European Journal of Operational Research 319 (2024) 611–624

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Optimal trading with regime switching: Numerical and analytic techniques applied to valuing storage in an electricity balancing market

Paul Johnson [a,*], Dávid Zoltán Szabó [b], Peter Duck [a]

[a] The University of Manchester, Oxford Rd, M13 9PL, Manchester, United Kingdom
[b] Budapesti Corvinus Egyetem, Fővám tér 8., Budapest, 1093, Hungary

ARTICLE INFO

Keywords:
Perpetual options
Regime switching
Stochastic optimal control
Partial differential equations

ABSTRACT

Accurately valuing storage in the electricity market recognizes its role in enhancing grid flexibility, integrating renewable energy, managing peak loads, providing ancillary services and improving market efficiency. In this paper we outline an optimal trading problem for an Energy Storage Device trading on the electricity balancing (or regulating) market. To capture the features of the balancing (or regulating) market price we combine stochastic differential equations with Markov regime switching to create a novel model, and outline how this can be calibrated to real market data available from NordPool. By modelling a battery that can be filled or emptied instantaneously, this simplifying assumption allows us to generate numerical and quasi analytic solutions.

We implement a case study to investigate the behaviour of the optimal strategy, how it is affected by price and underlying model parameters. Using numerical (finite-difference) techniques to solve the dynamic programming problem, we can estimate the value of operating an Energy Storage Device in the market given fixed costs to charge or discharge. Finally we use properties of the numerical solution to propose a simple quasi-analytic approximation to the problem. We find that analytic techniques can be used to give a benchmark value for the storage price when price variations during the day are relatively small.

1. Introduction

In order to tackle climate change, energy markets around the globe are being incentivized to transfer to clean energy. The share of energy produced by renewable sources in the EU is expected to rise to more than 60% by 2030 (European Commission, 2023). This transition entails new challenges as renewable sources are more unpredictable in their generation compared to fossil-based sources. That is, renewable energy sources, such as biomass, geothermal resources, sunlight, wind and water are prone to be seasonable and/or intermittent. Therefore, due to the ongoing shift in energy markets, in order to supply energy in potential shortage periods, energy storage facilities are increasingly being utilized to address the issue of unpredictability and variability. In order to supply energy in potential shortage periods, it is of utmost importance to deploy energy storage at scale.

Besides high seasonal variability, the variability of price within a day can also be significant, and this high intra-day variability can threaten the transmission network security, necessitating the more widespread usage of ancilliary reserves as discussed by Chattopadhyay (2014). Modelling this variability and the opportunities that may arise is a popular topic for academics, and one way to consider this is through regime switching. A stochastic process is regime switching if its dynamics is determined by different regimes under different periods. In fact, intra-day and regime-switching dynamics have been considered in several papers addressing various energy markets including Bierbrauer et al. (2004), Karakatsani and Bunn (2008), Mount et al. (2006), Tiwari and Menegaki (2019), Weron et al. (2004). Recently in the Danish renewable electricity sector Shah et al. (2021) described a Markov switching model has been employed by distinguishing a low and a high volatility regime.

Supported by the aforementioned articles and on the empirical analysis we conduct in this paper, we will also incorporate regime switching into our energy price modelling. In this paper we seek the fair value of optimally operating a storage facility, and we claim that the inclusion of a regime switching feature bolsters the framework. Modelling with a single regime has the potential to underestimate the potentially huge but infrequent profits available when things go wrong on the market and spikes associated with high volatility regimes occur, and at the same time overestimate the profits available under normal circumstances when the underlying volatility is lower. The cost of storage facilities inevitably has a significant influence on the strategy of

* Corresponding author.
E-mail addresses: paul.johnson-2@manchester.ac.uk (P. Johnson), davidzoltan.szabo@uni-corvinus.hu (D.Z. Szabó), Peter.Duck@manchester.ac.uk (P. Duck).

https://doi.org/10.1016/j.ejor.2024.06.026
Received 21 September 2023; Accepted 18 June 2024
Available online 20 June 2024

balancing market participants. The recent study Rahman et al. (2020) provides an overview of different energy storage technologies. Due to technological advances it is reasonable to assume that the most influential technical parameters affecting (adding to) the cost and environmental performances such as lifetime, round-trip efficiency, and cycle length are expected to decrease in the future.

In the present paper we will focus on a battery storage operator who makes transactions on the Nordic power market, where a number of closely linked markets are maintained (Nord Pool AS, 2022). In particular, there is a standalone *NordPool Day-Ahead market*, *NordPool Intraday market* and *NordPool Regulating market*.

The *NordPool Day-Ahead market* is the so called spot market (Elspot), in which the price of energy is determined by supply and demand. On this spot market, bidding closes at noon for deliveries from midnight and 24 h ahead. The *NordPool Regulating market* operates as a real-time market and its main function is to counteract imbalances related to planned day-ahead operation. As quite a long time period remains between close of bidding in the day-ahead market, and the regulating market, the *NordPool Intraday market* is introduced as an 'in between market', where participants in the day-ahead market can trade bilaterally. On this intraday market, the product traded is the one-hour long power contract. The intraday market thus works together with the day-ahead market to help secure the necessary balance between supply and demand, as one trades closer to the physical delivery within the intraday markets.

In this study, we will focus on trading on the regulating market. First, this is supposed to be the ultimate resort to counteract imbalances, second the mathematical framework, described in Section 2, would not be applicable to the day-ahead or intraday markets due to a sizeable delay between bids and deliveries. As outlined, the regulating market is maintained to level out imbalances in the physical trade on the spot and intraday markets in order to ensure a balance between production and consumption, and to provide power grid stability. Totalling the deviations from bid volumes at the spot and intraday markets gives a net imbalance for that hour in the system that needs to be settled. If the grid is congested, the market breaks up into area markets, and equilibrium must be established in each area. The regulating market is therefore the ultimate resource for correcting such imbalances, and it provides the necessary physical trade and accounting in the liberalized Nordic electricity system.

