Optimal strategies for operating energy storage in an arbitrage market

Lisa Flatley, Robert S MacKay , Michael Waterson

No 1048

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

Optimal strategies for operating energy storage in an arbitrage market

Lisa Flatley[1,3], Robert S MacKay[1], Michael Waterson[2]

June 9, 2014

Abstract

We characterise profit-maximising 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 and general cost functions. In the special case where the operating cost of a store depends only on its instantaneous power ouput (or input), we present an algorithm to determine the optimal strategies. A key feature is that this algorithm is localised 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 $[t, \tau] \subset [t, T]$.

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. Renewable energy, however, 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 if the system is not flexible enough to respond. Curtailment often comes at a high cost: in April 2011, for example, National Grid resorted to payments of up to £800/MWh in order to reduce electricity production from Scottish wind farms (compared to £28/MWh for the cheapest coal power station not to run). Flexibility of supply, therefore, is key to the successful and cost-effective integration of renewable generation into the energy system.

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 eventually become commonplace strategies for reducing carbon emissions [1]. If all environmental targets are met on time, then National Grid's Gone Green scenario predicts a 10% increase in Great Britain's peak demand, from current levels, by 2030; whilst meeting targets ahead of time, as represented by the Accelerated Growth scenario, could result in a 25% peak demand increase [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, on the other hand, is likely to distort our daily demand patterns. Being a relatively new technology, it is not clear to what extent electric transport will be deployed over the coming years and, moreover, there is the question of whether smart appliances will help to shift battery charging to off-peak times, or whether daily price variations will be too small to influence people's behaviour.

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 balancing and the availability of reserve. In the UK, electricity storage currently comes only in the 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

[1]Mathematics Institute, University of Warwick, Gibbet Hill Road, Coventry, CV4 7AL.
[2]Economics Department, University of Warwick, Gibbet Hill Road, Coventry, CV4 7AL.
[3]Corresponding author: L.Flatley@warwick.ac.uk

replenish their supply during the night, when electricity prices are at their lowest [4]. Hence, there is an underlying daily periodicity in the operation of such a store, which results from the periodicity and predictability of supply, demand and prices.

As discussed above, however, our future supply and demand patterns are set to look very different from today and it seems unlikely that an electricity store in a few decades' time will face such strong periodicity of prices. An important question is whether storage could be a financially viable option in this new setting.

Much research has been dedicated to this question over the last few years. Some papers focus on the social benefits which a store could bring to the electricity system (see, for example, [6, 7, 8, 9]). 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 optimisation 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. [10, 11]). Another approach is to implement Dynamic Programming techniques (e.g. [12, 13, 14]). A drawback of this latter approach, however, is that in order to evaluate the maximum profit, one needs information about prices over the entire time horizon over which we wish to optimise. Even with approximation methods to reduce the number of computations required, this is still a computationally heavy technique if we are considering the entire lifetime of the store. 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. The main assumption is that the store can predict an electricity price function $p : [0, T] \to \mathbb{R}$, where $[0, T]$ is the period of time over which we wish to optimise. This is not an unreasonable assumption, since most electricity today is traded through "over-the-counter" bilateral contracts, which are agreed ahead of operating time. If a store is not able to make a profit in this competitive market, then it is unlikely that other sources of revenue (such as contracts with National Grid to assist with real-time balancing of the system) will be sufficient to allow a storage company to survive.

Our method is derived from standard Calculus of Variations techniques and is intended as an extension to [5]. Our work differs from [5] in several respects. Firstly, we present our model in a continuous time setting, rather than the discrete setting employed in [5]. Even if prices are declared in discrete time intervals, in accordance with our current market system, this approach allows for the input of piecewise constant prices but with continuous variation in the operation of the store. Secondly, [5] allow 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 a new motor on if additional power is required. Finally, 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.

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. Finally, in Section 4, we prove existence of an optimal strategy for a particular (and practically useful) case.

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 optimise the actions of the store, for some $T > 0$. We denote by $U \subset \mathbb{R}$ the set of admissible power outputs associated with the store. The operator of the store may then choose an *operating strategy* $q : [0, T] \to U$ which allocates, at each time t, the amount of power $q(t)$ that is to enter the store (using the convention that if $q(t) < 0$, then the store discharges power $-q(t)$ at time t). 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

$$\frac{d}{dt}\ell[q](t) = -\alpha\ell[q](t) + q(t). \tag{1.1}$$

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

$$\ell[q](t) = e^{-\alpha t}\ell_0 + e^{-\alpha t}\int_0^t e^{\alpha s}q(s)ds. \tag{1.2}$$

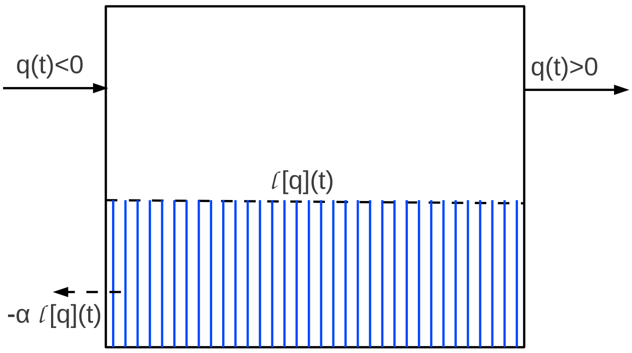

We will sometimes refer to a power output function $q : [0, T] \to U$ as a "strategy," and will only consider those strategies which are piecewise continuous functions of time. This is a physically reasonable assumption, since any strategy which has an accumulation point of discontinuities or a non-jump 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 $u : I \to \mathbb{R}$, for any subset $I \subset \mathbb{R}$.

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

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

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.

The total cost incurred by the store, given an operating strategy q, is denoted $C[q] \in \mathbb{R}$. Here, C is a functional mapping $\mathcal{C}_{pw}([0, T])$ into $\mathbb{R}$, and can incorporate running costs, warming-up costs, costs of storing, etc, 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 minimise 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 and optimal strategies). *Let $\ell_0, \ell_T \in [0, M]$.*

1. *We say that $q : [0, T] \to U$ is an $(\ell_0, \ell_T)-$admissible strategy if q is piecewise continuous and satisfies that*

$$\ell[q](0) - \ell_0 = \ell[q](T) - \ell_T = 0$$

and

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

We denote by $X(\ell_0, \ell_T)$ the set of all $(\ell_0, \ell_T)-$admissible strategies.

2. *We say that q is an $(\ell_0, \ell_T)-$optimal strategy if*

$$C[q] \le C[\tilde{q}] \qquad \forall \tilde{q} \in X(\ell_0, \ell_T). \tag{1.3}$$

Observe that, if no end level ℓ_T is specified, then the optimal strategy will always leave the store empty at time T. Hence, an optimal strategy in this case is in fact $(\ell_0, 0)-$optimal.

2 Main result: characterization of optimal strategies

Throughout this section, we fix initial and terminal conditions $\ell_0, \ell_T \in [0,T]$ and use the short-hand $X = X(\ell_0, \ell_T)$. The following proposition characterizes $(\ell_0, \ell_T)-$admissible strategies in terms of a "reference price" function μ.

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

(i) The strategy q is a minimiser of

$$C[\tilde{q}] - \int_0^T e^{\alpha t}\mu(t)\tilde{q}(t)dt \tag{2.1}$$

over all piecewise continuous functions $\tilde{q} : [0,T] \to U$.

