad (7)$$

(iii) If μ^* is not differentiable at $t \in [0, T]$, then the following ‘‘jump’’ complementary slackness conditions are satisfied:

$$\mu^*(t^+) - \mu^*(t^-) = 0 \quad \text{if } E^-(t) < \ell[q^*](t) < E^+(t), \quad (8)$$

$$\mu^*(t^+) - \mu^*(t^-) \geq 0 \quad \text{if } \ell[q^*](t) = E^+(t), \quad (9)$$

$$\mu^*(t^+) - \mu^*(t^-) \leq 0 \quad \text{if } \ell[q^*](t) = E^-(t), \quad (10)$$

where $\mu^*(t^-)$ and $\mu^*(t^+)$ are the left and right limits respectively of μ^* at t .

Then q^* is an optimal strategy.

Proof. Let $q \in X$ and let $0 = a_1 < \dots < a_n < a_{n+1} = T$ be a partition such that $\cup_{i=1}^n [a_i, a_{i+1}] = [0, T]$ and q^*, q are continuous and μ^* is differentiable over each (a_i, a_{i+1}) . Notice that (1) can be equivalently written as

$$\frac{d}{dt} e^{\alpha t} \ell[q](t) = e^{\alpha t} q(t) \quad \forall t \in [0, T].$$

Hence, (4) implies

$$\begin{aligned} C[q^*] - C[q] &\leq \int_0^T e^{\alpha t} \mu^*(t) (q^*(t) - q(t)) dt + \underbrace{\rho^* (S[q^*] - S[q])}_{=0} \\ &\stackrel{(1)}{=} \sum_{i=1}^n \int_{a_i}^{a_{i+1}} \mu^*(t) \left\{ \frac{d}{dt} (e^{\alpha t} \ell[q^*](t) - e^{\alpha t} \ell[q](t)) \right\} dt \\ &= \sum_{i=1}^n \left[e^{\alpha a_{i+1}} \mu^*(a_{i+1}^-) (\ell[q^*](a_{i+1}) - \ell[q](a_{i+1})) - e^{\alpha a_i} \mu^*(a_i^+) (\ell[q^*](a_i) - \ell[q](a_i)) \right] \\ &\quad - \sum_{i=1}^n \int_{a_i}^{a_{i+1}} e^{\alpha t} \frac{d\mu^*}{dt}(t) (\ell[q^*](t) - \ell[q](t)) dt \\ &= \sum_{i=2}^n e^{\alpha a_i} (\mu^*(a_i^-) - \mu^*(a_i^+)) (\ell[q^*](a_i) - \ell[q](a_i)) \\ &\quad - \sum_{i=1}^n \int_{a_i}^{b_i} e^{\alpha(t)} \frac{d\mu^*}{dt}(t) (\ell[q^*](t) - \ell[q](t)) dt \end{aligned}$$

where we have used the fact that, since q and q^* are admissible, then $S[q^*] = S[q] = 0$. The second equality follows from integration by parts and final equality results from a rearrangement of the first sum on the previous line, together with the condition that $\ell[q](0) = \ell[q^*](0) = E^-(0)$ and $\ell[q](T) = \ell[q^*](T) = E^-(T)$ if $q \in X$. Applying the complementary slackness conditions (5)-(10) to the final equality above, we obtain $C[q^*] - C[q] \leq 0$. \square

The reference price μ^* provides a reference cost per unit of power. Roughly, if the cost per unit of power for operating the store is lower than the reference price, then this indicates the store should charge at that time (and similarly for discharging).

It is worth mentioning here that exactly the same result follows if we apply the Karush-Kuhn-Tucker (KKT) conditions to the capacity constraints of the original minimization problem (3). By following such an approach, one instead looks for minimizers of the relaxed functional

$$C[q] - \int_0^T \left(\lambda_1(t) (\ell[q](t) - E^-(t)) + \lambda_2(t) (E^+(t) - \ell[q](t)) \right) dt + \rho^* S[q]$$

over all piecewise-continuous $q : [0, T] \rightarrow \mathbb{R}$ such that $q(t) \in U(t)$ for all $t \in [0, T]$, where $\lambda_1, \lambda_2 : [0, T] \rightarrow \mathbb{R}$ are piecewise-continuous functions which satisfy the complementary slackness conditions

$$\lambda_1(t) (\ell[q](t) - E^-(t)) = \lambda_2(t) (E^+(t) - \ell[q](t)) = 0 \quad \forall t \in [0, T].$$

The relation between the two approaches is that the reference price function μ^* in Proposition 2.1 can be expressed as an integral of the KKT functions:

$$\mu^*(t) = \int_t^T (\lambda_1(s) - \lambda_2(s)) ds$$

and, in particular, $(d\mu^*/dt)(t) = -\lambda_1(t) + \lambda_2(t)$.

The following lemma shows that, in certain cases, we should expect the determination of minimizers q^* to reduce to a sequence of smaller minimization problems which are localized in time.

Lemma 2.2. *[Localization of the minimization problem] Let $E : [0, T] \rightarrow \mathbb{R}$ be the admissible domain, $U(t) \subset \mathbb{R}$ be the set of admissible power outputs at each time $t \in [0, T]$ and let C be a cost functional of the form*

$$C[q] = \int_0^T L(t, q(t)) dt \tag{11}$$

for some running cost function $L : \mathcal{A} \rightarrow \mathbb{R}$, where \mathcal{A} is the set of time-power pairs $(t, x) \in [0, T] \times \mathbb{R}$ such that $x \in U(t)$. For each $t \in [0, T]$, let $\ell_1(t) = E^-(t)$ and $\ell_2(t) = E^+(t)$. Assume that there exists an optimal strategy which satisfies the conditions of Proposition 2.1 and suppose that the optimal level of the store at some time $\sigma_0 \in [0, T]$ is known to be $\ell = \ell_i(\sigma_0)$, with $i \in \{1, 2\}$. For any $\sigma \in (\sigma_0, T]$, let q_σ solve the minimization problem

$$\int_{\sigma_0}^\sigma L(t, q_\sigma(t)) dt \leq \int_{\sigma_0}^\sigma L(t, q(t)) dt \tag{12}$$

over all piecewise continuous $q : [\sigma_0, \sigma] \rightarrow \mathbb{R}$ such that $q(t) \in U(t)$ for all $t \in [0, T]$ and such that the end conditions $\ell[q](\sigma) = \ell_i(\sigma)$ and $\ell[q](\sigma_0) = \ell$ hold. If there exists a time $t_0 \in (\sigma_0, \sigma)$ such that $\ell[q_\sigma](t_0) = \ell_j(t_0)$, with $j \in \{1, 2\} \setminus \{i\}$, then there exists an optimal strategy $q^* \in X$ which coincides exactly with q_σ over the interval (σ_0, t_0) .

Proof. Assume first that $\ell = \ell_1(\sigma_0)$ so that $\ell[q_\sigma](t_0) = \ell_2(t_0)$ and $\ell[q_\sigma](\sigma) = \ell_1(\sigma)$. Let $q^* \in X$ be an optimal strategy, let q_σ satisfy the hypotheses above and suppose that the claim of the lemma is not true. Then $\ell[q^*](\sigma) > \ell_1(\sigma)$ (since otherwise we could replace q^* with q_σ over $[0, \sigma]$ to either reduce the total cost or keep the cost the same). Thus there must exist at least one interval $(a, b) \subset (\sigma_0, \sigma)$, with $b < \sigma$, such that $\ell[q^*](a) - \ell[q_\sigma](a) = \ell[q^*](b) - \ell[q_\sigma](b) = 0$ and $\ell[q^*](t) \neq \ell[q_\sigma](t)$ for all $t \in (a, b)$. However, this implies that both q^* and q_σ must be minimizers of $\int_a^b L(t, q(t)) dt$ over all admissible $q : [a, b] \rightarrow U$ such that $\ell[q](a) = \ell[q^*](a)$ and $\ell[q](b) = \ell[q^*](b)$. In particular, we may adapt q^* if necessary so that $q^*(t) = q_\sigma(t)$ for all $t \in (a, b)$. Repeating this argument for all such intervals (a, b) completes the proof that the optimal strategy can be chosen to coincide with q_σ over (σ_0, t_0) . A similar argument holds for the case where $\ell = \ell_2(\sigma_0)$. □

The significance of the above lemma is that q_σ depends only on knowledge of a previous time σ_0 when the store is either full or empty, and on the action of the cost functional over the time interval $[\sigma_0, \sigma]$. In particular, starting at time 0, one sets $\ell = E^-(0) = E^+(0) = 0$ and finds a time $\sigma \in (0, T]$ such that the minimizer q_σ of (12) hits the upper capacity constraint at some time in $t_0 \in (0, \sigma)$. The optimal strategy $q^* \in X$ can thus be chosen to coincide with q_σ over the interval $[0, t_0]$. We then set $\ell = E^+(t_0)$ and $\sigma_0 = t_0$ and repeat the process. Continuing in this way, we eventually construct a global optimizer q^* .