Stylized facts of the electricity market include mean reversion, seasonality, extreme volatility and spikes as discussed by Gudkov and Ignatieva (2021). We discussed above that several papers have modelled energy markets with switching processes that can successfully incorporate the effects of seasonality, extreme volatility and spikes. Moreover, mean-reverting phenomena on energy prices is also observed. Periodic shifts in demands push prices up, which entails the entering of more expensive generators on the supply side that pushes back the price to its former level. Studies such as Weron (2007) and Escribano et al. (2011) included this additional stylized fact, moreover Janczura and Weron (2012) considered the parameter estimation of mean-reverting processes combined with Markov regime-switching using energy prices.

We have recently seen several papers considering stochastic optimal controls within an energy business context. For a detailed summary we refer the reader to Nadarajah and Secomandi (2023) and also Dai et al. (2021). A real options approach for natural gas storage optimization by controlling the stored amount of gas in the storage facility has been implemented in Thompson et al. (2009) and a solution has been achieved via the Hamilton–Jacobi–Bellman equations. The optimal timing of reserve electricity reserve facilities has been considered in Moriarty and Palczewski (2017) who obtained analytical results assuming the dynamics of the spot energy price is a Wiener process. A paper concerning optimal switching with an energy storage facility in a setting where the spot price dynamics is a mean-reverting Ornstein–Uhlenbeck process, Szabó et al. (2020) obtained numerical results as a solution of the corresponding Hamilton–Jacobi–Bellman equations. Allowing for negative energy prices, Zhou et al. (2016) considered a Markovian decision problem for fast storage facilities by controlling the stored energy in the inventory and achieved analytical and numerical results for the corresponding Bellman equation via backward dynamic programming. Regarding Markovian regime switching processes, Wahab and Lee (2011) provided numerical methods to value swing options and presented numerical examples for gasoline prices. Considering oil markets and allowing for jump diffusions, regime switches and mean reversions, Harikae et al. (2021) valued projects via a real options approach. They approximated the underlying stochastic process depicting the project value with a generalized implied binomial tree. In our understanding, there is a lack of articles considering stochastic control problems on energy markets with Markovian regime switching underlying processes.

Therefore, we contribute to the literature by examining optimal switching problems corresponding to the operation of energy storage facilities on a market where the energy price goes through regime switches and in both regimes they exhibit mean-reverting phenomena. In our model, we derive the optimal control for the timing of charges/discharges, where the energy storage facility is assumed to be always full or empty. We are able to show that the derived optimal control strategies depend primarily on the dynamics of the current regime, with only small adjustments made to account for regime switching. This means we can solve each regime independently with analytic approximations and then recombine them to provide an estimate of the value. When applying the models to real data through calibration and testing, we find that our regime switching model is a good fit to the data, and so the optimal strategies that we produce will perform better and provide more realistic estimates for the value of an investment. Our quasi analytic approximations can be solved efficiently and used to provide the optimal strategies for operation as well as rough estimates of the value, although more detailed time varying models (and numerical solutions) could be required in some cases. This technique in which we solve regimes independently and recombine them, could then be used to solve other optimal control problems with regime switching.

2. The regime switching optimal trading problem

2.1. Regime-switching diffusion processes

In line with the Nord-Pool data study to be detailed in Section 3, we introduce a regime-switching diffusion process, which will be used to model the difference between regulating and day-ahead market prices. The stability and optimal control of regime-switching diffusion processes are of great interest and there are numerous studies in the literature, for example Korn et al. (2017), Mao and Yuan (2006), Shao (2015), Shao and Xi (2013) among others. Regime-switching diffusion processes comprise a number of diffusion processes modulated by a random switching process. To provide a precise definition, we follow the methodology of Shao and Xi (2013). The regime-switching diffusion process is a two-component process $(X(t), \Lambda(t))$, where $X(t)$ describes the dynamics of the diffusion process itself and $\Lambda(t)$ describes the random switching process. We will work on a complete probability space $(\Omega, \mathbb{A}, \mathbb{P})$, where Ω is the sample space, the σ-algebra $\mathbb{A}$ denotes the event space and $\mathbb{P}$ is the corresponding probability measure. This $(\Omega, \mathbb{A}, \mathbb{P})$ carries a standard Wiener process W and N^I and N^{II} Poisson processes. We set

$$\mathcal{F}_t^0 \triangleq \sigma(W_s, N_s^I, N_s^{II} \, : \, s \in [0,t]) \tag{1}$$

and denote by $\mathcal{F}_t$ the completion of $\mathcal{F}_t^0$. According to Korn et al. (2017) we know that $\mathcal{F} = \{\mathcal{F}_t\}_{t\geq 0}$ is right continuous, thus the filtered probability space $(\Omega, \mathbb{A}, \mathcal{F}, \mathbb{P})$ satisfies the usual conditions.

Our model for electricity prices follows the general regime-switching diffusion process framework, where the $X(t)$ component satisfies the following stochastic differential equation (SDE):

$$dX(t) = \sigma(X(t), \Lambda(t))dW(t) + b(X(t), \Lambda(t))dt, \quad X(0) = x_0, \tag{2}$$

and b and σ are Lipschitz continuous functions of linear growth in the first component for all $i \in \mathbb{S}$, and $\mathbb{S}$ denotes the finite state space of the $\Lambda(t)$ second component. In this particular study, the Markov chain will contain two elements– $\mathbb{S} := \{I, II\}$, such that

$$\mathbb{P}\{\Lambda(t+\delta) = l | \Lambda(t) = k\} = \begin{cases} p_{k,l}\delta + o(\delta) & if \quad k \neq l \\ 1 + p_{k,k}\delta + o(\delta) & if \quad k = l \end{cases} \tag{3}$$

provided $\delta \downarrow 0$. The P-matrix is irreducible and conservative meaning that $p_{k,k} = -\sum_{j\neq k} p_{k,j}$, $k \in \mathbb{S}$. Based on Shao and Xi (2014) the existence of a nonexplosive solution $(X(t), \Lambda(t))$ of (2)–(3) is ensured.