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

$$\dot{\mu}(t) = 0 \quad \text{if } E^-(t) < \ell[q](t) < E^+(t), \tag{2.2}$$
$$\dot{\mu}(t) \ge 0 \quad \text{if } \ell[q](t) = E^+(t), \tag{2.3}$$
$$\dot{\mu}(t) \le 0 \quad \text{if } \ell[q](t) = E^-(t). \tag{2.4}$$

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

$$\mu(t^+) - \mu(t^-) = 0 \quad \text{if } E^-(t) < \ell[q](t) < E^+(t), \tag{2.5}$$
$$\mu(t^+) - \mu(t^-) \ge 0 \quad \text{if } \ell[q](t) = E^+(t), \tag{2.6}$$
$$\mu(t^+) - \mu(t^-) \le 0 \quad \text{if } \ell[q](t) = E^-(t), \tag{2.7}$$

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 $\tilde{q} \in X$ and let $0 = a_1 < \ldots < a_n < a_{n+1} = T$ be a partition such that $\cup_{i=1}^n [a_i, a_{i+1}] = [0,T]$ and $q, \tilde{q}$ are continuous and μ is differentiable over each (a_i, a_{i+1}). Notice that (1.1) can be equivalently written as

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

Hence, (2.1) implies

$$\begin{aligned}
C[q]-C[\tilde q] &\le \int_0^T e^{\alpha t}\mu(t)\Big(q(t)-\tilde q(t)\Big)dt\\
&\stackrel{(1.1)}{=} \sum_{i=1}^{n}\int_{a_i}^{a_{i+1}} \mu(t)\left\{\frac{d}{dt}\Big(e^{\alpha t}\ell[q](t)-e^{\alpha t}\ell[\tilde q](t)\Big)\right\}dt\\
&= \sum_{i=1}^{n}\Big[e^{\alpha a_{i+1}}\mu(a_{i+1}^-)\left(\ell[q](a_{i+1})-\ell[\tilde q](a_{i+1})\right)-e^{\alpha a_i}\mu(a_i^+)\left(\ell[q](a_i)-\ell[\tilde q](a_i)\right)\Big]\\
&\qquad\qquad -\sum_{i=1}^{n}\int_{a_i}^{a_{i+1}} e^{\alpha t}\dot\mu(t)\left(\ell[q](t)-\ell[\tilde q](t)\right)dt\\
&= \sum_{i=2}^{n} e^{\alpha a_i}\left(\mu(a_i^-)-\mu(a_i^+)\right)\left(\ell[q](a_i)-\ell[\tilde q]a_i)\right)-\sum_{i=1}^{n}\int_{a_i}^{b_i} e^{\alpha(t)}\dot\mu(t)\left(\ell[q](t)-\ell[\tilde q](t)\right)dt
\end{aligned}$$

where 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 $\ell[\tilde q](0)=\ell[q](0)$ and $\ell[\tilde q](T)=\ell[q](T)$ if $\tilde q\in X$. Applying the complementary slackness conditions (2.2)-(2.7) to the final equality above, we obtain $C[q]-C[\tilde q]\le 0$. □

It is worth mentioning here that exactly the same result follows if we apply the Karush-Kuhn-Tucker (KKT) conditions to the original minimisation problem (1.3). By following such an approach, one instead looks for minimisers of the relaxed functional

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

over all piecewise-continuous $\tilde q:[0,T]\to U$, where $\lambda_1,\lambda_2:[0,T]\to\mathbb{R}$ are piecewise-continuous functions which satisfy the complementary slackness conditions

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

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:

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

and, in particular,

$$\dot\mu(t)=-\lambda_1(s)+\lambda_2(s).$$

3 An algorithm to determine the optimal strategies of Proposition 2.1

In this section, we present an algorithm to determine μ, and consequently the optimal strategy q, for the class of functionals C which take the form

$$C[u]=\int_0^T L(t,u(t))dt \tag{3.1}$$

for some $L:[0,T]\times U\to\mathbb{R}$. 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 $\lambda\in\mathbb{R}$.

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

$$q=\arg\min_{u\in\mathcal{C}_{pw}([0,T])}\int_0^T\Big(L(t,u(t))-\mu(t)u(t)\Big)dt \tag{3.2}$$

Our algorithm is a generalisation of that provided in [5]. We remark that in the general case, we prove nothing about existence or uniqueness of the optimal solution. We state only that, if the algorithm does not terminate early, then it does indeed provide an optimal strategy q. In Section 4, we apply the algorithm to a more concrete setting, under which we can prove the existence of an optimal solution and the completion of the algorithm.

3.1 Preliminary results and definitions

By Proposition 2.1, we aim to find a piecewise differentiable $\mu : [0,T] \to \mathbb{R}$ such that the optimal strategy $q \in X$ solves (3.2). Moreover, μ should be constant over intervals of time where the level of the store is away from 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

$$u(t,\lambda) = \arg\min_{w\in\mathbb{R}} \Big(L(t,w) - e^{\alpha t}\lambda w \Big). \tag{3.3}$$

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 minimisers associated with λ. With this in mind, we denote by $\mathcal{M}_t \subset \mathbb{R}$ the set of λ which admit multiple minimisers in (3.3). Note that, for the majority of the analysis in [5], L is assumed to be strictly convex and strictly increasing in its second argument. In that setting, minimisers of (2.1) 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 a solution of (3.3). 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 $\lambda \to u(t,\lambda)$ 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,

$$\begin{aligned} L(t,u_1) - e^{\alpha t}\lambda_2 u_1 &= L(t,u_1) - e^{\alpha t}\lambda_1 u_1 - e^{\alpha t}(\lambda_2 - \lambda_1)u_1 \\ &\le L(t,u_2) - e^{\alpha t}\lambda_1 u_2 - e^{\alpha t}(\lambda_2 - \lambda_1)u_1 \\ &= L(t,u_2) - e^{\alpha t}\lambda_2 u_2 - e^{\alpha t}(\lambda_2 - \lambda_1)(u_1 - u_2) \\ &< L(t,u_2) - e^{\alpha t}\lambda_2 u_2. \end{aligned}$$

The first inequality follows from the definition of u_1 and u_2 as solutions to (3.3), 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 $t \in [\tau,T]$.

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 minimisers $u(t,\lambda)$ lie away from discontinuities of L, then they must satisfy that

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

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 $x(\lambda), y(\lambda) \in U\backslash\partial U$ to (3.3), 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 $\lim_{n\to\infty} \lambda_n = \lambda_0$, and set $x(\lambda_0) = x_0$ and $y(\lambda_0) = y_0$. Assume wlog 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 wrt 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

$$\begin{aligned} g'(\lambda_0) &= L_u(t, x_0)x'(\lambda_0) - x'(\lambda_0)\lambda_0 - x_0 \\ &= (L_u(t, x_0) - \lambda_0)\, x'(\lambda_0) - x_0 \\ &= -x_0 \end{aligned}$$

and similarly, $h'(\lambda_0) = -y_0$. Hence, if x_0 and y_0 are both global minimisers solving (3.3) 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 minimisers 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 (2.2) motivates the following definition.

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

$$x(t) = \arg\min_{w \in \mathbb{R}} \Big(L(t, w) - e^{\alpha t} \lambda w \Big). \tag{3.4}$$

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 (3.3) at each $t \in [0, T]$.

Note that, in general, a $(\lambda, \tau)-$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 $(\lambda, \tau)-$admissible strategies in order to construct the optimal strategy $q \in X$. The construction will require the choice of λ to be updated at a discrete subset of times in $[0, T]$. In the special case that one can find a $(\lambda, 0)-$admissible strategy which is also admissible, then of course we are done, since such a strategy satisfies condition (2.2) at every $t \in [0, T]$.

Definition 3.3. *Let $\tau \in [0, T]$, and let $x : [0, T] \to U$ be $(\lambda, \tau)-$admissible, for some $\lambda \in \mathbb{R}$. For any $m \in [0, M]$, we categorise x into sets $X(\tau, m)$ and $X'(\tau, m)$ as follows:*

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

$$m + \ell[x](t) = E^-(t),$$

and there exists $\varepsilon_0 > 0$ such that

$$m + \ell[x](t) + \varepsilon < E^-(t + \varepsilon) \qquad \forall \varepsilon \in (0, \varepsilon_0)$$

and

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

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.