3 Determination of optimal strategies with no additional constraints

We assume in this section that the only constraints faced by the store are its capacity and power constraints (i.e. $\Gamma = \emptyset$). For simplicity of notation, we assume from now on that there is a fixed set of power constraints $U \subset \mathbb{R}$ such that $U(t) = U$ for all $t \in [0, T]$. All of the results extend naturally to the more general case where U is allowed to vary with time. We present an algorithm to determine μ^* , and consequently the optimal strategy q^* , for the class of functionals C which take the form (11) for some $L : [0, T] \times U \rightarrow \mathbb{R}$ with $L(t, 0) = 0$ for all $t \in [0, T]$. We assume that, for each $t \in [0, T]$, the map $\lambda \mapsto L(t, \lambda)$ is piecewise differentiable and that the associated partial derivative $\partial L(t, \lambda)/\partial \lambda$ has a continuous inverse at almost every $\lambda \in U$. This class of functionals contains costs which depend only on the time t and the power output of the store at that time. The piecewise differentiability assumption allows for jumps in the cost. These jumps may relate to the cost of switching on an additional motor in order to provide a higher power output, for example.

Proposition 2.1 states that, if we can find the appropriate reference price function μ^* , then the optimal strategy $q^* \in X$ solves

$$\mathcal{L}[q^*, \mu^*] \leq \mathcal{L}[q, \mu^*] \quad \forall q : [0, T] \rightarrow U, \quad (13)$$

where for any functions $q, \mu : [0, T] \rightarrow \mathbb{R}$, we define the relaxed functional

$$\mathcal{L}[q, \mu] = \int_0^T \left(L(t, q(t)) - e^{\alpha t} \mu(t) q(t) \right) dt. \quad (14)$$

The regularity assumptions on L , q and μ imply that (13) reduces to a minimization problem which is pointwise in time. Specifically, for almost all $t \in [0, T]$, we search for $q^*(t) \in U$ which solves

$$L(t, q^*(t)) - e^{\alpha t} \mu^*(t) q^*(t) \leq L(t, w) - e^{\alpha t} \mu^*(t) w \quad \forall w \in U. \quad (15)$$

Our algorithm is a generalization of that provided in [13]. An important result is Proposition 3.5, which states that if the algorithm does not terminate early, then it does indeed provide an optimal strategy q^* . In particular, in the special case that the cost functional C is strictly convex, then there exists a unique optimal strategy, and this coincides exactly with q^* .

3.1 Preliminary results and definitions

By Proposition 2.1, we aim to find a piecewise differentiable $\mu^* : [0, T] \rightarrow \mathbb{R}$ such that the optimal strategy $q^* \in X$ solves (13). Moreover, μ^* should be constant over intervals of time where the level of the store is away from the capacity constraints. This motivates the following construction:

Given $t \in [0, T]$ and $\lambda \in \mathbb{R}$, we want to define a quantity $u(t, \lambda) \in \mathbb{R}$ as a solution of

$$u(t, \lambda) = \arg \min_{w \in U} \left(L(t, w) - e^{\alpha t} \lambda w \right). \quad (16)$$

If we know that $\mu \equiv \lambda$ over some connected interval $I \subset [0, T]$, then we hope that $q^*(t) = u(t, \lambda)$ for all $t \in I$. Care needs to be taken, however, because there may be multiple minimizers associated with

λ . With this in mind, we denote by $\mathcal{M}_t \subset \mathbb{R}$ the set of λ which admit multiple minimizers in (16). Note that, if L is assumed to be strictly convex in its second argument, then minimizers of (4) are of course unique, implying that each $\mathcal{M}_t = \emptyset$.

Lemma 3.1. *For each $(t, \lambda) \in [0, T] \times \mathbb{R}$, let $u(t, \lambda)$ be any solution of (16). Then, at each $t \in [0, T]$, the mapping $\lambda \mapsto u(t, \lambda)$ is monotone increasing and piecewise continuous. Moreover, the set \mathcal{M}_t is finite and discontinuities in $\lambda \mapsto u(t, \lambda)$ occur only at points in \mathcal{M}_t .*

Proof. Let $t \in [0, T]$ and suppose that the map $\lambda \mapsto u(t, \lambda)$ is not monotone increasing. Then, there exists $\lambda_1 < \lambda_2$ such that $u_1 := u(t, \lambda_1) > u(t, \lambda_2) =: u_2$. But then,

$$L(t, u_1) - e^{\alpha t} \lambda_2 u_1 \leq L(t, u_2) - e^{\alpha t} \lambda_1 u_2 - e^{\alpha t} (\lambda_2 - \lambda_1) u_1 < L(t, u_2) - e^{\alpha t} \lambda_2 u_2.$$

The first inequality follows from the definition of u_1 and u_2 as solutions to (16), and the second inequality follows from the supposition. However, the above contradicts the definition of u_2 , and we conclude that the map $\lambda \mapsto u(t, \lambda)$ is monotone increasing at all $t \in [0, T]$.

The piecewise continuity of the map $\lambda \mapsto u(t, \lambda)$ follows immediately from the regularity assumptions on L . Precisely, if for almost all $\lambda \in \mathbb{R}$, the minimizers $u(t, \lambda)$ lie away from discontinuities of L , then they must satisfy that

$$\frac{\partial L}{\partial \lambda}(t, \lambda) - e^{\alpha t} \lambda = 0 \quad \text{a.e. } \lambda \in \mathbb{R}.$$

The continuous invertibility assumption on the partial derivative of L therefore implies the piecewise continuity of the map $\lambda \mapsto u(t, \lambda)$. If, on the other hand, there is a non-degenerate subset $W \subset \mathbb{R}$ such that if $\lambda \in W$ then $u(t, \lambda)$ lies at a discontinuity of L , then the above monotonicity property and the piecewise continuity of L imply that the map $\lambda \mapsto u(t, \lambda)$ must be piecewise constant over W . In either case, the map $\lambda \mapsto u(t, \lambda)$ is piecewise continuous.

Finally, we prove the finiteness of the set \mathcal{M}_t by supposing that the opposite is true. To this end, let $\mathcal{A} \subset \mathbb{R}$ be an infinite set such that, for each $\lambda \in \mathcal{A}$, there are two distinct local solutions $x(\lambda), y(\lambda) \in U \setminus \partial U$ to (16), where ∂U denotes the boundary of the set U . Let $\lambda_0 \in \mathcal{A}$ be such that there exists a sequence $(\lambda_n)_{n \in \mathbb{N}}$ in \mathbb{R} with $\lim_{n \rightarrow \infty} \lambda_n = \lambda_0$, and set $x(\lambda_0) = x_0$ and $y(\lambda_0) = y_0$. Assume without loss of generality that

$$L_u(t, x(\lambda)) = L_u(t, y(\lambda)) = e^{\alpha t} \lambda$$

is satisfied at $\lambda = \lambda_0$, where $L_u(t, \cdot)$ denotes the partial derivative of L with respect to the second argument, and that $x(\lambda)$ and $y(\lambda)$ lie away from any discontinuity of L . Setting $g(\lambda) := L(t, x(\lambda)) - \lambda x(\lambda)$ and $h(\lambda) := L(t, y(\lambda)) - \lambda y(\lambda)$, we obtain

$$g'(\lambda_0) = L_u(t, x_0) x'(\lambda_0) - x'(\lambda_0) \lambda_0 - x_0 = (L_u(t, x_0) - \lambda_0) x'(\lambda_0) - x_0 = -x_0$$

and similarly, $h'(\lambda_0) = -y_0$. Hence, if x_0 and y_0 are both global minimizers solving (16) but $x_0 \neq y_0$, then there exists $N \in \mathbb{N}$ such that $g(\lambda_n) \neq h(\lambda_n)$ for all $n \geq N$. In particular, $x(\lambda_n)$ and $y(\lambda_n)$ cannot both be minimizers for $n \geq N$. This contradicts the assumption that there exists such a set \mathcal{A} . \square

The definitions which follow will be employed in the algorithm, and are written under the assumption that there exists a pair (μ^*, q^*) which satisfies the complementary slackness conditions of Proposition 2.1. Condition (5) motivates the following definition.