In light of the above, the continuous-time Markov chain Λ alternates between regimes I and II according to events in the Poisson processes N^I and N^{II} with intensities $p_{I,II}$ and $p_{II,I}$. That said, if $\Lambda(t) = I$ then $\inf_{s\geq t}\{N^I(s) > N^I(t)\}$ corresponds to the time point when Λ will change to state II, and if $\Lambda(t) = II$ then $\inf_{s\geq t}\{N^{II}(s) > N^{II}(t)\}$ corresponds to the time point when Λ will change to state I. Thus, the X diffusion process is only affected by exogenous forces, meaning that whenever the random environment jumps from one state to the other one, the diffusion process starts to behave according to the dynamics of the new environment.

In particular, we consider a regime-switching Ornstein–Uhlenbeck process, meaning that $X(t)$ satisfies the mean-reverting SDE:

$$dX(t) = \kappa_{\Lambda(t)}(\theta_{\Lambda(t)} - X(t))dt + \sigma_{\Lambda(t)}dW(t), \quad X(0) = x_0 \tag{4}$$

We also assume that both $\Lambda(t)$ and $W(t)$ are observable, so the battery operator will know all the subsets of the $\mathcal{F}_t$ σ-algebra at time t. The goal of the battery operator will be to maximize the expected values of total net profits. Even though they have no control over the X process, they are aware of its dynamics, and so they are bidding to exploit expected future price movements. They will alternate between buying and selling energy on the regulating market, to completely empty or fully charge the storage facility, therefore their adjoining transactions will be of different signs.

2.2. Impulse control problem

In order to mathematically express the optimization problem associated with the actions of the battery operator, we consider an impulse control problem and so introduce additional variables by following the study of Korn et al. (2017). First, we denote by $\hat{E} = \{E, F\}$ whether the current state of the battery is empty or full, which can be controlled by the operator through buy/sell transactions. Since we assume that the battery storage operator is trading on the regulating market, we consider the day ahead market price of electricity $P(t)$ as a deterministic function of t known for the next 24 h trading period. It is the regulating price R that we treat as stochastic, and the time variable $t \in \mathbb{R}^+$ affects the regulating price though the equation

$$R(X, t) = X(t) + P(t). \tag{5}$$

By taking the Cartesian product of $\hat{E}$, $\mathbb{S}$ and $\mathbb{R}^+$, let $\mathcal{E} = \hat{E} \times \mathbb{S} \times \mathbb{R}^+ = \{(E, I, t), (E, II, t), (F, I, t), (F, II, t)\}$ represent the space of possible regimes of the impulse control problem, containing four distinct regimes for a fixed t. $\mathcal{E}$ is a vector-valued stochastic process, and its second component is the Λ Markov-chain process. The vector valued I_t^k process, defined as $I_t^k(t) \triangleq (k, \Lambda(t), t)$, runs uncontrolled for $k = \{E, F\}$. As the state of the variable k does not have an impact on the X difference[1] process, so we can also write that the difference process has the uncontrolled $X^{I_t^k}(t) \triangleq X^{\Lambda(t)}(t)$ dynamics, and it is the unique solution of

$$X(t) = x_0 + \int_0^t \kappa_{I_s}(\theta_{I_s} - X(s))ds + \int_0^t \sigma_{I_s} dW(s) \tag{6}$$

where $X(0) = x_0$.

Within the optimization strategy, the operator controls the timing of the buy/sell transactions and thus has control to decide when to charge and discharge the battery and therefore when to switch between the values of $\mathcal{E}$. Nonetheless, the operator has no control over the $X(t)$ and $\Lambda(t)$ process values, which have uncontrolled switching dynamics.

That said, the state-dependent constraints on interventions are modelled with the help of a set-valued function:

$$\mathcal{A} : \mathcal{E} \to 2^{\mathcal{E}} \quad \text{such that} \quad i \in \mathcal{A}(i) \quad \text{for all} \quad i \in \mathcal{E} \tag{7}$$

where $\mathcal{A}(i)$ represents the set of regimes that are attainable by an intervention in regime $i \in \mathcal{E}$.

Since at each regime only one possible intervention is available, the intervention function only contains one element, and we define it as follows:

$$\mathcal{A}(\hat{E}, \mathbb{S}, t) = \begin{cases} \{(F, I, t)\} & \hat{E} = E \text{ and } \mathbb{S} = I \\ \{(F, II, t)\} & \hat{E} = E \text{ and } \mathbb{S} = II \\ \{(E, I, t)\} & \hat{E} = F \text{ and } \mathbb{S} = I \\ \{(E, II, t)\} & \hat{E} = F \text{ and } \mathbb{S} = II \end{cases} \tag{8}$$

and hence $\mathcal{A}$ here reflects that the battery operator can only control how they charge and discharge the battery.

Next, we discuss the intervention costs of the admissible regime switches. Note that there is no instantaneous running cost or profit, meaning that staying at the same battery state does not generate or incur any value. The intervention costs are determined by the energy price on the regulating market at the time of the transaction, as well as some financial penalty incurred at the time of trading. The operator always pays or is paid the regulating price $R(x,t)$ whenever market trade occurs, and an additional penalty function $C(x,t)$ should be included to take account of any market transaction fees or marginal losses incurred by the battery degrading/maintenance. Thus, for each admissible regime switch, $R(x,t) - C(x,t)$ or $-R(x,t) - C(x,t)$ is the cash-flow for the storage operator. Nonetheless, a buy/sale restriction is also affected by the sign of the difference process X, meaning that $\text{sign}(X(t))$ determines which side of the transaction is viable for traders on the regulating market at time t.[2] We can describe how to incorporate this restriction to the impulse control problem once costs for switches have been introduced.