(ii) Similarly, we write $x \in X'(\tau, m)$ if there exists $t \in (\tau, T)$ such that

$$m + \ell[x](t) = E^+(t),$$

and there exists $\varepsilon_0 > 0$ such that

$$m + \ell[x](t) + \varepsilon > E^+(t + \varepsilon) \qquad \forall \varepsilon \in (0, \varepsilon_0)$$

and

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

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.

(iii) If $x \in X(\tau, m) \cup X'(\tau, m)$, then we denote by $s[x](\tau, m) \in (\tau, T]$ the first time at which a capacity constraint is broken.

Definition 3.4. *For any $\tau \in [0,T]$ and $m \in [0,M]$, we define*

$$\Lambda(\tau,m) := \{\lambda \in \mathbb{R} : x \in X(\tau,m) \text{ for some } (\lambda,\tau)-\text{admissible } x\}$$

and

$$\Lambda'(\tau,m) := \{\lambda \in \mathbb{R} : x \in X'(\tau,m) \text{ for some } (\lambda,\tau)-\text{admissible } x\}.$$

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

$$\sup \Lambda(\tau,m) \leq \inf \Lambda'(\tau,m) \qquad \forall(\tau,m) \in [0,T] \times [0,M]. \tag{3.5}$$

Thus, the interiors of the two sets are disjoint:

$$\text{int}(\Lambda(\tau,m)) \cap \text{int}(\Lambda'(\tau,m)) = \emptyset \qquad \forall(\tau,m) \in [0,T] \times [0,M]. \tag{3.6}$$

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 wlog 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 $\dot{\mu}(t) = 0$ 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 wlog that $\tau = \tau_k$, for some $k \in \{1,\ldots,N\}$ (otherwise, again, move τ backwards until this is true).

Step 1: If there exists $\lambda \in \mathbb{R}$ and a $(\lambda,\tau_k)-$admissible x such that $x \notin X(\tau_k,m) \cup X'(\tau_k,m)$, then 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 $(\tau_k,\sigma_k] \subset (\tau_k,T]$ and let $s[x] := s[x](\tau_k,m)$.

a) $\lambda \in \Lambda(\tau,m) \setminus \Lambda'(\tau,m)$: Choose a $(\lambda,\tau_k)-$admissible x such that $m + \ell[x](\sigma_k) = E^+(\sigma_k)$ for some $\sigma_k \in [\tau_k, s[x])$, 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. Proceed to step 3.

b) $\lambda \in \Lambda'(\tau,m) \setminus \Lambda(\tau,m)$: Choose a $(\lambda,\tau_k)-$admissible x such that $m + \ell[x](\sigma_k) = E^-(\sigma_k)$ for some $\sigma_k \in [\tau_k, s[x])$, 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. Proceed to step 3.

c) $\lambda \in \Lambda(\tau,m) \cap \Lambda'(\tau,m)$: Choose a $(\lambda,\tau_k)-$admissible x such that either $m + \ell[x](\sigma_k) = E^+(\sigma_k)$ or $m + \ell[x](\sigma_k) = E^-(\sigma_k)$ for some $\sigma_k \in [\tau_k, s[x])$, 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. Proceed to step 3.

Step 3: Let τ_{k+1} be the first time in $[\sigma_k,T]$ such that, on replacing τ_k with τ_{k+1}, the algorithm would not terminate at step 2. 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 $\dot{\mu} = 0$. 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

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

and choose μ so that (i) $q(t)$ is $(\mu(t),t)-$admissible and (ii) $\mu(\sigma_i^+) \geq \mu(\sigma_i^-)$ if $\ell[q](\sigma_i) = E^+(\sigma_i)$, or $\mu(\sigma_i^-) \leq \mu(\sigma_i^-)$ if $\ell[q](\sigma_i) = E^-(\sigma_i)$.

The algorithm is complete if (on varying the choice of μ in Step 4 if needed) the following conditions hold at each $i \in \{1,\ldots,N\}$:

- If $\ell[q](\sigma_i) = E^+(\sigma_i)$, then μ is non-decreasing over $(\sigma_i,\tau_{i+1}]$ and $\mu(\tau_{i+1}^-) \leq \mu(\tau_{i+1}^+)$.

- If $\ell[q](\sigma_i) = E^-(\sigma_i)$, then μ is non-increasing over $(\sigma_i, \tau_{i+1}]$ and $\mu(\tau_{i+1}^-) \geq \mu(\tau_{i+1}^+)$.

If these conditions do not hold, then the algorithm terminates here and yields no solution q.

End of algorithm

Remark 3.5. *The above algorithm implicitly reduces the original optimisation problem to a series of new optimisation problems which are localised in time. At each time τ_k, one only needs to look ahead to time σ_k to know how to operate the store over (τ_k, σ_k). Hence, whilst we originally assumed knowledge of the entire price function $p : [0,T] \to [0,\infty)$, in practise 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 3.7 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 3.6. *If the above algorithm yields a strategy $q : [0,T] \to U$, then q is an optimal strategy.*

Proof. Suppose that the pair (μ, q) of Proposition 2.1 exists and is known up to time $\tau_k \in [0,T)$. At each $i \in \{1, \dots, N\}$, the algorithm yields

$$\sup \Lambda(\tau_i, \ell[q](\tau_i)) \leq \mu(\tau_i^+) \leq \inf \Lambda'(\tau_i, \ell[q](\tau_i)). \tag{3.7}$$

If the conditions of step 1 are satisfied at τ, 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 (2.2)-(2.7) are satisfied by μ.

To this end, let σ_k be as defined in the algorithm. By construction, we have ensured that the conditions of the proposition 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

$$\mu(\sigma_k -) \in \Lambda(\sigma_k, E^+(\sigma_k)). \tag{3.8}$$

Thus, if $\sigma_k = \tau_{k+1}$ then, together with (3.7), this implies

$$\mu(\sigma_k^-) \leq \sup \Lambda(\sigma_k, E^+(\sigma_k)) \leq \mu(\sigma_k^+) \tag{3.9}$$

and condition (2.6) 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 $\mu(\sigma_k^-) \leq \mu(\sigma_k^+)$. It is always possible to find a μ which satisfies conditions (i) and (ii) of step 3 due to the monotonicity property of Lemma 3.1.

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

$$\mu(\sigma_k) \in \Lambda'(\sigma_k, E^-(\sigma_k)). \tag{3.10}$$

Thus, if $\sigma_k = \tau_{k+1}$ then, together with (3.7), this implies

$$\mu(\sigma_k) \geq \inf \Lambda'(\sigma_k, E^-(\sigma_k)) \geq \mu(\sigma_k^+) \tag{3.11}$$

and condition (2.7) 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 $\mu(\sigma_k^-) \geq \mu(\sigma_k^+)$. As above, it is always possible to find a μ which satisfies conditions (i) and (ii) of step 3 due to the monotonicity property of Lemma 3.1.

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

Proposition 3.7. *If the algorithm terminates early, then there is no pair (μ, q) satisfying the conditions of Proposition 2.1.*

Proof. Suppose that there is a pair $(\overline{\mu}, \overline{q})$ which satisfies the conditions of Proposition 2.1, and let (μ, q) be a pair generated by the algorithm. 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_i, σ_i). We write $m_i := \ell[q](\tau_i)$ and

$$\lambda_i := \mu(t) = \sup \Lambda(\tau_i, m_i) \qquad \forall t \in (\tau_i, \sigma_i).$$

We may assume that there is no $k \in \{1, \dots, N\}$ such that $\sigma_k = \tau_{k+1}$ and $\lambda_k = \lambda_{k+1}$, since if this were the case, then the intervals (τ_k, σ_k) and $(\tau_{k+1}, \sigma_{k+1})$ could be replaced by the single interval (τ_k, σ_{k+1}), and we could remove the times σ_k and τ_{k+1} from the sequences.