Definition 3.2 ((λ, τ) -admissibility). *For any choice of $\lambda \in \mathbb{R}$ and $\tau \in [0, T]$, we say a piecewise continuous function $x : [0, T] \rightarrow U$ is (λ, τ) -admissible if $x \equiv 0$ over $[0, \tau]$ and if, for each $t \in [0, T]$, $x(t)$ is a solution of*

$$x(t) = \arg \min_{w \in U} \left(L(t, w) - e^{\alpha t} \lambda w \right). \quad (17)$$

In particular, the mapping $t \mapsto x(t)$ coincides with $t \mapsto u(t, \lambda)$ over $[\tau, T]$, for some choice of map u which solves (16) at each $t \in [0, T]$.

Note that, in general, a (λ, τ) -admissible strategy x is not admissible: for a general choice of λ , one of the capacity constraints is likely to be broken at some time in $(\tau, T]$. The idea of the algorithm is to piece together (λ, τ) -admissible strategies in order to construct the optimal strategy $q^* \in X$. The reference price μ^* attains the appropriate value of λ over intervals of time when μ^* is constant. In the special case that one can find a $(\lambda, 0)$ -admissible strategy which is also admissible, then of course we are done, since we may set $\mu^* \equiv \lambda$. Such a strategy satisfies condition (5) at every $t \in [0, T]$.

The key feature of the algorithm is thus to determine: (i) the intervals of time over which μ^* is constant and (ii) the value that μ^* attains over these intervals. Intuitively, we expect the store to completely fill whenever costs are sufficiently low. The store will then wait for costs to rise before fully discharging. Correspondingly, the reference price should increase over this waiting period, from λ_1 say to λ_2 . If we did not allow for this increase in reference price, then the associated strategy would cause the store to begin discharging too soon and to empty too much, thus breaking the lower capacity constraint in order to take advantage of the low costs. This behaviour is an indication that λ_1 is the correct choice of reference price over the charging period. This motivates the following characterization of strategies.

Definition 3.3 (Characterization of (λ, τ) -admissible strategies). *Let $\tau \in [0, T]$, $\lambda \in \mathbb{R}$ and $m \in [0, M]$. We characterize each (λ, τ) -admissible $x : [0, T] \rightarrow U$ according to sets $X(\tau, m)$ and $X'(\tau, m)$ as follows:*

(i) *We write $x \in X(\tau, m)$ if there exists $t \in (\tau, T)$ and $\varepsilon_0 > 0$ such that*

$$m + \ell[x](t) = E^-(t) \quad \text{and} \quad m + \ell[x](t + \varepsilon) < E^-(t + \varepsilon)$$

for all $\varepsilon \in (0, \varepsilon_0)$, with

$$E^-(t') \leq m + \ell[x](t') \leq E^+(t') \quad \forall t' \in (\tau, t). \quad (18)$$

In other words, if the level of the store at time τ is m , and if the store adopts the strategy x over $(\tau, T]$, then the first violation of a capacity constraint occurs at the lower bound of E . We write $t = t[x, m]$.

(ii) *Similarly, we write $x \in X'(\tau, m)$ if there exists $t \in (\tau, T)$ and $\varepsilon_0 > 0$ such that condition (18) holds and*

$$m + \ell[x](t) = E^+(t) \quad \text{and} \quad m + \ell[x](t + \varepsilon) > E^+(t + \varepsilon)$$

for all $\varepsilon \in (0, \varepsilon_0)$. In other words, if the level of the store at time τ is m , and if the store adopts the strategy x over $(\tau, T]$, then the first violation of a capacity constraint occurs at the upper bound of E . We write $t = t[x, m]$.

(iii) *We say $\lambda \in \Lambda(\tau, m)$ if there exists a (λ, τ) -admissible $x \in X(\tau, m)$. Similarly, $\lambda \in \Lambda'(\tau, m)$ if there exists a (λ, τ) -admissible $x \in X'(\tau, m)$.*

Lemma 3.1 implies that $\Lambda(\tau, m)$ and $\Lambda'(\tau, m)$ are each connected subintervals of \mathbb{R} which satisfy

$$\sup \Lambda(\tau, m) \leq \inf \Lambda'(\tau, m) \quad \forall (\tau, m) \in [0, T] \times [0, M]. \quad (19)$$

In particular, the interiors of the two sets are disjoint:

$$\text{int}(\Lambda(\tau, m)) \cap \text{int}(\Lambda'(\tau, m)) = \emptyset \quad \forall (\tau, m) \in [0, T] \times [0, M]. \quad (20)$$

3.2 The algorithm

We are now in a position to outline the algorithm. To this end, fix $\tau \in [0, T]$ and suppose that μ^* and q^* are known over $[0, \tau]$. We assume without loss of generality that $m := \ell[q^*](\tau) \in \{E^-(\tau), E^+(\tau)\}$ (otherwise, move τ backwards until this is satisfied). The algorithm will identify an $N \in \mathbb{N}$ and two increasing sequences of times $\{\tau_i\}_{i=1}^N, \{\sigma_i\}_{i=1}^N$ such that $\tau_i \leq \sigma_i$ for each i , and such that $(d/dt)\mu^*(t) = 0$ whenever $t \in (\tau_i, \sigma_i)$. The reference price μ^* is only allowed to jump in value at times of the form τ_i or σ_i . We may assume without loss of generality that $\tau = \tau_k$, for some $k \in \{1, \dots, N\}$ (otherwise, again, move τ backwards until this is true). Recall from Definition 3.3 that $t[x, m]$ is the time at which a strategy x first breaks a capacity constraint, given that $x(0) = m$.

Step 1: If there exists $\lambda \in \mathbb{R}$ and a (λ, τ_k) -admissible x such that $x \notin X(\tau_k, m) \cup X'(\tau_k, m)$, then set $\sigma_k = T$. Set $\mu^*(t) = \lambda$ for all $t \in (\tau, T]$ and define the restriction of q^* to $(\tau_k, T]$ to be x . **The algorithm is complete.**

If there is no λ satisfying these conditions, proceed to step 2.

Step 2: Set $\lambda := \sup \Lambda(\tau_k, m)$. There are three cases. For each case, we define the pair (μ^*, q^*) restricted to an interval $(\tau_k, \sigma_k] \subset (\tau_k, T]$.

- a) If $\lambda \in \Lambda(\tau, m) \setminus \Lambda'(\tau, m)$: Choose a (λ, τ_k) -admissible x such that $m + \ell[x](\sigma_k) = E^+(\sigma_k)$ for some $\sigma_k \in [\tau_k, t[x, m])$, and select the latest such σ_k . Set $\mu^*(t) = \lambda$ for all $t \in (\tau_k, \sigma_k]$ and define the restriction of q^* to $(\tau_k, \sigma_k]$ to coincide with x . If $\sigma_k < T$, proceed to step 3; if $\sigma_k = T$, proceed to step 4.
- b) If $\lambda \in \Lambda'(\tau, m) \setminus \Lambda(\tau, m)$: Choose a (λ, τ_k) -admissible x such that $m + \ell[x](\sigma_k) = E^-(\sigma_k)$ for some $\sigma_k \in [\tau_k, t[x, m])$, and select the latest such σ_k . Set $\mu^*(t) = \lambda$ for all $t \in (\tau_k, \sigma_k]$ and define the restriction of q^* to $(\tau_k, \sigma_k]$ to coincide with x . If $\sigma_k < T$, proceed to step 3; if $\sigma_k = T$, proceed to step 4.
- c) If $\lambda \in \Lambda(\tau, m) \cap \Lambda'(\tau, m)$: Choose a (λ, τ_k) -admissible x such that either $m + \ell[x](\sigma_k) = E^+(\sigma_k)$ or $m + \ell[x](\sigma_k) = E^-(\sigma_k)$ for some $\sigma_k \in [\tau_k, t[x, m])$, and select the latest such σ_k . Set $\mu^*(t) = \lambda$ for all $t \in (\tau_k, \sigma_k]$ and define the restriction of q^* to $(\tau_k, \sigma_k]$ to coincide with x . If $\sigma_k < T$, proceed to step 3; if $\sigma_k = T$, proceed to step 4.