The cost function is defined as

$$K : \mathcal{E}^2 \times \mathbb{R} \to \mathbb{R}, \tag{9}$$

where the last component represents the value of the price process X modelling the difference between the regulating and spot markets. Note that for a given (i, j, x) realization of $(\mathcal{E}, \mathcal{E}, X(t))$, the $K(i, j, x)$ values are finite in cases where $j \in \mathcal{A}(i)$. For all other $j \neq i$ values, the intervention costs are infinitely large implying that an optimal trader would never choose to switch from i to j. This allows us to neatly capture physical or market constraints such as only being able to place bids when $X(t) < 0$ and asks for $X(t) > 0$. Clearly, staying at a particular regime or in other words not making an intervention is always feasible and there is no associated cost to it.

Having outlined the possible regime switches we proceed by expressing the associated intervention costs. Intervention costs are fully determined by the current state of the battery, the time and the regulating price process, thus we can disregard other components of $\mathcal{E}$.

[1] We can rewrite (5) as $X(t) = R(X,t) - P(t)$, indeed being a difference process. We only directly model X stochastically, thus R the regulating price is the sum of a stochastic and a deterministic component.

[2] The sign of $X(t)$ expresses whether the power gets regulated upwards or downwards. In case $X < 0$, market participants are incentivized to sell energy on the regulating market for a price higher than the spot price; and in case $X > 0$, market participants are incentivized to buy energy on the regulating market for a price lower than the spot price.

Now we can write the cost (negative cash-flow) $K_x^t : \hat{E}^2 \to \mathbb{R}$ functional under x and at time t, as a matrix:

$$K_x^t(\cdot,\cdot) = \begin{pmatrix} \infty & \infty \\ R(x,t) + C(x,t) & \infty \end{pmatrix} \quad \text{for} \quad x \geq 0 \tag{10}$$

and

$$K_x^t(\cdot,\cdot) = \begin{pmatrix} \infty & -R(x,t) + C(x,t) \\ \infty & \infty \end{pmatrix} \quad \text{for} \quad x \leq 0. \tag{11}$$

Note that in the above matrices the order of the elements of set $\hat{E}$ follows $\hat{E} = \{E, F\}$ and the infinite cost ensures that such a strategy will never be chosen. Of course, when solving such a problem numerically we can simply build those constraints into the algorithm by not considering those switches that would result in infinite costs.

Remark. As a simplification, we assume that the battery is always empty or full, thus only two states are considered. As further discussed in Section 4, we consider a battery with a capacity of 1 MWh. Besides this, there is a minimum bid size restriction to trade on the Nordic regulating market. The minimum granularity will be 1 MW in the Nordic market after the European implementation project for the creation of the European mFRR platform is conducted (Nordic Balancing Model, 2021). Thus, due to the minimum size restriction, we rule out trades with a fraction of energy. Note, that more frequent trading with less than full capacity would also entail more frequent penalties and overcharges associated with the $C(x,t)$ penalty function.

Next we must define the intervention operators. For a given $x \in \mathbb{R}$ and for a given bounded continuous function $\phi : \mathcal{E} \times \mathbb{R} \to \mathbb{R}$, the intervention operator $\mathcal{M}$ is defined via

$$\mathcal{M}\phi(i,x) \triangleq \sup_{i' \in \mathcal{A}(i)} \{\phi(i',x) - K(i,i',x)\}. \tag{12}$$

Let the initial regime $i \in \mathcal{E}$, the initial value of the regime-switching process $X(0) = x_0$ and the initial Markov chain value $\Lambda(0)$ be given.

An impulse control strategy $\mathcal{IC}$ consists of a non-decreasing sequence of $(\tau_\ell)_{\ell \geq 1}$ stopping times, and a sequence of actions $(\iota_\ell)_{\ell \geq 1}$ where (ι_ℓ) is an $\hat{E}$ valued $\mathcal{F}_{\tau_\ell}$- measurable random variable for each $\ell \geq 1$. Given a strategy $\mathcal{IC} = \{\tau_\ell, \iota_\ell\}_{\ell \geq 1}$ we can construct the corresponding controlled regime process $I^{\mathcal{IC}}$ iteratively by setting $\tau_0 = 0$, $\iota_0 = i$ and by defining $I^{\mathcal{IC}}$ for $t \in [\tau_\ell, \tau_{\ell+1})$ as:

$$I_t^{\mathcal{IC}} \triangleq I_t^{\iota_\ell}$$

The above construction is feasible provided that the strategy $\mathcal{IC} = \{\tau_\ell, \iota_\ell\}_{\ell \geq 1}$ is admissible, i.e. the following two conditions are satisfied:

1. $0 \leq \tau_\ell \leq \tau_{\ell+1} \to \infty$ a.s. as $\ell \to \infty$
2. $\{\iota_\ell, \Lambda(\tau_\ell), \tau_\ell\} \in \mathcal{A}(I_{\tau_\ell}^{\iota_{\ell-1}})$ for $\ell \geq 1$

In this study since our defined Λ runs uncontrolled, and the operator has no influence over its value, we can define the expanded impulse control strategy as $\mathcal{IC}^{EXP} \triangleq \{\iota_\ell, \Lambda(\tau_\ell), \tau_\ell\}_{\ell \geq 1}$. In this way we can rewrite condition 2. as

2′. $\mathcal{IC}_{\tau_\ell}^{EXP} \in \mathcal{A}(I_{\tau_\ell}^{\iota_{\ell-1}})$ for $\ell \geq 1$

We are now ready to define the impulse control problem as:

$$V(i,x_0) \triangleq \sup_{\mathcal{IC}^{EXP}} \mathcal{V}(\mathcal{IC}^{EXP}; i, x_0) \tag{13}$$

where

$$\mathcal{V}(\mathcal{IC}^{EXP}; i, x_0) \triangleq \mathbb{E}\Big[-\sum_{\ell=1}^{\infty} e^{-r\tau_\ell} K_{X(\tau_\ell)}(I_{\tau_\ell}^{\iota_{\ell-1}}, \mathcal{IC}_{\tau_\ell}^{EXP})\Big] \tag{14}$$

given initial regime $i = \{\iota_0, \Lambda(0), 0\}$ and $X(0) = x_0$, which enables us to invoke the findings of Korn et al. (2017).