Analogously, let $\{\overline{\tau}_i\}_{i=1}^{\overline{N}}$ and $\{\overline{\sigma}_i\}_{i=1}^{\overline{N}}$ be increasing sequences, which define the end-times of all maximal intervals (τ_i, σ_i) over which $\overline{\mu}$ is constant (where by "maximal" we mean that if $(\tau, \sigma) \supset (\tau_k, \sigma_k)$ is a strictly larger interval, then μ is not constant over this interval). Again, we may write $\overline{m}_i := \ell[\overline{q}](\overline{\tau}_i)$ and and

$$\overline{\lambda}_i := \overline{\mu}(t) \qquad \forall t \in (\overline{\tau}_i, \overline{\sigma}_i).$$

Throughout this proof, we will use the notation $\lambda \sim_k \lambda'$ if the set of $(\lambda, \tau_k)-$admissible functions coincides exactly with the set of $(\lambda', \tau_k)-$admissible functions over the interval (τ_k, σ_k). (In fact, Lemma 3.1 implies that it suffices to find only one function x which is both $(\lambda, \tau_k)-$admissible and $(\lambda', \tau_k)-$admissible, in order to conclude that $\lambda \sim_k \lambda'$.) We introduce the ordering $\lambda \prec_k \lambda'$ if $\lambda \nsim_k \lambda'$ and $\lambda < \lambda'$.

The proof consists of proving the following two statements:

1. If $k \in \{1, \dots, \overline{N}\}$, then $\overline{\lambda}_k \sim_k \sup \Lambda(\overline{\tau}_k, \overline{m}_k)$ and $\overline{q}$ is $(\lambda_k, \tau_k)-$admissible over the interval (τ_k, σ_k).
2. For each $k \in \{1, \dots, N\}$, we have $(\overline{\tau}_k, \overline{\sigma}_k) = (\tau_k, \sigma_k)$.

Together, these statements will complete the proof: The second statement implies that the algorithm uncovers all intervals $(\tau_k, \sigma_k]$ over which $\overline{\mu}$ is constant, so that the algorithm does not terminate at steps 1 to 3. The first statement implies that we may set $\mu = \overline{\mu}$ over each interval of the form (σ_k, τ_{k+1}), so that the remaining conditions of step 4 are satisfied. To see this, suppose that $\overline{\mu}$ is increasing over (σ_k, τ_{k+1}). Then, statement 2 implies that $\overline{q}(\sigma_k^-) = q(\sigma_k^-)$. Thus, if $t \in (\sigma_k, \tau_{k+1})$ is such that $\overline{\mu}(t) \geq \overline{\mu}(\sigma_k)$, then Lemma 3.1 implies that $\overline{q}(t) \geq \overline{q}(\sigma_k) = q(\sigma_k)$. Another application of Lemma 3.1 thus implies that $\overline{\mu}(t) \geq \mu(\sigma_k^-)$. Thus, by taking $\mu = \overline{\mu}$, we ensure that μ is increasing over (σ_k, τ_{k+1}), with $\mu(\sigma_k^+) - \mu(\sigma_k^-) \geq 0$. Similar arguments apply at τ_{k+1} and for the case where $\overline{\mu}$ is decreasing over (σ_k, τ_{k+1}).

Proof of statement 1: Let $k \in \{1, \dots, \overline{N}\}$. It is clear from the definition of $(\overline{\lambda}_k, \tau_k)-$admissibility and from the requirement that $\overline{q}$ must be a minimiser of (2.1), that $\overline{q}$ must be $(\overline{\lambda}_k, \tau_k)-$admissible over the interval $(\overline{\tau}_k, \overline{\sigma}_k)$. We need to show that

$$\overline{\lambda}_k \sim_k \lambda := \sup \Lambda(\overline{\tau}_k, \overline{m}_k).$$

Suppose first that $\overline{\lambda}_k \prec_k \lambda$. Then, $\overline{\lambda}_k \in \Lambda(\overline{\tau}_k, \overline{m}_k) \setminus \Lambda'(\overline{\tau}_k, \overline{m}_k)$. If $m + \ell[\overline{q}](\overline{\sigma}_k) = E^+(\overline{\sigma}_k)$, then recall that the functions $u_t(\cdot)$ of Lemma 3.1 are piecewise continuous and monotone increasing. The functions are constructed in such a way that one can always find a suitable $u_t(\overline{\lambda}_k)$ to coincide with $\overline{q}$ over $(\overline{\tau}_k, \overline{\sigma}_k)$. The monotonicity property of u_t implies that $\lambda' \in \Lambda'(\overline{\tau}_k, \overline{m}_k)$ for all $\lambda' \succ_k \lambda$. Together with the piecewise continuity property of u_t implies that $\overline{\lambda}_k \sim_k \lambda$, which contradicts the assumption that $\overline{\lambda}_k \prec_k \lambda$. If, on the other hand, $m + \ell[\overline{q}](\overline{\sigma}_k) = E^-(\overline{\sigma}_k)$, then Proposition 2.1 requires that $\overline{\mu}(\overline{\sigma}_k^+) - \overline{\mu}(\overline{\sigma}_k^-) \leq 0$. If $\overline{\mu}$ is constant over an interval $(\sigma, \tau') \subset (\sigma, T]$, then the monotonicity property of Lemma 3.1 implies that the lower capacity constraint will be broken by $\overline{q}$, and no admissible solution will be found. If $\overline{\mu}$ is not constant over such an interval, then we must have $\dot{\overline{\mu}}(t) \leq 0$ and $u_t(\overline{\mu}(t)) = (d/dt)E^-(t)$ for all $t \in (\sigma, \tau')$. However, since u_t in monotone increasing, this is only possible if $\dot{\overline{\mu}}(t) > 0$ at some $t \in (\sigma, \tau')$, which contradicts Proposition 2.1. Hence, we must have either $\overline{\lambda}_k \sim_k \lambda$ or $\overline{\lambda}_k \succ_k \lambda$.

Suppose now that $\overline{\lambda}_k \succ_k \lambda$. Then, $\overline{\lambda}_k \in \Lambda'(\overline{\tau}_k, \overline{m}_k) \setminus \Lambda(\overline{\tau}_k, \overline{m}_k)$. If $m + \ell[\overline{q}](\overline{\sigma}_k) = E^-(\overline{\sigma}_k)$, then a similar argument as above implies that $\lambda' \in \Lambda'(\tau, m)$ for all $\lambda' \succ_k \lambda$, so that $\overline{\lambda}_k \sim_k \lambda$, which contradicts the assumption that $\overline{\lambda}_k \succ_k \lambda$. If, on the other hand, $m + \ell[\overline{q}](\overline{\sigma}_k) = E^+(\overline{\sigma}_k)$, then Proposition 2.1 requires that $\overline{\mu}(\overline{\sigma}_k^+) - \overline{\mu}(\overline{\sigma}_k^-) \geq 0$. If $\overline{\mu}$ is constant over an interval $(\sigma, \tau') \subset (\sigma, T]$, then the monotonicity property of Lemma 3.1 implies that the lower capacity constraint will be broken by $\overline{q}$, and no admissible solution will be found. If $\overline{\mu}$ is not constant over such an interval, then we must have $\dot{\overline{\mu}}(t) \geq 0$ and $u_t(\overline{\mu}(t)) = (d/dt)E^+(t)$ for all $t \in (\sigma, \tau')$. However, since u_t in monotone increasing, this is only possible if $\dot{\overline{\mu}}(t) < 0$ at some $t \in (\sigma, \tau')$, which contradicts Proposition 2.1. We conclude that $\overline{\lambda}_k \sim_k \lambda$, and the proof of statement 1 is complete.