Step 3: Let τ_{k+1} be the first time in $[\sigma_k, T]$ such that, on relabeling τ_{k+1} as τ_k and $m + \ell[x](\tau_{k+1})$ as m , and on determining the new associated σ_k from steps 1 and 2, we have $\sigma_k \neq \tau_k$. Relabel τ_{k+1} as τ_k and return to step 1.

If there is no such τ_{k+1} , then we have found all intervals (τ_i, σ_i) over which $(d/dt)\mu^* = 0$. Proceed to step 4.

Step 4: Over each interval $(\sigma_i, \tau_{i+1}]$, corresponding to the times found in the steps above, and for each $t \in (\sigma_i, \tau_{i+1}]$, define

$$q^*(t) = \begin{cases} \frac{d}{dt}E^+(t) + \alpha E^+(t) & \text{if } \ell[q](\sigma_i) = E^+(\sigma_i) \\ \frac{d}{dt}E^-(t) + \alpha E^-(t) & \text{if } \ell[q](\sigma_i) = E^-(\sigma_i). \end{cases}$$

The algorithm is complete if, for each $i \in \{1, \dots, N\}$, there exists a piecewise differentiable $\mu_i : (\sigma_i, \tau_{i+1}) \rightarrow \mathbb{R}$ such that the following conditions hold:

- $q^*(t)$ is a minimizer of $L(t, x) - \mu_i(t)x$ over all $x \in \mathbb{R}$.
- If $\ell[q^*](\sigma_i) = E^+(\sigma_i)$, then μ_i is non-decreasing over (σ_i, τ_{i+1}) and satisfies that $\mu_i(\sigma_i^+) \geq \mu^*(\sigma_i^-)$ and $\mu_i(\tau_{i+1}^-) \leq \mu^*(\tau_{i+1}^+)$.
- If $\ell[q^*](\sigma_i) = E^-(\sigma_i)$, then μ_i is non-increasing over (σ_i, τ_{i+1}) and satisfies that $\mu_i(\sigma_i^+) \leq \mu^*(\sigma_i^-)$ and $\mu_i(\tau_{i+1}^-) \geq \mu^*(\tau_{i+1}^+)$.

If these conditions do not hold, the algorithm ends here and yields no solution q^* .

End of algorithm

The above algorithm implicitly reduces the original optimization problem to a series of new optimization problems which are localized in time. At each time τ_k , one only needs to look ahead to a time t_k to know how to operate the store over (τ_k, σ_k) , where t_k is defined as $t[x, m]$ for the appropriate choice of x and m from step 2 of the algorithm (or as the minimum value of $t[x, m]$ whenever there is a choice of x). Hence, whilst we originally assumed knowledge of the entire price function $p : [0, T] \rightarrow [0, \infty)$, in practice we may not need so much information.

The following result states that, provided the algorithm does not terminate early, then it does indeed give an optimal strategy q^* . Proposition 3.5 then states the converse: if there is an optimal strategy, which satisfies the conditions of Proposition 2.1, then the algorithm will find it (in other words, the algorithm will not terminate early).

Proposition 3.4. *If the above algorithm yields a strategy $q^* : [0, T] \rightarrow U$, then q^* is an optimal strategy.*

Proof. Let (μ^*, q^*) be the pair which is determined through the algorithm and let $\{\tau_i\}_{i=1}^N$ and $\{\sigma_i\}_{i=1}^T$ be the associated sequences of time, so that μ^* is constant over each interval (τ_i, σ_i) . Assume that the conditions of Proposition 2.1 are satisfied up until time $\tau_k \in [0, T]$. At each $i \in \{1, \dots, N\}$, we have

$$\sup \Lambda(\tau_i, \ell[q^*](\tau_i)) \leq \mu^*(\tau_i^+) \leq \inf \Lambda'(\tau_i, \ell[q^*](\tau_i)). \quad (21)$$

If the conditions of step 1 are satisfied at τ_k , then q^* is clearly admissible and satisfies the properties of the proposition. Hence, q^* is optimal and the algorithm is complete.

Assume therefore that the conditions of step 1 do not hold at τ_k so that $\sup \Lambda(\tau_k, m) = \inf \Lambda'(\tau_k, m)$. It is clear that the strategy defined by the algorithm is admissible, and it remains to check that conditions (5)-(10) are satisfied by μ^* .

To this end, note that, the algorithm ensures that the conditions are satisfied over the intervals (τ_k, σ_k) and $(\sigma_k, \tau_{k+1}]$, and it remains to check that the conditions hold at σ_k . Suppose first that case a) of step 2 is satisfied at time τ_k , so that $\ell[q^*](\sigma_k) = E^+(\sigma_k)$. Then, by construction, we have

$$\mu^*(\sigma_k^-) \in \Lambda(\sigma_k, E^+(\sigma_k)). \quad (22)$$

Thus, if $\sigma_k = \tau_{k+1}$ then, together with (21), this implies

$$\mu^*(\sigma_k^-) \leq \sup \Lambda(\sigma_k, E^+(\sigma_k)) \leq \mu^*(\sigma_k^+) \quad (23)$$

and condition (9) is satisfied at σ_k . If, on the other hand, $\sigma_k < \tau_{k+1}$, then the construction of μ^* through the algorithm immediately implies that $\mu^*(\sigma_k^-) \leq \mu^*(\sigma_k^+)$.

Similarly, if case b) of step 2 holds at τ_k , then $\ell[q^*](\sigma_k) = E^-(\sigma_k)$ and

$$\mu^*(\sigma_k) \in \Lambda'(\sigma_k, E^-(\sigma_k)). \quad (24)$$

Thus, if $\sigma_k = \tau_{k+1}$ then, together with (21), this implies

$$\mu^*(\sigma_k) \geq \inf \Lambda'(\sigma_k, E^-(\sigma_k)) \geq \mu^*(\sigma_k^+) \quad (25)$$

and condition (10) is satisfied at σ_k . If, on the other hand, $\sigma_k < \tau_{k+1}$, then the construction of μ^* through the algorithm immediately implies that $\mu^*(\sigma_k^-) \geq \mu^*(\sigma_k^+)$.

Finally, if case c) holds at τ , then the above two cases together imply that the conditions of Proposition 2.1 are satisfied. \square

Proposition 3.5. *If the algorithm terminates early, then there is no pair (μ^*, q^*) which satisfies the conditions of Proposition 2.1.*

Proof. Let (μ^*, q^*) be a pair generated by the algorithm. (If the algorithm terminates early, at a time T_0 say, then μ^* and q^* are both considered as functions over the domain $[0, T_0]$.) Let $\{\tau_i\}_{i=1}^{N^*}$ and $\{\sigma_i\}_{i=1}^{N^*}$ be the sequences of time generated by the algorithm, so that μ^* is constant over each (τ_k, σ_k) . We write $m_k^* := \ell[q^*](\tau_k)$ and assume that each interval (τ_k, σ_k) is maximal in the sense that if $(\tau, \sigma) \supset (\tau_k, \sigma_k)$ is a strictly larger interval, then μ^* is not constant over this interval.

Suppose also that there is a pair (μ, q) which satisfies the conditions of Proposition 2.1. Analogously, let $\{t_i\}_{i=1}^N$ and $\{s_i\}_{i=1}^N$ be increasing sequences which define the end-times of all maximal intervals (t_k, s_k) over which μ is constant. We write $m_k := \ell[q](t_k)$ and

$$\lambda_k := \mu(t) \quad \forall t \in (t_k, s_k).$$

We make the following additional assumptions on the pair (μ, q) : Suppose there is a pair of times $\tau, \tau' \in [0, T]$ such that $\tau < \tau'$ and $\ell[q](t) \in \{E^-(t), E^+(t)\}$ for each $t \in \{\tau, \tau'\}$. If there exists a pair (μ', q') such that (i) the conditions of Proposition 2.1 are satisfied, (ii) $\ell[q'](\tau) = \ell[q](\tau)$ and $\ell[q'](\tau') = \ell[q](\tau')$, and (iii) $(d/dt)\mu'(t) = 0$ for all $t \in (\tau, \tau')$, then we assume that (μ, q) agrees with (μ', q') over the larger interval $[0, \tau']$. In other words, we always take μ to be constant whenever there is the choice.