Theorem (*Theorem 3.1 of Korn et al., 2017*). *Assume that the value function $V : \mathcal{E} \times \mathbb{R} \to \mathbb{R}$ satisfies:*

1. *V is non-negative, bounded and continuous*
2. *$V \geq \mathcal{M}V$ on $\mathbb{R}$*
3. *V satisfies the Dynamic Programming Principle*

Then V is the value function of the impulse control problem of Eq. (13). The optimal control strategy $\hat{\mathcal{IC}}$ can be constructed as follows:

- *$\hat{\tau}_0 \triangleq 0$ and $\hat{\iota}_0 \triangleq \iota_0$*
- *Given the optimal controlled regime process $I^{\hat{\mathcal{IC}}}$, and the process X on $[0, \hat{\tau}_k)$ and an intervention $\hat{\iota}_k$ at $\hat{\tau}_k$ we set*

$$\hat{\tau_{k+1}} \triangleq \inf\{t \geq \hat{\tau}_k : V(I_t^{\hat{\tau}_k, \hat{\iota}_k}, X(t)) = \mathcal{M}V(I_t^{\hat{\tau}_k, \hat{\iota}_k}, X(t))\} \tag{15}$$

Then the optimal controlled regime process $I^{\hat{\mathcal{IC}}}$ for $[\hat{\tau}_k, \hat{\tau}_{k+1})$ is defined as $I_t^{\hat{\tau}_k, \hat{\iota}_k}$.

Remark. Even though $\lim_{x \to \infty} R(x,t) = \infty$, we will confirm with the help of the numerical results that $V^i(x)$ is bounded for all $x \in \mathbb{R}$ and $i \in \mathcal{E}$. Global boundedness of V in the proof of Theorem 3.1 of Korn et al. (2017) is solely used to claim that the maximizer of $\mathcal{M}V$ is always measurable. This property trivially holds for the problem of this paper as only a maximum of one regime switch is attainable for each $i \in \mathcal{E}$.

We now turn to discussing the dynamic programming principle. This is satisfied for all bounded functions $V \in C^2$ with bounded derivatives that satisfy the quasi-variational inequalities:

$$\max\left(\mathbb{L}^i V^i - rV^i - \sum_{j \in \mathcal{E}} p_{i,j}(V^j - V^i)\ ,\ \mathcal{M}V^i - V^i \right) = 0 \quad i \in \mathcal{E} \tag{16}$$

Here $p_{i,j} \neq 0$ only in case $i = (\cdot, I, \cdot)$ and $j = (\cdot, II, \cdot)$; or in case $j = (\cdot, I, \cdot)$ and $i = (\cdot, II, \cdot)$, meaning that sum reduces to just one element. We can trivially identify $p_{i,j}$ values, when the three dimensional $i, j \in \mathcal{E}$ values are given by only considering the second entries. In general we do not expect to have C^2 functions as strong solutions of impulse control problems. Therefore we assume that there exists a unique viscosity solution of the quasi-variational inequalities, and we seek the viscosity solution to this problem (see Chapter 9 of Øksendal & Sulem, 2005, for a discussion of viscosity solutions).

2.3. Numerical scheme

We now discuss the numerical scheme used to solve this dynamic programming principle. The infinitesimal operator of X is determined by the state of the Markov chain, thus for an $i = (\cdot, S, \cdot)$ we can equivalently write $\mathbb{L}^i = \mathbb{L}^S$. As for a fixed t time point, there are four distinct regimes (E, I, t), (E, II, t), (F, I, t), and (F, II, t) leading to four quasi-variational inequalities, one for each regime. Given a point in the domain (x, S, k, t), the function $V_S^k(x,t)$ denotes the value of the problem for a battery in state k, S Markov chain state, x and t time point. Then we can write a general form for the variational inequality as:

$$\max\left(\mathbb{L}^S V_S^k - rV_S^k - p_{S,S'}(V_{S'}^k - V_S^k), \mathcal{M}V_S^k - V_S^k\right) = 0, \tag{17}$$

where the infinitesimal operators of the Ornstein–Uhlenbeck processes for state S is given as

$$\mathbb{L}^S V_S^k = \frac{\partial}{\partial t} V_S^k + \frac{\sigma_S^2}{2} \frac{\partial^2}{\partial x^2} V_S^k - \kappa_S x \frac{\partial}{\partial x} V_S^k \tag{18}$$

to generate the four equation problem to solve when $k = E, F$ and $S = I, II$. Note that $S' = II$ if $S = I$ and $S' = I$ if $S = II$.

To obtain a representative value for trading over a single day, we solve the above problem with a terminal condition

$$V_S^k(x, t = T; T) = 0$$

for large T and repeat the daily cycle of the Elspot price according to the relation

$$P(t + 24) = P(t).$$

In the limit as $T \to \infty$, we should expect to find that a time-periodic solution emerges in all states such that

$$\lim_{T\to\infty}\left[V_S^k(x,t+24;T) - V_S^k(x,t;T)\right] = 0.$$

To discretize the problem, we choose a sufficiently large domain $(-x_{\max}, x_{\max})$, and divide into $2iMax+1$ equally spaced grid points so

$$x_i = i\Delta x \text{ where } -iMax \le i \le iMax \text{ and } \Delta x = \frac{x_{\max}}{iMax}.$$

For the time variable, we discretize the repeating period $[0,24]$ so

$$t_j = j\Delta t \text{ where } 0 \le j \le jMax \text{ and } \Delta t = \frac{24}{jMax}.$$