Proof of statement 2: We will show first that the algorithm is well-defined. This reduces to proving the following: suppose that the algorithm uniquely selects the first $k-1$ times of the sequences $\{\tau_i\}_{i=1}^N$ and $\{\sigma_i\}_{i=1}^N$. Then, by construction, τ_k is uniquely defined by the algorithm, and we need to show that the selection of σ_k is also unique. In other words, we must prove that the selection of σ_k is independent of the choice of x made at step 2 of the algorithm, at time τ_k.

To this end, suppose that there are two $(\lambda_k, \tau_k)-$admissible functions x^1, x^2 which satisfy the conditions of step 2. Let $m = \ell[q](\tau_k)$ be the level of the store determined by the algorithm at time τ_k. For $i \in \{1,2\}$, let N^i correspond to the N of the algorithm, if $x = x^i$ is selected at time τ_k, and let $\{\tau_j^i\}_{j=k}^{N^i}$ and $\{\sigma_j^i\}_{j=k}^{N^i}$ be the corresponding sequences of times. Then,

$$m^i := m + \ell[x^i](\sigma_k^i) = \begin{cases} E^+(\sigma_k^i) & \text{if } x^i \in X(\tau_k, m) \\ E^-(\sigma_k^i) & \text{if } x^i \in X'(\tau_k, m) \end{cases}$$

Assume wlog that $\sigma_k^1 \leq \sigma_k^2$ and let t^* be the first intersection time of the two strategies over the interval $[\sigma_k^1, T]$, so that $\ell[x^1](t^*) = \ell[x^2](t^*)$. Given $i, j \in \{1,2\}$, we can construct a new $(\lambda_k, \tau_k)-$admissible strategy $x_{i,j}^*$ as follows:

$$x_{i,j}^* := \begin{cases} x^i(s) & \text{if } s \in [0, t^*) \\ x^j(s) & \text{if } s \in [t^*, T]. \end{cases}$$

If $\sigma_k^1 \leq t^* \leq \sigma_k^2$, then $x_{1,2}^*(t) = x^2(t)$ for all $t \in [\sigma_k^2, T]$. In particular, if we choose $x = x^1$ at step 2, then the algorithm yields $(\tau_{k+1}^1, \sigma_{k+1}^1) = (\sigma_k^1, \sigma_k^2)$. Hence, since $x_{1,2}^*$ is $(\lambda_k, \tau_k)-$admissible, we deduce that

$$\sup \Lambda(\tau_k, m) = \sup \Lambda(\tau_{k+1}^1, m^1) = \sup \Lambda(\tau_k^2, m)$$

and the two choices x^1 and x^2 yield the same interval $(\tau_k, \sigma_k^2) = (\tau_k, \sigma_k^1] \cup [\tau_{k+1}^1, \sigma_{k+1}^1)$ over which μ is constant.

If $\sigma_k^1 \leq \sigma_k^2 < t^*$, then if we can find another $(\lambda_k, \tau_k)-$admissible function x^3 whose first intersection with x^1 or x^2 over the interval $(\tau_k, T]$ occurs at a time $s^* \in [\sigma_k^1, \sigma_k^3]$ (where σ_k^3 is defined analogously to σ_k^1 and σ_k^2), then a similar argument as above proves that μ defined over $[\tau_k, \sigma_k^3]$ is independent of the choice of x. If there is no such x^3, then we have $x_{1,2}^* \in X(\tau_k, m)$ and $x_{2,1}^* \in X'(\tau_k, m)$ (switching the labels $1,2$ if necessary). Neither $x_{1,2}^*$ or $x_{2,1}^*$ satisfies the conditions of step 1 or step 2 of the algorithm and, consequently, there is no solution and the algorithm would terminate early. This completes the proof that the algorithm is well-defined and that the sequences of times $\{\tau_k\}_{k=1}^N$ and $\{\sigma_k\}_{k=1}^N$ are independent of the choices made at step 2.

Suppose now that $\tau \in [0, T)$ and that μ coincides exactly with $\overline{\mu}$ (up to $\sim_k$ equivalences) up until time τ. The proof is complete on showing that $\tau = \tau_k$ if and only if $\tau = \overline{\tau}_k$, and that if $\tau = \tau_k = \overline{\tau}_k$, then $\sigma_k = \overline{\sigma}_k$.

Suppose that $\tau \neq \tau_k$ but $\tau = \overline{\tau}_k$. Then the conditions of steps 1 and 2 do not hold at time τ. Statement 1 states that $\overline{\lambda}_k \sim_k \sup \Lambda(\overline{\tau}_k, \overline{m}_k)$. If $\overline{q}$ coincides with some x over $(\overline{\tau}_k, \overline{\sigma}_k)$ such that $x \in X(\overline{\tau}_k, \overline{m}_k)$, then Proposition 2.1 implies that either $\overline{\sigma}_k = T$ or $\ell[\overline{q}](\overline{\sigma}_k) = E^-(\overline{\sigma}_k)$ (since the conditions of step 2 do not hold). If $\overline{\sigma}_k = T$, then clearly $\tau = \tau_k$, since we may set $\mu(t) = \overline{\mu}(t)$ for all $t \in [\tau, T]$, and the conditions of step 1 are satisfied. If $\overline{\sigma}_k < T$, then $\ell[\overline{q}](\overline{\sigma}_k) = E^-(\overline{\sigma}_k)$ and Proposition 2.1 implies that either (i) $\overline{\sigma}_k = \overline{\tau}_{k+1}$ and $\overline{\mu}(\overline{\sigma}_k^+) - \overline{\mu}(\overline{\sigma}_k^-) \leq 0$, or (ii) $\overline{\sigma}_k < \overline{\tau}_{k+1}$ and $\overline{\mu}(t) - \overline{\mu}(\overline{\sigma}_k^-) \leq 0$, for some $t \in [\overline{\sigma}_k, \overline{\tau}_{k+1}]$. In case (i), the monotonicity property of Lemma **??** results in breaking the lower capacity constraint. In case (ii), since x crosses the E^-, we must have that $\overline{q}(s) < (d/dt)E^-(s)$ for all s in some interval $I \subset [\overline{\sigma}_k, \overline{\tau}_{k+1}]$. However, since we require that $\overline{q}(s) \equiv (d/dt)E^-(s)$ for all $s \in (\overline{\sigma}_k, \overline{\tau}_{k+1})$, this means that we must have $\dot{\overline{\mu}}(s) > 0$ for some $s \in I$. This contradicts the conditions of Proposition 2.1, and we conclude that $\tau = \tau_k$. The same result follows from similar arguments if $\overline{q}$ coincides with some x over $(\overline{\tau}_k, \overline{\sigma}_k)$ such that $x \in X'(\overline{\tau}_k, \overline{m}_k)$.

Finally, suppose that $\tau = \tau_k$ but $\tau \neq \overline{\tau}_k$. Then $\tau \in (\sigma_i, \tau_{i+1})$ for some $i \in \{1, \ldots, \overline{N}-1\}$ and $\ell[\overline{q}](\tau) = m \in \{E^-(\tau), E^+(\tau)\}$. Suppose that $m = E^+(\tau)$ so that, by Proposition 2.1, we have $\overline{\mu}$ is increasing over (σ_i, τ_{i+1}). Since $\tau = \tau_k$, we can find a $(\lambda_k, \tau)-$admissible x which satisfies the conditions of step 2. In particular, the conditions of step 2 imply that there are times $\tau' < \tau''$ such that $\overline{q}(t) \leq (d/dt)E^+(t)$ for all $t \in I_1$ and $\overline{q}(t) \geq (d/dt)E^-(t)$ for all $t \in I_2$, with strict inequality somewhere in each interval. By the monotonicity property of Lemma 3.1, this implies that either (i) $\dot{\mu}(t) < 0$ for some $t \in I_2$, or (ii) $\dot{\mu}(t) = 0$ for all $t \in I_1 \cup I_2$. If case (i) holds, then this contradicts condition (2.3) of Proposition 2.1, whilst if case (ii) holds, then this contradicts the assumption that $\tau \neq \overline{\tau}_k$. Similar arguments result in similar contradictions if $m = E^-(\tau)$ and we conclude that $\tau = \overline{\tau}_k$. □

4 Existence of an optimal strategy for a typical storage model example participating in the wholesale market

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

$$U = [-q^-_{\max}, -q^-_{\min}] \cup \{0\} \cup [q^+_{\min}, q^+_{\max}], \tag{4.1}$$

with $-q^-_{\max} < -q^-_{\min} < 0 < q^+_{\min} < q^+_{\max}$. Thus, the store is equipped with maximum and minimum power outputs, on both the charge and discharge side. We denote by $M > 0$ the maximum capacity of the store and assume that the minimum amount of stored energy which the store can hold is 0 (it is not difficult to adapt our results for a more general lower capacity bound $m \in (0, M)$, but we use 0 here for ease of notation).