Throughout this proof, we will use the notation $\lambda \sim_k \lambda'$ if the set of (λ, t_k) -admissible functions coincides exactly with the set of (λ', t_k) -admissible functions over the interval (t_k, s_k) . This assumption is only required to handle the degenerate cases where q is identically zero over (t_k, s_k) . In these cases, there may be a choice of values to assign to μ , each of which would result in the same (λ_k, t_k) -admissible q over (t_k, s_k) . We introduce the ordering $\lambda \prec_k \lambda'$ if $\lambda \asymp_k \lambda'$ and $\lambda < \lambda'$.

The proof consists of proving that the following statements are true at each $k \in \{1, \dots, N\}$:

1. $\lambda_k \sim_k \sup \Lambda(t_k, m_k)$ and q is (λ_k, t_k) -admissible over the interval (t_k, s_k) .
2. There exists $j \in \{1, \dots, N^*\}$ such that $t_k = \tau_j$.
3. If $j \in \{1, \dots, N^*\}$ and $t_k = \tau_j$, then $s_k = \sigma_j$ and $\ell[q](t_k) = \ell[q](\sigma_k)$.

The first statement is required to prove statement 2, and implies that the algorithm determines the correct value of μ over each interval (t_k, s_k) . Together, statements 2 and 3 then complete the proof, since they imply that the algorithm uncovers all intervals (t_k, s_k) over which μ is constant. Moreover, statement 3 implies that the strategy uncovered by the algorithm agrees with q^* at the end point of each of these intervals. Hence, in between each σ_k and τ_{k+1} , the strategies q and q^* necessarily lie on the capacity constraints and so must coincide. In particular, the existence of the pair (μ, q) implies that the algorithm does not terminate early but instead generates a pair (μ^*, q^*) which are defined over $[0, T]$.

Proof of statement 1: Let $k \in \{1, \dots, N\}$. It is clear from the definition of (λ_k, t_k) -admissibility and from the requirement that q must be a minimizer of (15), that q must be (λ_k, t_k) -admissible over the interval (t_k, s_k) . We need to show that

$$\lambda_k \sim_k \sup \Lambda(t_k, m_k) =: \theta_k.$$

Recall that the maps $\lambda \mapsto u(t, \lambda)$ of Lemma 3.1 are piecewise continuous and monotone increasing at each time $t \in [0, T]$. These maps are constructed in such a way that one can always find a suitable $u(t, \lambda_k)$ to coincide with q over (t_k, s_k) . Thus, the monotonicity of u with respect to its second argument implies the relations:

$$\beta \prec_k \theta_k \quad \Rightarrow \quad \beta \in \Lambda(t_k, m_k) \setminus \Lambda'(t_k, m_k) \quad (26)$$

$$\beta \succ_k \theta_k \quad \Rightarrow \quad \beta \in \Lambda'(t_k, m_k) \setminus \Lambda(t_k, m_k). \quad (27)$$

Suppose now that $\lambda_k \prec_k \theta_k$. If $m_k + \ell[q](s_k) = E^+(s_k)$, then this immediately implies that $\lambda_k \sim_k \theta_k$ since, otherwise, the continuity and monotonicity of u would imply the existence of β such that

$\lambda_k \prec_k \beta \prec_k \theta_k$ and a (β, t_k) -admissible q' such that $m_k + \ell[q'](t) \geq E^+(t)$ for all $t \in [t_k, s_k]$, with strict inequality at $t = s_k$. This would imply that $\beta \in \Lambda'(t_k, m_k)$, which contradicts (26).

If, on the other hand, $m_k + \ell[q](s_k) = E^-(s_k)$, then Proposition 2.1 requires that $\mu(s_k^+) - \mu(s_k^-) \leq 0$. If μ is constant over an interval $(s_k, t') \subset (s_k, T]$, then the monotonicity property of Lemma 3.1 implies that the lower capacity constraint will be broken by q , and no admissible solution will be found. If μ is not constant over any such interval, then Proposition 2.1 requires that $q(t) = u(t, \mu(t)) = (d/dt)E^-(t) + \alpha E^-(t)$ for all $t \in (s_k, t')$. By the monotonicity of u , this is only possible if $(d/dt)\mu(t) > 0$ at some $t \in (s_k, t')$ or if $\mu(t'^+) - \mu(t'^-) > 0$, since otherwise the fact that $\lambda_k \in \Lambda(t_k, m_k)$ implies that q will break the lower capacity constraint. However, this contradicts condition (7) of the proposition. Hence, we must have $\lambda_k \sim_k \theta_k$ or $\lambda_k \succ_k \theta_k$. If $\lambda_k \succ_k \theta_k$, then similar arguments as above arrive at analogous contradictions to Proposition 2.1 and we conclude that we must therefore have $\lambda_k \sim_k \theta_k$. This completes the proof of statement 1.

Proof of statement 2: Suppose first that $\tau = t_k$ but that there is no $j \in \{1, \dots, N^*\}$ such that $\tau = \tau_j$. Then, statement 1 implies that $\lambda_k \sim_k \sup \Lambda(t_k, m_k) = \theta_k$ and that q must coincide with some (t_k, λ_k) -admissible x over (t_k, σ_k) . However, since τ does not coincide with any τ_j , the conditions of steps 1 and 2 of the algorithm do not hold at time τ . If $x \in X(t_k, m_k)$, then Proposition 2.1 implies that $s_k < T$ and $\ell[q](s_k) = E^-(s_k)$ (since the conditions of step 2 do not hold). Thus, Proposition 2.1 implies that either (i) $s_k = t_{k+1}$ and $\mu(s_k^+) - \mu(s_k^-) \leq 0$, or (ii) $s_k < t_{k+1}$ and $\mu(t) - \mu(s_k^-) \leq 0$, for all $t \in [s_k, t_{k+1}]$. In case (i), the monotonicity property of Lemma 3.1 implies that q breaks the lower capacity constraint. In case (ii), since x breaks the lower capacity constraint, we must have that $x(s) < (d/ds)E^-(s) + \alpha E^-(s)$ for all s in some interval $I \subset [s_k, t_{k+1}]$. However, since we require that $q(s) = (d/ds)E^-(s) + \alpha E^-(s)$ for all $s \in (s_k, t_{k+1})$, this means that we must have $(d/ds)\mu(s) > 0$ for all s in some subset $I' \subset I$. This contradicts the conditions of Proposition 2.1, and we conclude that there exists $j \in \{1, \dots, N^*\}$ such that $\tau = \tau_j$. The same result follows from similar arguments if $x \in X'(t_k, m_k)$. This completes the proof of statement 2.

Proof of statement 3: Let $\tau = \tau_j = t_k$. First notice that the proof of statement 2 yielded a contradiction when we assumed that q coincided with a (λ_k, t_k) -admissible $x \in X(t_k, m_k)$ such that $s_k < T$ and $\ell[q](s_k) = E^-(s_k)$. Since statement 1 implies that q must coincide with some (λ_k, t_k) -admissible x over (t_k, s_k) , it follows that, if $x \in X(t_k, m_k)$, then either $s_k = T$ or $\ell[q](s_k) = E^+(s_k)$. Similarly, if $x \in X'(t_k, m_k)$, then either $s_k = T$ or $\ell[q](s_k) = E^-(s_k)$. The assumption that the intervals (t_k, s_k) are maximal thus implies that s_k coincides with σ_j and that $\ell[q](s_k) = \ell[q^*](\sigma_j)$, provided that σ_j is independent of any choices made at steps 1 or 2 of the algorithm. This independence is not immediately clear since the general assumptions on the cost function L mean that there may be more than one choice of (λ_k, τ) -admissible strategy.