By writing the value function at a particular grid point as

$$v_{S,i,j}^k = V_S^k(x_i, t_j)$$

the Ornstein–Uhlenbeck operator with a Crank–Nicolson style approximation at $(x_i, t_j + 1/2\Delta t)$ becomes

$$\begin{aligned}\mathbb{L}^S v_{S,i,j}^k =& \frac{v_{S,i,j+1}^k - v_{S,i,j}^k}{\Delta t} \\ &+ \frac{\sigma_S^2}{2}\frac{v_{S,i-1,j+1}^k - 2v_{S,i,j+1}^k + v_{S,i+1,j+1}^k + v_{S,i-1,j}^k - 2v_{S,i,j}^k + v_{S,i+1,j}^k}{2(\Delta x)^2} \\ &- \kappa_S x_i \frac{v_{S,i+1,j+1}^k - v_{S,i-1,j+1}^k + v_{S,i+1,j}^k - v_{S,i-1,j}^k}{4\Delta x}\end{aligned} \tag{19}$$

for $k = E, F$ and $S = I, II$. Now we use a combined explicit/implicit procedure to capture regime/optimal switching respectively, so that the final scheme may be written

$$\max\left(\mathbb{L}^S v_{S,i,j}^k - r\frac{1}{2}(v_{S,i,j}^k + v_{S,i,j+1}^k) - p_{S,S'}(v_{S',i,j+1}^k - v_{S,i,j+1}^k), \mathcal{M}v_{S,i,j}^k - v_{S,i,j}^k\right) = 0. \tag{20}$$

For large positive and large negative x, we solve

$$\max\left(\frac{v_{S,i,j+1}^k - v_{S,i,j}^k}{\Delta t} - \kappa_S x_i \frac{v_{S,i+1,j+1}^k - v_{S,i-1,j+1}^k + v_{S,i+1,j}^k - v_{S,i-1,j}^k}{4\Delta x} - r\frac{1}{2}(v_{S,i,j}^k + v_{S,i,j+1}^k) - p_{S,S'}(v_{S',i,j+1}^k - V_{S,i,j+1}^k), \mathcal{M}V_{S,i,j}^k - V_{S,i,j}^k\right) = 0. \tag{21}$$

This boundary condition is well known to work well on these types of problems (see Szabó et al., 2020, for example). Essentially it assumes that

$$\left|\frac{\sigma_S^2}{2}\frac{\partial^2 V}{\partial x^2}\right| \ll \frac{\partial V}{\partial t} - \kappa_S x \frac{\partial V}{\partial x} - rV$$

holds for either large and positive x or large and negative x. In this case the second order term can be ignored leading to the reduced PDE (21) and what we have in essence is a Robin-type condition implemented at the boundary.

Numerous grid checks, convergence analysis and verification against analytic solutions in special cases have been carried out to ensure the validity and integrity of the numerical results.

3. Calibrating from NORD pool data

We now detail our methods concerning the parameter estimation for NordPool markets. Provided by Nord Pool AS (2022) we have access to historical data for the following Nordic areas: NO1, NO2, NO3, NO4, NO5, SE1, SE2, SE3, SE4, FI, DK1, DK2; where NO abbreviates Norway, SE abbreviates Sweden, FI abbreviates Finland and DK abbreviates Denmark. We downloaded hourly spot prices and regulating prices for the time period 2018–2021.

The day-ahead market is the less volatile market, and it contains peak (09:00 to 20:00) and off peak (21:00 to 08:00) periods, indicating the presence of daily seasonality. Nord-Pool day-ahead prices have been previously studied (Kristiansen, 2012; Weron & Misiorek, 2008), where price forecast models have been proposed applying daily and weekly seasonality effects. Fig. 1 illustrates Elspot and Regulating prices for FI and SE4 areas. The range of the regulating price is much wider than the range of the Elspot price for both areas, besides we occasionally observe spikes in the regulating price process, when the level reaches extreme outliers. As the Elspot price is more confined than the regulating price, we will model the Elspot price as a deterministic periodic function with the help of Fourier series. This choice still captures the observed seasonality effect, but simplifies the model by avoiding stochastic terms. Nonetheless, as Fig. 1 indicates, there is significant deviation in the Elspot prices between different days. Energy prices indeed tend to have weekly and monthly seasonality effects, moreover the levels of supply and demand depend on additional unpredictable factors. That said, we approximate the Elspot price with a Fourier series fit separately for each day. Denoting the Elspot price by P, we use the following approximation:

$$P(t) = \frac{a_0}{2} + \sum_{j=1}^{2} a_j \cos(\omega j t) + b_j \sin(\omega j t) \tag{22}$$

where $\omega = \frac{2\pi}{24}$.

We illustrate the goodness of the Fourier series approximation in Fig. 2, where for a randomly chosen week we plot the historical Elspot price values along with the approximated values for both the FI and SE4 areas.

As noted earlier, we can examine the X difference process by simply calculating $X = R - P$. The X difference process for both FI and SE4 areas is shown at the bottom of Fig. 1. We can see 0-reversion on the bottom two graphs, meaning although the difference of actual regulating and Elspot prices can be quite large at times, this difference often crosses level 0. In addition, by inspecting the graphs we can also see periods – in the vicinity of spikes – when the volatility and so the actual values increase. This can be contrasted with calm periods when the difference level stays much closer to 0. These high and low volatility periods are clustered. Whenever the process takes unusually large values, it usually remains hectic for a quite long period of time, whilst in contrast the process also remains quiet for many consecutive settlement hours. This suggests that process X switches between states, and whenever a state change happens, the process follows new dynamics. Allied with these observations, the regime-switching diffusion process we introduced in Section 2.1 seems an appropriate way to model the price difference between the regulating and day-ahead price (X). This choice is supported by the corresponding literature, as several recent studies observed regime-switching energy prices on various commodity markets including Alizadeh et al. (2008) and Scarcioffolo and Etienne (2021). Moreover, not restricted to energy markets, De Grauwe and Vansteenkiste (2007), Frömmel et al. (2005), Schaller and Norden (1997) found evidence of switching behaviour between low and high volatility regimes driven by a two-state Markov process. Therefore, we consider the case scenario in which the energy economy follows a two-regime model ('calm and turbulent' periods). We continue by using regime-switching Ornstein–Uhlenbeck process, meaning that $X(t)$ satisfies the SDE outlined in (4):

$$dX(t) = \kappa_{\Lambda(t)}(-X(t))dt + \sigma_{\Lambda(t)}dW(t), \quad X(0) = x_0 \tag{23}$$

where $\Lambda(t)$ is a continuous time Markov chain. $\Lambda = I$ indicates low volatility regimes and $\Lambda = II$ high volatility regimes. In line with the discussion above, we assumed 0-reversion for both regimes I and II, thus we set $\theta_I = \theta_{II} = 0$. In order to distinguish points corresponding to regime I and regime II, we use a change-point analysis technique. Detecting changes in time series is of interest in different research areas, including finance (Lenardon & Amirdjanova, 2006; Thies & Molnár, 2018; Ye et al., 2012). In particular, to conduct this task, we use the changepoint R package (see Killick & Eckley, 2014) and its