We look for $(\ell_0, \ell_T)-$optimal strategies, for any choice of $\ell_0, \ell_T \in [0, M]$. Thus, the admissible energy domain E is the unique closed domain in $\mathbb{R}^2$ bounded by the region $[0, T] \times [0, M]$ and the lines

$$\begin{aligned} y &= q^+_{\max} t + \ell_0, & y &= -q^+_{\max}(T - t) + \ell_T, \\ y &= -q^-_{\max} t + \ell_0 & y &= q^-_{\max}(T - t) + \ell_T. \end{aligned}$$

Notice that in the special case that $q_{\max}- = q^+_{\max} = \infty$ (i.e. there are no maximum power constraints), then $E = [0, T] \times [0, M]$.

The cost functional C for this storage model takes the special form (3.1), with

$$L(t, q(t)) := w(q(t))p(t)q(t) + k(t) \qquad \forall q \in X(\ell_0, \ell_T). \tag{4.2}$$

Here, $w : U \to (0, \infty)$ 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. Here, we assume that w takes on two values $w_1 \geq 1$ and $w_2 \in (0, 1]$, so that

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

The function $k : [0, T] \to [0, \infty)$ is piecewise continuous and represents the basic cost of running the store (including staff costs, rent payments and so on). In fact, it is clear in this case, that we would obtain the same optimal strategy on replacing (w_1, w_2, p) with $(1, w_2/w_1, w_1 p)$, thus reducing the number of parameters by one.

With a cost functional of this form, a simple result is immediately available, as stated in the following lemma: if the price of electricity is constant over all time, then it is not worth investing in storage.

Lemma 4.1. *If $\dot{p} \equiv 0$ and if $\ell_T \geq \ell_0$, then there is no non-zero strategy $q \in X$ which yields a positive profit.*

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

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

□

We now prove that, for this particular example, the algorithm always yields an optimal strategy. We begin by assuming that $w_1 > 1$ and $0 < w_2 < 1$, and will later prove existence for the case where $w_1 \geq 1$ and $0 < w_2 \leq 1$ by finding a sequence of strategies which converges to the optimal one.

By Proposition 2.1, if the algorithm yields a differentiable $\mu : [0, T] \to \mathbb{R}$, then the optimal strategy $q \in X$ minimises

$$\int_0^T \Big(w(\tilde{q}(t))p(t) - e^{\alpha t}\mu(t)\Big)\tilde{q}(t)dt + \int_0^T k(t)dt$$

over all piecewise continuous $\tilde{q} : [0, T] \to U$. Hence, assuming we know μ, we obtain that the optimal strategy is of the form:

$$q(t) = \begin{cases} q^+_{\max} & \text{if } p(t) < e^{\alpha t}\mu(t)/w_1 \\ -q^-_{\max} & \text{if } p(t) > e^{\alpha t}\mu(t)/w_2 \\ 0 & \text{if } e^{\alpha t}\mu(t)/w_1 < p(t) < e^{\alpha t}\mu(t)/w_2. \end{cases}$$

Suppose that we have run the algorithm up until a time $\tau \geq 0$ and assume, without loss of generality, that either $\tau = 0$ or τ corresponds to one of the times σ_k determined by the algorithm. Then, the store is either full or empty at time τ, and we denote by $m \in \{E^-(\tau), E^+(\tau)\}$ the level of the store at time τ.

If there exists $\lambda \in \mathbb{R}$ which satisfies the conditions of step 1 of the algorithm, then we are done. If there is no such λ, then we need to apply step 2. We will show that, under the conditions that $w_1, w_2 \neq 1$, we have that $\tau = \sigma_k = \tau_{k+1}$. In particular, $[0, T] = \cup_{i=1}^N [\tau_i, \sigma_k]$ and μ is piecewise constant. The algorithm will therefore always yield the solution $\mu(t) = \sup \Lambda(\tau_i, \ell[q](\tau_i))$ for each $i \in (\tau_i, \sigma_i)$.

To this end, let $\lambda := \sup \Lambda(\sigma_k, m)$. In Lemma 4.2, we will prove that an appropriate $(\lambda_k, \sigma_k)-$admissible x can always be chosen to satisfy the conditions of step 2 of the algorithm. Before stating the lemma, we introduce some notation. Define the candidate strategy x_λ by $x_\lambda(t) = 0$ if $t \in [0, \sigma_k]$ and

$$x_\lambda(t) = \begin{cases} q^+_{\max} & \text{if } p(t) < e^{\alpha t}\lambda/w_1 \\ -q^-_{\max} & \text{if } p(t) < e^{\alpha t}\lambda/w_2 \\ 0 & \text{if } e^{\alpha t}\lambda/w_1 < p(t) < e^{\alpha t}\lambda/w_2 \end{cases} \tag{4.3}$$

if $t \in (\sigma_k, T]$.

In the degenerate case where $p(t)$ coincides with $p(t) = e^{\alpha t}\lambda/w_1$ or $p(t) = e^{\alpha t}\lambda/w_2$ for all t in some interval in $[\sigma_k, T]$, we have a choice of values to assign to the x at time t, and some extra care needs to be taken to ensure that the correct choice is made. For this reason, we will actually consider all x of the form

$$x_\lambda(t) + f_\lambda(t) \qquad \forall t \in I,$$

for some function f_λ such that

$$\begin{aligned} &f_\lambda(t) \in \{0, q^+_{\max}\} && \text{if } p(t) = e^{\alpha t}\lambda/w_1 \text{ and } t \in (\sigma_k, T] \\ &f_\lambda(t) \in \{-q^-_{\max}, 0\} && \text{if } p(t) = e^{\alpha t}\lambda/w_2 \text{ and } t \in (\sigma_k, T] \\ &f_\lambda(t) = 0 && \text{otherwise.} \end{aligned} \tag{4.4}$$

The purpose of f_λ is simply to adjust x_λ, if needed, so that it complies with the conditions of the algorithm i.e. either the upper or lower capacity constraint is reached but not crossed. We identify two particular choices:

$$f^+_\lambda(t) = \begin{cases} q^+_{\max} & \text{if } p(t) = e^{\alpha t}\lambda/w_1 \text{ and } t \in (\sigma_k, T] \\ 0 & \text{otherwise} \end{cases}$$

and

$$f^-_\lambda(t) = \begin{cases} -q^-_{\max} & \text{if } p(t) = e^{\alpha t}\lambda/w_2 \text{ and } t \in (\sigma_k, T] \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, $\lambda \in \Lambda(\sigma_k, m)$ if and only if $x_\lambda + f^-_\lambda \in X(\sigma_k, m)$; and $\lambda \in \Lambda'(\sigma_k, m)$ if and only if $x_\lambda + f^+_\lambda \in X'(\sigma_k, m)$.