To prove the independence of σ_j from the choice of strategy made in the algorithm at steps 1 or 2, suppose that there are two (λ_k, τ) -admissible strategies x^1 and x^2 which satisfy the conditions of step 1 or 2 of the algorithm. Let σ^i denote the time σ_j if strategy x^i is chosen at time τ and assume without loss of generality that $\sigma^1 \leq \sigma^2$. The monotonicity property of Lemma 3.1 makes it immediately clear that, if we choose x^1 then the algorithm will select $\mu^*(t) = \lambda_k$ for all $t \in (\sigma^1, \sigma^2)$. Hence, μ^* is constant over (τ, σ^2) . In particular, if there are several choices of strategy to choose at time τ , and if σ^2 is the latest time σ_j associated with any of the choices, then the assumption that each interval (τ_j, σ_j) is maximal implies that $\sigma_j = \sigma^2$, regardless of the strategy chosen. This completes the proof of statement 3. □

4 A simple storage model

We introduce here a simple but practically relevant model of a store, which operates purely within a single wholesale market. The set of admissible powers U is taken to be of the form $U = [-q_{\max}^-, q_{\max}^+]$, with $-q_{\max}^- < 0 < q_{\max}^+$. The store is assumed to have a fixed maximum energy capacity $M > 0$, so that the admissible domain satisfies $E^-(t) = 0$ and $E^+(t) = M$ for all $t \in (0, T)$. We further assume

that we require the store to be empty at both the start and end times, so that $E^-(0) = E^+(0) = E^-(T) = E^+(T) = 0$. The cost functional C is defined by

$$C[q] = \int_0^T w(q(t))p(t)q(t), \quad (28)$$

where $p(t) \in \mathbb{R}$ is the price for one unit of power at time t . Here, $w : U \rightarrow (0, \infty)$ is a piecewise continuous function which acts as a weighting to the cost of buying and selling electricity, and should satisfy that $w(x) \geq 1$ if $x \geq 0$; and $0 < w(x) \leq 1$ if $x \leq 0$. More intuitively, w can be thought of as the result of inefficiencies during the charging and discharging processes of the store: in order to fill the store at a rate $q(t)$ at time t , the storage operator would actually need to purchase a greater amount of electricity, since power is lost during the charging process. Similarly, whenever an amount $q(t)$ is discharged from the store, a lesser amount is actually available to sell on the market. For simplicity, we assume that w takes on two values $w_1 \geq 1$ and $w_2 \in (0, 1]$, so that

$$w(x) = \begin{cases} w_1 & \text{if } x \geq 0 \\ w_2 & \text{if } x < 0. \end{cases} \quad (29)$$

In the special case where there are no additional operating constraints (i.e. $\Gamma = \emptyset$), this cost functional is convex and thus an optimal strategy exists. Proposition 3.5 states moreover that an optimal strategy can be determined via the algorithm of Section 3.2. Some simple results are immediately available, as stated in the following lemma.

Lemma 4.1. *Define the parameters of a store as above and let the cost functional C be given by (28). Assume that $p(t) \geq 0$ for all $t \in [0, T]$ and let $\ell_0, \ell_T \geq 0$ be the required levels of stored energy at times 0 and T respectively. The following properties hold:*

1. *If $\ell_T \geq \ell_0$ and $\dot{p} \equiv 0$, then there is no non-zero strategy $q \in X$ which yields a positive profit.*
2. *The minimum cost, over all admissible strategies, is a decreasing function of w_2, M, q_{\max}^- and q_{\max}^+ . It is an increasing function of α and w_1 .*

Proof. If $q \in X$ is not the zero function, and if $p(t) = p$ is constant for all $t \in [0, T]$, then

$$C[q] = p \int_0^T w(q(t))q(t) dt \geq p \int_0^T q(t) dt = p(\ell_T - \ell_0) + \alpha \int_0^T \ell[q](t) dt \geq 0.$$

This proves statement 1. To prove statement 2, let q^* be the optimal strategy associated with the parameters $\theta = (\alpha, w_1, w_2, M, q_{\max}^-, q_{\max}^+)$ and write $V(\theta) = C[q^*]$. It is easy to see that q^* is also admissible with respect to the parameters $\theta + \delta\gamma$, for any $\delta > 0$ and any $\gamma = (\gamma_1, \dots, \gamma_6)$ with $\gamma_1 = 0$, $\gamma_2 \in \{-1, 0\}$ and $\gamma_3, \dots, \gamma_6 \in \{0, 1\}$. Thus, if $w_\gamma(x)$ is defined by $w_\gamma(x) = \theta_2 + \delta\gamma_2$ whenever $x \geq 0$ and $w_\gamma(x) = \theta_3 + \delta\gamma_3$ whenever $x < 0$, then

$$V(\theta + \delta\gamma) \leq \int_0^T w_\gamma(q^*(t))p(t)q^*(t) dt \leq \int_0^T w(q^*(t))p(t)q^*(t) dt = V(\theta).$$

This proves the required growth conditions of the minimum cost with respect to w_1, w_2, M, q_{\max}^- and q_{\max}^+ . To complete the proof, let $\alpha' < \alpha$ and denote by $\ell'[q](t)$ the level of stored energy at time t associated with the strategy q and the leakage rate α' . Set $q'(t) = q^*(t)$ whenever $q^*(t) \leq 0$ or $\ell'[q'](t) < M$. For all other $t \in [0, T]$, set $q'(t) = \min(\alpha M, q^*(t))$. Then q' is admissible with respect to the new parameter vector $\theta' = \theta + (\alpha' - \alpha, 0, 0, 0, 0, 0)$. Hence,

$$V(\theta) = C[q^*] = C[q'] + \int_0^T w(q(t))p(t)(q^*(t) - q'(t)) dt \geq C[q'] \geq V(\theta')$$

and the minimum cost function is increasing with respect to the leakage rate α . \square

4.1 A storage model with minimum switching times

Throughout this section, we employ the notation t^+ and t^- to mean $\max\{t, 0\}$ and $\min\{t, T\}$ respectively for any $t \in \mathbb{R}$. We investigate the impact of introducing a non-empty set of additional constraints Γ to the simple storage model introduced above. In this setting, Lemma 2.2 no longer holds in general. It is therefore not clear that there will exist a method with similar time-localization properties as the algorithm of Section 3.2. We demonstrate this problem explicitly in this section. In particular, we impose a minimum switching time constraint $\tau > 0$ on the simple storage model introduced above. This is the minimum amount of time for which the store must be idle when switching between the charging and discharging modes of operation. Hence the additional constraint set Γ consists of the requirement that any power output q must satisfy that

$$q(t)q(s) \geq 0 \quad \forall t \in [0, T] \text{ and } s \in (t - \tau, t + \tau) \cap [0, T]. \quad (30)$$

The minimization problem of Proposition 2.1 is difficult to solve in its current form, but we may rewrite the problem by adapting the Lagrangian \mathcal{L} of (14) to incorporate the minimum switching time constraint as follows:

$$\Lambda[q, \mu, \lambda] := \underbrace{\int_0^T \left(w(q(t))p(t) - e^{\alpha t} \mu(t) \right) q(t) dt}_{=\mathcal{L}[q, \mu]} - \int_0^T \int_t^{(t+\tau)^-} \lambda(t, s) q(s) q(t) ds dt$$

for all piecewise continuous $q, \mu : [0, T] \rightarrow U$ and $\lambda : [0, T]^2 \rightarrow [0, \infty]$. In order to apply Proposition 2.1, we need to reduce Λ into a form which is compatible with (4). To do this, we carefully choose λ to depend on q and μ as follows: For a given μ , let $q_\mu : [0, T] \rightarrow U$ solve

$$\left(w(q_\mu(t))p(t) - e^{\alpha t} \mu(t) \right) q_\mu(t) \leq \left(w(x)p(t) - e^{\alpha t} \mu(t) \right) x \quad \forall x \in U. \quad (31)$$

Since the efficiency function w takes the simple form (29), it is clear that $q_\mu(t)$ takes a value in $\{-q_{\max}^-, 0, q_{\max}^+\}$ whenever $w_1 p(t) - e^{\alpha t} \mu(t) \neq 0$ and $w_2 p(t) - e^{\alpha t} \mu(t) \neq 0$ at $t \in [0, T]$. In these cases $q_\mu(t)$ is uniquely defined. If one of these relations is an equality instead, however, then $q_\mu(t)$ may take any value in U . Hence, there may be multiple solutions associated with μ . Let now λ minimize

$$\int_0^T \left| w(q_\mu(t))p(t) - e^{\alpha t} \mu(t) - \int_{(t-\tau)^+}^{(t+\tau)^-} \lambda(t, s) q_\mu(s) ds \right| dt \quad (32)$$

over all $\lambda : [0, T]^2 \rightarrow [0, \infty]$ such that (i) $\lambda(t, s) = 0$ whenever $q_\mu(s)q_\mu(t) \geq 0$ and (ii) the inequality

$$\left(w(q_\mu(t))p(t) - e^{\alpha t} \mu(t) - \int_{(t-\tau)^+}^{(t+\tau)^-} \lambda(t, s) q_\mu(s) ds \right) \left(w(q_\mu(t))p(t) - e^{\alpha t} \mu(t) \right) \geq 0 \quad (33)$$

holds for all $t \in [0, T]$. We write $\lambda = \lambda(q_\mu)$.