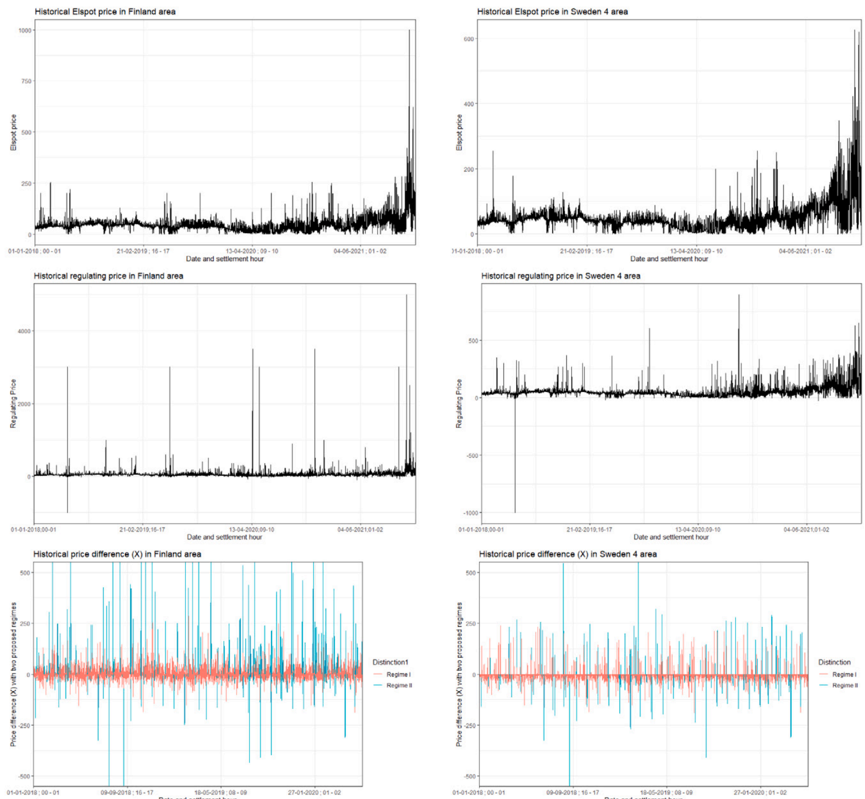

Fig. 1. Historical prices for FI and SE4 areas. On the top row we can see P Elspot prices, on the middle row we can see R Regulating prices, while on the bottom row we can see $X = R - P$ price difference values.

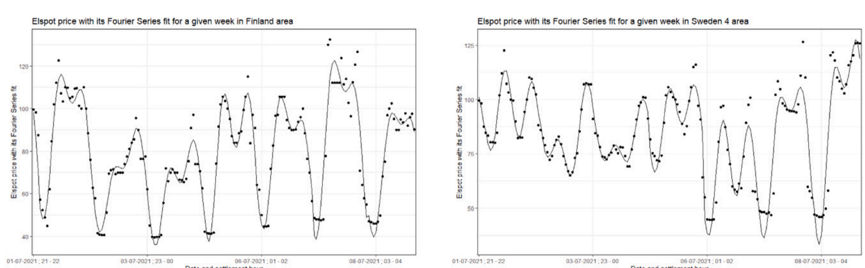

Fig. 2. Historical Elspot prices for FI and SE4 areas with the corresponding daily Fourier Series fits described by Eq. (22).

cpt.meanvar R function, which detects regime switches based on the changes in the means and variances of a given time series.[3] As an output we obtain a set of intervals.

Next, we classify all the intervals into low and high volatility regimes by calculating the standard deviations of the observations in each obtained interval and make a decision based on these calculated standard deviations. In case the standard deviation of the particular interval is greater than Tn times the mean of the standard deviations of all intervals, the interval is part of regime II. Otherwise, the interval is part of regime I (Tn is the chosen threshold number). Fig. 1 shows the result of this classification for both FI and SE4 areas for $Tn = 1$. This figure indicates the goodness of the separation as red and green lines show different characteristics in terms of volatility and range. The periods associated with spikes are almost exclusively observed in regime II, while difference values stay in a confined region in regime I. It should also be noted that both regimes exhibit mean reversion behaviour. Once our data points have successfully been separated into different regions, we can discuss the estimation of the parameters of the Ornstein–Uhlenbeck process within a region. We used 'the '*fitsde*" built in R function with its Euler pseudo-likelihood estimator and "Nelder–Mead" optimization method specification. Numerous variations of specifications have been conducted to ensure the robustness of this choice. Table 2 shows all the estimated parameters for both Finland and Sweden 4 areas using three different scenarios associated with three threshold numbers ($Tn = 0.5$, $Tn = 1$ or $Tn = 2$). The three scenarios show substantial differences within the estimated parameters, but as a common trait for each fixed scenario there is a significant difference between σ_I and σ_{II} and moderate difference between κ_I and κ_{II}. For the sake of this case study we will use just $Tn = 1$, which according to Fig. 1 shows a good representation of the point separation.

[3] As for calibrating the cpt.meanvar R function, Bayesian information criterion was specified for the penalty input. We checked and the output value is not substantially impacted by the choice of this penalty input. As indicated, we used binary segmentation method. Moreover, changepoints can distinguish regimes and thus might ultimately be associated with the event of external shocks in the economy. Therefore, we constrained the frequency and number of the changes by setting the minimal segment length input to 12 and the maximum number of changepoints input to 200.