Lemma 4.2. *Using the notation of the algorithm, let $\lambda = \sup \Lambda(\sigma_k, m)$. If $\lambda \in \Lambda(\sigma_k, m)$, then either $x_\lambda + f^+_\lambda$ is admissible, or there exists a function f_λ and a time $\sigma' \in (\sigma_k, T]$ such that f_λ satisfies conditions (4.4), and the strategy $x_\lambda + f_\lambda$ is admissible over some interval I such that $(\sigma_k, \sigma'] \subsetneq I \subset (\sigma_k, T]$ and $\ell(\sigma', x_\lambda + f_\lambda) = M$.*

Similarly, if $\lambda \in \Lambda'(\sigma_k, m)$, then either $x_\lambda + f^-_\lambda$ is admissible, or there exists a function f_λ and a time $\sigma' \in (\sigma_k, T]$ such that f_λ satisfies conditions (4.4), and the strategy $x_\lambda + f_\lambda$ is admissible over some interval I such that $(\sigma_k, \sigma'_k] \subsetneq I \subset (\sigma_k, T]$ and $\ell(\sigma', x_\lambda + f_\lambda) = 0$.

In particular, in either case we have $\sigma_k = \tau_{k+1}$ and $\sigma' = \sigma_{k+1}$.

Proof. Note first the identity

$$x_{\lambda_1}(t) + f_{\lambda_1}(t) \le x_{\lambda_2}(t) + f_{\lambda_2}(t)$$

for all $\lambda_1, \lambda_2 \in \mathbb{R}$ with $\lambda_1 < \lambda_2$, all f_{λ_1} and f_{λ_2} satisfying (4.4), and all $t \in (\sigma_k, T]$. This follows immediately from Lemma 3.1 and from the conditions of (4.4).

Assume first that $\lambda \in \Lambda(\sigma_k, m)$ and let $s = s[x + f_\lambda^-](\sigma_k, m)$ be the first time at which the lower capacity constraint is broken by the strategy $x_\lambda + f_\lambda^-$. For any $\delta > 0$ and any $\lambda' \in \mathbb{R}$, we construct $f_{\lambda'}^\delta$ as follows: Let $S_1 \subset [\sigma_k, s]$ be the set of times t such that $p(t) = e^{\alpha t}\lambda/w_1$, and $S_2 \subset [\sigma_k, s]$ the set of times t such that $p(t) = e^{\alpha t}\lambda/w_2$. Choose any subsets $J_1 \subset S_1$ and $J_2 \subset S_2$ such that $|J_1| = \delta_1/q_{\max}^+$ and $|J_2| = \delta_2/q_{\max}^-$, with $\delta_1 + \delta_2 = \delta$, and set

$$f_{\lambda'}(t) = \begin{cases} q_{\max}^+ & \text{if } t \in J_1 \\ 0 & \text{if } t \in J_2 \\ f_\lambda^- & \text{otherwise.} \end{cases}$$

This construction is always possible provided $q_{\max}^+|S_1| + q_{\max}^-|S_2| \ge \delta$, and we have the identity

$$\int_{\sigma_k}^t f_{\lambda'}^\delta(u)du = \int_{\sigma_k}^t f_{\lambda'}^-(u)du + \delta \le \int_{\sigma_k}^t f_{\lambda'}^+(u)du \qquad \forall t \in (\sigma_k, T]. \tag{4.5}$$

The right-hand inequality becomes equality when $\delta = q_{\max}^+|S_1| + q_{\max}^-|S_2|$.

If $\lambda \in \Lambda(\sigma_k, m)$ and $\lambda \in \Lambda'(\sigma_k, m)$, then the strategy $x_\lambda + f_\lambda^+$ breaks the upper capacity constraint at a time $s' \in (\sigma_k, T]$ but is admissible over the interval (σ_k, s'). Hence, the relation given by (4.5) implies that we can choose a suitable $\delta \in (0, q_{\max}^+|S_1| + q_{\max}^-|S_2|)$, and suitable J_1 and J_2, to ensure that $\ell[x_\lambda + f_\lambda^\delta](\sigma') = M$ at some $\sigma' \in (\sigma_k, s')$. The claim therefore holds true if $\lambda \in \Lambda(\sigma_k, m)$ and $\lambda \in \Lambda'(\sigma_k, m)$.

If, on the other hand, $\lambda \in \Lambda(\sigma_k, m)$ and $\lambda \notin \Lambda'(\sigma_k, m)$, then suppose that the claim is not true so that there exists a time $s' \in (\sigma_k, T]$ such that $x_\lambda + f_\lambda^+$ breaks the lower capacity constraint at time s', but $\ell[x_\lambda + f_\lambda^+](t) < M$ for all $t \in (\sigma_k, s')$. Let $\varepsilon > 0$ be small enough so that there are no intervals of time $J \subset (\sigma_k, T]$ such that $p(t) = e^{\alpha t}(\lambda + \varepsilon)/w_1$ or $p(t) = e^{\alpha t}(\lambda + \varepsilon)/w_2$ for all $t \in J$. This is always possible due to the piecewise continuity of p. The piecewise continuity of p also implies the existence of $\delta(\varepsilon) > 0$ which satisfies that $\delta(\varepsilon) \to 0$ as $\varepsilon \to 0$ and

$$\begin{aligned} \# \left\{ t \in (\sigma_k, s') : e^{\alpha t}\lambda/w_1 < p(t) \ < e^{\alpha t}(\lambda + \varepsilon)/w_1 \right\} \\ + \# \left\{ t \in (\sigma_k, s') : e^{\alpha t}\lambda/w_1 < p(t) < e^{\alpha t}(\lambda + \varepsilon)/w_1 \right\} \le \delta(\varepsilon). \end{aligned} \tag{4.6}$$

Hence, we obtain

$$0 \le \ell[x_{\lambda+\varepsilon} + f_{\lambda+\varepsilon}^-](t) - \ell[x_\lambda + f_\lambda^+](t) \le \delta(\varepsilon)(q_{\max}^- + q_{\max}^+) \qquad \forall t \in (\sigma_k, T]. \tag{4.7}$$

If the store is empty at time σ_k so that $m = 0$, then inequality (4.7) implies that

$$\ell[x_{\lambda+\varepsilon} + f_{\lambda+\varepsilon}^-](t) < M \qquad \forall t \in (\sigma_k, s]$$

and, in particular, $\lambda + \varepsilon \in \Lambda(\sigma_k, m)$, which contradicts the definition $\lambda = \sup \Lambda(\sigma_k, m)$. If, on the other hand, $m = M$, then we need to appeal to the assumption that $w_1 > 1$ and $0 < w_2 < 1$. Since $\ell[x_\lambda + f_\lambda^+](t) < M$ for all $t \in (\sigma_k, s)$, there exists $\kappa_0 > 0$ such that $p(t) \ge e^{\alpha t}\lambda/w_2$ for all $t \in (\sigma_k, \sigma_k + \kappa_0)$. Thus, there exists $\kappa \in (0, \kappa_0)$ such that $p(t) > e^{\alpha t}(\lambda + \varepsilon)/w_1$ for all $t \in (\sigma_k, \sigma_k + \kappa)$, provided $\varepsilon > 0$ is sufficiently small. This implies that $x_{\lambda+\varepsilon}(t) + f_{\lambda+\varepsilon}(t) \le 0$ for all $t \in (\sigma_k, \sigma_k + \kappa)$ and, in particular, $0 < \ell[x_{\lambda+\varepsilon} + f_{\lambda+\varepsilon}](t) \le M$ for all $t \in (\sigma_k, \sigma_k + \kappa)$. Together with inequality (4.7), this implies that $\lambda + \varepsilon \in \Lambda(\sigma_k, m)$ which, again, contradicts the definition $\lambda = \ sup\Lambda(\sigma_k, m)$.

This completes the proof of the claim for the case $\lambda \in \Lambda(\sigma_k, m)$. If $\lambda \in \Lambda'(\sigma_k, m)$, then the proof of the claim follows similarly.