The significance of the above construction is that an optimal strategy can now be found by searching for minimizers of Λ . Specifically, if a strategy q^* satisfies that

$$\Lambda[q^*, \mu^*, \lambda^*] \leq \Lambda[q, \mu^*, \lambda^*] \quad (34)$$

for all piecewise continuous $q : [0, T] \rightarrow U$, and if (i) (q^*, μ^*) satisfies the conditions (5)–(10), (ii) λ^* is of the form $\lambda^* = \lambda(q_{\mu^*})$ and (iii) q^* is admissible, then q^* is an optimal strategy. To see this, choose ρ^* large enough that the relation

$$\rho^* \geq q_{\max}^+ q_{\max}^- \int_0^T \int_t^{(t+\tau)^-} \lambda^*(t, s) ds dt \quad (35)$$

holds. It is clear from the assumption that q^* is admissible that $S[q^*] = 0$ and it follows from (34) that q^* minimizes $\mathcal{L}[q, \mu^*] + \rho^* S[q]$ over all admissible $q \in X$. Assumption (35) moreover implies that q^* minimizes $\mathcal{L}[q, \mu^*] + \rho^* S[q]$ over all piecewise continuous $q : [0, T] \rightarrow U$ because, for any non-admissible q , we have

$$\mathcal{L}[q, \mu^*] + \rho^* \underbrace{S[q]}_{=1} \geq \mathcal{L}[q, \mu^*] - \int_0^T \int_t^{(t+\tau)^-} \lambda^*(t, s) q(s) q(t) ds dt = \mathcal{L}[q^*, \mu^*] + \rho^* \underbrace{S[q^*]}_{=0}.$$

Thus, Proposition 2.1 implies that q^* is optimal.

Condition (iii) of the above argument requires that the minimizer q^* of (34) is admissible, in order for q^* to be an optimal strategy. As in the algorithm of Section 3.2, a suitable choice of μ^* will ensure that the capacity constraints are satisfied, whilst the power rating constraints are immediately absorbed into the minimization problem (34). The following lemma ensures that, as long as a minimizer q^* exists, then we may select q^* so that the switching time constraints are also satisfied.

Lemma 4.2. *Let $\mu : [0, T] \rightarrow \mathbb{R}$ be piecewise differentiable and let λ be of the form $\lambda = \lambda(q_\mu)$ for some minimizer q_μ of (31). Let Y be the set of piecewise continuous $q^* : [0, T] \rightarrow U$ such that $\Lambda[q^*, \mu, \lambda] \leq \Lambda[q, \mu, \lambda]$ for all piecewise continuous $q : [0, T] \rightarrow U$. If $Y \neq \emptyset$, then there exists $q^* \in Y$ which satisfies the switching time constraints (30).*

Proof. Suppose the claim is not true. Then there exists $t, s \in [0, T]$ such that $|t - s| < \tau$ and $q^*(t)q^*(s) < 0$. Thus $\lambda(t, s)q^*(t)q^*(s) \leq 0$. Note that

$$f[q, \mu, \lambda](t) := \frac{\delta \Lambda}{\delta q(t)}[q, \mu, \lambda] = w(q(t))p(t) - e^{\alpha t} \mu(t) - \int_{(t-\tau)^+}^{(t+\tau)^-} \lambda(t, s) q(s) ds,$$

where $\delta \Lambda / \delta q(t)$ denotes the functional derivative of Λ with respect to $q(t)$. In particular, the term $\lambda(t, s)$ only influences the value of q^* at times t and s . The condition that λ minimizes (32) subject to the constraint (33) implies that $f[q^*, \mu, \lambda](t) = 0$ or $f[q^*, \mu, \lambda](s) = 0$. In particular, we may always choose a q^* so that $q^*(t) = 0$ or $q^*(s) = 0$. \square

An interpretation of the role of λ^* in the minimization problem (34) is that the function acts to adjust q_{μ^*} to a new strategy which satisfies the switching time constraints. It does this by finding all pairs of times (t, s) such that $|s - t| < \tau$ and $q_{\mu^*}(s)q_{\mu^*}(t) < 0$, and adjusting $q_{\mu^*}(t)$ or $q_{\mu^*}(s)$ to 0. The condition that $\lambda^*(t, s) = 0$ whenever $q_{\mu^*}(t)q_{\mu^*}(s) \geq 0$ ensures q_{μ^*} is only adjusted at times when the minimum switching time constraint (30) is broken. Choosing λ^* to minimize (32) subject to the constraint (33) ensures that λ^* makes these adjustments in such a way that, if $a(t)$ is the absolute change in instantaneous cost from the original strategy q_{μ^*} at time t , then the accumulation of these cost changes $\int_0^T a(t) dt$ is minimal. Thus, an equivalent and generally easier method for determining the optimal strategy q^* , given μ^* , is to find a suitable subset of times $I \subset [0, T]$ and set $q^*(t) = 0$ if $t \in I$ and $q_{\mu^*}(t) = q_{\mu^*}(t)$ if $t \in [0, T] \setminus I$. The subset I should be chosen to minimize the associated absolute cost change integral $\int_0^T a(t) dt$.

If μ^* is given, then the minimization problem (34) is much easier to solve than (4). The main task in determining an optimal strategy q^* , therefore, is to choose the correct reference price function μ^* . As alluded to at the beginning of Section 4, this is a more difficult task now that the set of additional constraints Γ is non-empty. To illustrate the complexity of the problem, consider the following scenario:

Let the power ratings and capacity constraints of the store be given by $q_{\max}^- = q_{\max}^+ = 1$ and $M = 2$ and let the minimum switching time of the store be $\tau = 1$. Assume that there is no leakage and the store is perfectly efficient so that $\alpha = 0$ and $w_1 = w_2 = 1$. Suppose that $T = 3$ and that the price of electricity p is a function of time given by $p(t) = p_i$ for all $t \in [i - 1, i)$ and for each $i \in \{1, 2, 3\}$, where the values p_i satisfy $p_3 < p_1 < p_2$. Suppose further that $\frac{1}{2}(p_1 + p_2) > p_3$. It is not

difficult to see that the optimal strategy q^* satisfies $q^*(t) = 1$ if $t \in [0, \frac{1}{2}]$ and $q^*(t) = -1$ if $t \in [\frac{3}{2}, 2]$, with $q^*(t) = 0$ elsewhere. The corresponding reference price μ^* satisfies $\mu^*(t) = \mu_1 := \frac{1}{2}(p_1 + p_2)$ if $t \in [0, 2)$ and $\mu^*(t) = \mu_2 = p_3$ otherwise. However, the correct choice of μ_1 here relies on knowing the correct choice of μ_2 ; unlike in the algorithm of Section 3.2, the constant sections of the reference price function cannot be determined independently. To see this, we simply remark that, if we had guessed at a constant reference price function $\mu_0^* \equiv \mu_1$, then we would arrive at the corresponding solution q_0^* with $q_0^*(t) = 1$ if $t \in [0, 1) \cup [2, 3]$ and $q_0^*(t) = 0$ otherwise. This strategy never reaches the upper or lower capacity constraint after the initial time, and so there is no chance to update our reference price function in accordance with Proposition 2.1. Even though our guess μ_0^* is correct over the time interval $[0, 2)$, it is only possible to know this when considering the value that μ^* attains over the time interval $[2, 3]$.

This issue prevents the methods of Section 3.2 from working here. However, as long as the store is required to perform at least one full charge and discharge cycle then the determination of a suitable μ^* still reduces to a localized problem thanks to Lemma 2.2. Finally, we remark that, if μ^* is a reference price function and if $q_{\mu^*}^*$ is the associated strategy, constructed via the minimization (34), then it can easily be shown that at each time $t \in [0, T]$, the level of the store $\ell[q_{\mu^*}^*]$ is increasing in μ^* . Whilst the algorithm of Section 3.2 no longer works in this more complicated setting, we can still use this monotonicity in order to reduce the choice of μ^* at each time to a suitable range, using similar arguments as in the case with no additional constraints.

5 Illustrative results

The following graphs plot the optimal strategies for a store which follows the model set out in Section 4, with $q_{\max}^+ = q_{\max}^-$ and $M/q_{\max}^+ = 7$. We take $\alpha = 0$, $w_1 = 1$ and vary w_2 as labeled below. The bottom two figures plot the level of the store $\ell[q^*](t)$ at each time t , assuming that the optimal strategy q^* is followed. The top figures both plot the price of electricity $p : [0, T] \rightarrow [0, \infty)$ which the store faces in making its operating decision. Here, we have chosen hourly prices from the first 230 hours of N2EX's day-ahead auction in November 2013.

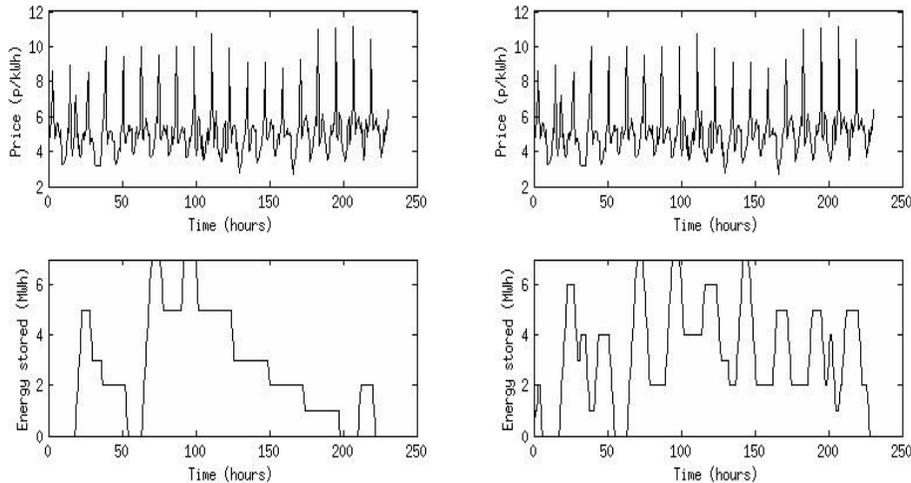


Figure 1: November 2013 hourly day-ahead auction prices (top left and top right); optimal fill level of the store with $w_2 = 0.5$ (bottom left) and $w_2 = 0.7$ (bottom right).

An immediate observation is that the store cycles more frequently as its efficiency w_2 increases. Less energy is lost during operation at higher efficiencies and the store is able to earn revenues over a larger proportion of elapsed time. We highlight here that, although the price of electricity follows

a reasonably periodic pattern (with periods of roughly a day in length), it does not follow that the level of the store is the same at the start and end of each day (see, for example, [16] for a proof of the potential non-periodicity of solutions in a periodic setting). This underlies an important motivation for choosing this method over a dynamic programming approach. Using a dynamic programming method, one works backwards in time and evaluates, at each $t \in [0, T]$, the value function

$$V(t, x) := \min_{q \in X} \int_t^T L(q(s), s) ds,$$

where the minimization is subject to the constraints $q(t) = x$. However, this requires information about all prices over the time $[0, T]$. If T is large (say, $T = 40$ years, the average life of a pumped hydro store), then the number of calculations becomes infeasible. One option is to split $[0, T]$ into a union of smaller disjoint intervals $I_1 \cup \dots \cup I_N = [0, T]$, for some $N \in \mathbb{N}$, and to find the optimal strategy over each of these smaller intervals. However, this requires setting an end-state for each interval, which may not be known. A reasonable guess in the above examples would be to assume that the store is empty at some off-peak time during each night but, as seen, such an assumption would lead to a sub-optimal solution. Using our approach of Section 3.2, the algorithm has the property that it implicitly reduces our problem to a series of new optimization problems, which are localized in time.

6 Conclusions

We have presented a method for determining the operating strategy for a store to maximize its arbitrage profits. Our setting allows for leakage, inefficiencies and general operating costs which are functions of the power output. This setting also allows for time-varying constraints on the power output of the store. A significant benefit associated with this method is the implicit localization in time of the solution. Moreover, we have shown that there only exists an optimal strategy of the form of Proposition 2.1 if the algorithm does not terminate early. As an extension, we have discussed the inclusion of more general operating constraints and proposed a method for determining optimal strategies in the presence of minimum switching time constraints.

We believe that the theory put forward in this paper serves as a good starting point for evaluating the profits available to a store. The assumption that prices can be predicted over suitable periods of time is a good approximation to the situation where a store trades through bilateral contracts, or through an auctioning market. A next step would be to consider the store as a larger player, whose actions have an impact on the price of electricity, and it is believed that the approaches given in this paper can be extended to this new setting.

Acknowledgments

This work is made possible by the IMAGES (Integrated Market-fit Affordable Grid-Scale Energy Storage) research group. The authors would like to thank the other members of the group for their suggestions and insight. Particular thanks go to Monica Giulietti, Jihong Wang, Xing Luo, Andrew Pimm and Seamus Garvey. The authors would also like to thank Stan Zachary, James Cruise and Richard Gibbens for their many helpful discussions related to this topic.

References

- [1] *Planning our electric future: a White Paper for secure, affordable and low-carbon electricity*, Department of Energy and Climate Change, 2011
- [2] *UK Future Energy Scenarios*, National Grid, 2014

- [3] *The future role for energy storage in the UK: main report*, Energy Research Partnership, 2011
- [4] D.J.C MacKay: *Sustainable Energy - without the hot air*, UIT Cambridge, 2009
- [5] *Energy Storage and Management Study*, AEA, 2010
- [6] *Options for low-carbon power sector flexibility to 2050 - a report to the Committee on Climate Change*, Pöyry, 2010
- [7] M. Black, G. Strbac: *Value of Bulk Energy Storage for Managing Wind Power Fluctuations*, IEEE Transactions on Energy Conversion, 22(1) (2007), pp197–205
- [8] M. Aunedi, N. Brandon, D. Jackravut, D. Predrag, D. Pujianto, R. Sansom, G. Strbac, A. Sturt, F. Teng, V. Yufit: *Strategic Assessment of the Role and Value of Energy Storage Systems in the UK Low Carbon Energy Future - Report for Carbon Trust*, 2012
- [9] J.P. Barton, D.G. Infield: *Energy Storage and Its Use With Intermittent Renewable Energy*, IEEE Transactions on Energy Conversion, 19(2) (2004), pp441–448
- [10] P. Grünewald, T. Cockerill., M. Contestabile., P. Pearson.: *The role of large scale storage in a GB low carbon energy future: Issues and policy challenges*, Energy Policy, 39 (2011), pp4807-4815
- [11] *Study of Compressed Air Energy Storage with Grid and Photovoltaic Energy Generation, Draft Final Report - for Arizona Public Service Company*, Arizona Research Institute for Solar Energy (AZTISE), 2010
- [12] N. Löhndorf, S. Minner: *Optimal day-ahead trading and storage of renewable energies - an approximate dynamic programming approach*, Energy Syst., 1 (2010), pp61–77
- [13] J. Cruise, L. Flatley, R. Gibbens, S. Zachary: *Optimal control of storage for arbitrage, with applications to energy systems* (2014) arXiv:1307.0800v1; under consideration with MOR
- [14] A.J. Pimm, S.D. Garvey: *The economics of hybrid energy storage plant*, International Journal of Environmental Studies, (2014), pp1-9
- [15] A.J. Pimm, A.D. Garvey, B.K. Kantharaj: *Economic analysis of a hybrid energy storage system based on liquid air and compressed air*, (2014) Under consideration with Energy Economics.
- [16] R.S. MacKay, S. Slijepčević, J. Stark: *Optimal scheduling in a periodic environment*, Nonlinearity, (2000) 13, pp257–297