□

The above lemma proves the assertion that the algorithm always yields an optimal strategy if $w_1 > 1$ and $0 < w_2 < 1$. If $w_1 = 1$ or $w_2 = 1$ then, for any small $\varepsilon > 0$, consider optimal strategy q_ε associated with the parameters $w_1 + \varepsilon$ and $w_2 - \varepsilon$ (whose existence is guaranteed by Lemma 4.2). Then, provided p, T and U are all bounded, there exists a constant $K > 0$ such that

$$\int_0^T w(q_\varepsilon(t))p(t)q_\varepsilon(t)dt \le \int_0^T p(t)\tilde{q}(t)dt + K\varepsilon$$

for all $\tilde{q} \in X(\ell_0, \ell_T)$. If U is bounded, then there exists a convergent subsequence ε_n, and a piecewise continuous $q_0 : [0,T] \to U$, such that $\varepsilon_n \to 0$ and $q_{\varepsilon_n} \to q_0$ as $n \to \infty$. In particular, taking limits of the above inequality, we obtain

$$\int_0^T w(q_0(t))p(t)q_0(t)dt \leq \int_0^T p(t)\tilde{q}(t)dt.$$

But the conditions on w also imply the relation

$$\int_0^T w(q_0(t))p(t)q_0(t)dt \geq \int_0^T p(t)q_0(t)dt.$$

Hence, we conclude that q_0 is an optimal strategy.

4.1 Illustrative results

The following graphs plot the optimal strategies for a store with $q^+_{\max} = q^-_{\max}$ and $q^+_{\max}/M = 7$ (the exact values of these parameters do not affect the solution, we need only know their relation to each other). We take $w_1 = 1$ and vary w_2 as labelled below. In each figure, the "blocky" blue lines plot the level of the store $\ell[q](t)$ at each time t, assuming that the optimal strategy q is followed. The green lines plot the price of electricity $p : [0,T] \to [0,\infty)$ which the store faces in making its operating decision. Here, we have chosen hourly prices from the first 250 hours N2EX's day-ahead auction in November 2013.

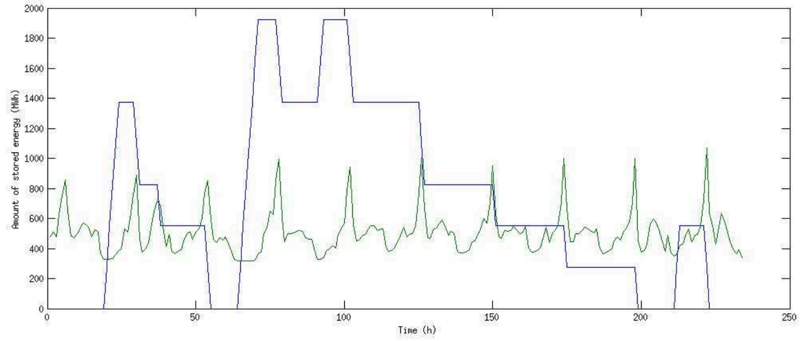

Figure 4.1: $w_2 = 0.5$

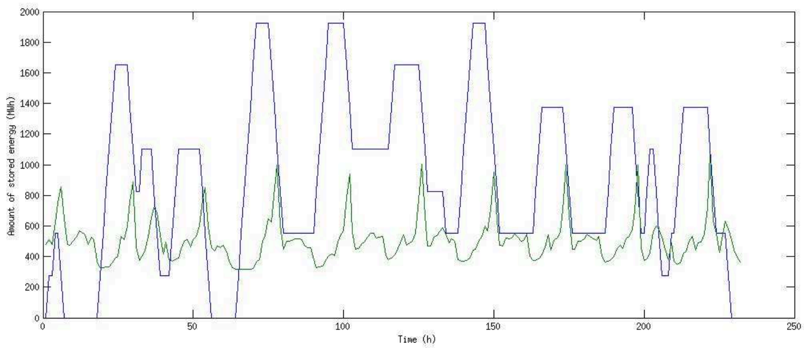

Figure 4.2: $w_2 = 0.7$

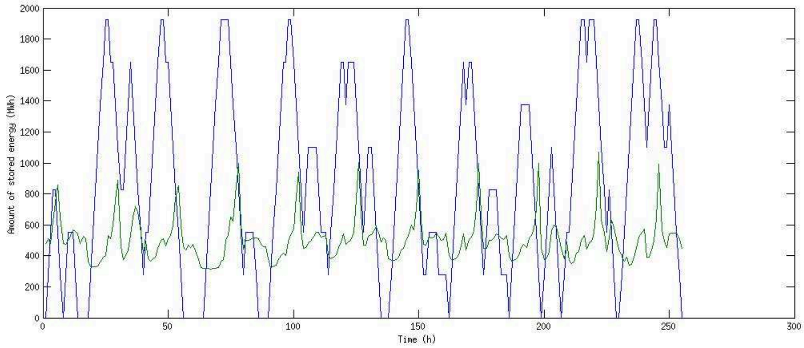

Figure 4.3: $w_2 = 0.9$

An immediate observation is that the store cycles more frequently as its efficiency increases. As an implication of this, we conjecture that profits are an increasing nonlinear non-negative function of efficiency. There are two intuitive reasons for this. Firstly, less energy is lost during operation at higher efficiencies. Secondly, as the efficiency increases, then the system cycles more often and is earning revenue over a larger proportion of elapsed time. However, this proportion of time will not increase linearly with efficiency: if the store charges to full capacity and then immediately discharges, for example, then increasing the efficiency by a small amount will not change this behaviour.

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. 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 minimisation 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 $I_1 \cup \ldots \cup I_N = [0, T]$, 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, as mentioned in Remark 3.5, the algorithm has the property that it implicitly reduces our problem to a series of new optimisation problems, which are localised in time. More precisely, if the store is either empty or full at a time $\tau \in [0, T]$, then there is a time $\sigma \in (\tau, T]$ such that, as long as we know the restriction of p over (σ, τ), then we can determine the optimal strategy over this interval.

5 Conclusions

We have presented a method for determining the operating strategy for a store to maximise its arbitrage profits. Our setting allows for leakage, inefficiencies and general operating costs which, in particular, are not required to be convex functions of power output. As illustrated in the example of Section 4, a significant benefit associated with this method is the implicit localisation in time of the solution.

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.

Finally, we should comment that in reality, it is likely that a store will need to secure revenue through a variety of markets if it is to be a profitable investment. In particular, it should not be ignored that a store could provide valuable balancing services to the system operator. In such a setting, we would need to introduce some element of uncertainty into our problem and, as such, new methods would probably be required.

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.

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

References

[1] *Planning our electric future: a White Paper for secure, affordable and low-carbon electricity,* Department of Energy and Climate Change (July 2011)

[2] *UK Future Energy Scenarios,* National Grid (September 2012)

[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] J. Cruise, L. Flatley, R. Gibbens, S. Zachary: Optimal control of storage for arbitrage, with applications to energy systems arXiv:1307.0800v1 (2013)

[6] *Energy Storage and Management Study,* AEA (2010)

[7] *Options for low-carbon power sector flexibility to 2050 - a report to the Committee on Climate Change,* Pöyry (October 2010)

[8] M. Black, G. Strbac: Value of Bulk Energy Storage for Managing Wind Power Fluctuations. *IEEE Transactions on Energy Conversion,* 22(1), p197–205 (2007)

[9] 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* (June 2012)

[10] J.P. Barton, D.G. Infield: Energy Storage and Its Use With Intermittent Renewable Energy. *IEEE Transactions on Energy Conversion,* 19(2), p441–448 (2004)

[11] 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, p4807-4815 (2011)

[12] *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) (August 2010)

[13] N. Löhndorf, S. Minner: Optimal day-ahead trading and storage of renewable energies - an approximate dynamic programming approach. *Energy Syst.,* 1, p61–77 (2010)

[14] *Smarter Network Storage, Six Monthly Report, December 2013,* UK Power Networks (February 2014)