Table 1
This table presents the $\mathbf{a} = (a_0, a_1, a_2)$ and $\mathbf{b} = (b_0, b_1, b_2)$ estimated parameters for the four representative days and for both regions.

Market	Date							
	20 Jan		16 May		8 Aug		25 Oct	
	a	**b**	**a**	**b**	**a**	**b**	**a**	**b**
FI	27.53	0	23.61	0	28.21	0	20.18	0
FI	−5.46	−2.6	1.56	−6.68	−10.05	−6.17	−4.31	−6.13
FI	0.61	−1.63	−9.66	2.39	−16.09	−1.97	−9.7	−4.01
SE4	27.53	0	22.87	0	19.46	0	16.92	0
SE4	−5.46	−2.6	0.69	−5.94	−1.43	−2.35	−2.44	−3.56
SE4	0.61	−1.63	−7.41	1.76	−0.87	0.76	−4.62	−0.53

Table 2
Here we present the estimated parameters to both Finland and Sweden 4 areas. The last two columns show the number of occurrences in Regime I and Regime II for the given area and based on the applied threshold number used for classification.

Market	Tn	κ_I	σ_I	κ_{II}	σ_{II}	$p_{I,II}$	$p_{II,I}$	# I	# II
FI	0.5	0.332	13.669	0.383	202.602	0.0016	0.0143	31 524	3540
SE4	0.5	0.283	4.487	0.371	29.721	0.0024	0.0036	20 949	14 115
FI	1	0.326	17.733	0.414	365.592	0.0008	0.0250	33 982	1082
SE4	1	0.354	7.902	0.360	43.838	0.0015	0.0076	29 385	5679
FI	2	0.343	22.037	0.404	611.698	0.0004	0.0380	34 722	342
SE4	2	0.386	12.125	0.350	64.926	0.0009	0.0155	33 198	1866

Table 3
Here we present the estimated parameters to both Finland and Sweden 4 areas by not classifying the sample into two regimes and using all 35 064 observations for the estimation of both countries, respectively.

Area	κ	σ
FI	0.38	64.994
SE 4	0.37	19.179

For four particular days representing each season in 2019, we present the Fourier series estimated parameters for both areas in Table 1. We can see that in Finland $P(t)$ over the 24 h period of 20th January will have the smallest range whilst 8th August has the largest (indicated by the larger magnitudes of a_1, a_2, b_1 and b_2). In Sweden SE4, the largest range is on 16th May and the smallest on 8th August. In Table 3 we show the fitted parameters if we assume that there is only one regime (and so no switching) and we can see the values sit somewhere in between the values we obtain for regime I and II.

Note that throughout this section we discussed the stochastic modelling of X and deterministic modelling of P. However, as also clarified in Section 2, the trader interacts on the regulating market and makes transactions linked with the regulating price R. As $R = X + P$, the predictability of the P day-ahead price is already implicitly incorporated into process R.

4. Numerical results for storage trading on the regulating market

We remind the reader that implicit to this model is that the operator is a price taker and does not influence prices by bidding or not bidding on the market. The result of this is that our solutions would trivially scale with battery size, at least in the mathematical sense rather than physical or economical, that according to the equations we write down the revenue generated by a battery that is twice the size (with suitably scaled discharge/charge rates) will generate twice the revenue. As a result all the valuations presented here are in Euros for a battery with a capacity of 1 MWh, and the strategies we derive should work well for smaller scale battery operators but not for larger grid-scale operators. We also make the assumption that the charge rate does not influence the strategy of the operator, and this would only be valid if the battery characteristics are such that the charge/discharge rates are of similar order to the capacity given the time scales of delivery. So for example, given the typical hourly or half hourly timescales in the market, we are restricted to 1 MWh batteries that can charge/discharge at a rate between 1 MW to 2 MW. This way the battery can be completely emptied or filled during the delivery period and our assumption of instantaneous empty or fill of the battery should not have too much impact.

Now we present numerical results when solving the set of PDEs (16) as outlined in Section 2.3. We consider two cases, that of FI and SE4, and in both cases we choose the $Tn = 1$ threshold parameter fits for the OU process in each regime. Starting our calculations with a terminal condition

$$V_S^k(x, T) = 0, \tag{24}$$

we search for the periodic solution as $T \to \infty$ by iterating through 24 h periods until

$$\|v^k_{S,i,0} - v^k_{S,i,jMax}\| < \text{tol}.$$

Now since we use a repeating 24 h set of electricity prices, the periodic solution can be interpreted as the net present value of all revenues for operating the battery in perpetuity. As we see later, the resulting solution does not vary significantly according to the states I, II, E, F or the variables x and t at least over the solution domain we have defined. This is because of the mean reverting property of the underlying stochastic process, and the fact that the initial starting point has little influence when valuing over an infinite time horizon. This leads us to define the equivalent fixed continuous yearly revenue rate for the revenue generated on this particular day by selecting a particular point in the grid and multiplying by the corresponding annual interest rate:

$$A = r \cdot v^I_{E,0,0}. \tag{25}$$

This gives us a result that is almost independent of the discount rate, given the time scale over which revenues are generated (hours) versus the time scale of the discount rate (years). Now using a representative set of example days from the year, such as those outlined in Table 1, we can estimate the expected yearly revenue rate by running our perpetual solution for all of those days and simply taking the average.

4.1. The effect of cost on the yearly revenue rates

We now demonstrate annuity calculations for our two case studies with a variety of different penalty functions. We choose the penalty function to be a fixed payment on using the battery (i.e. charging or discharging) so that $C(x,t) = C_0$, and we do not consider any other operating costs or subsidies. The focus of this study is on the numerical

Table 4
In this table the estimated A equivalent fixed continuous yearly revenue rates are shown according to the formula (25). These A rates are shown for the $P(t) \equiv 0$ simplified case and for the cases when $P(t)$ reflects the choice of the four particular days, respectively. Finally the last column shows the average of A rates over those four days. The A rates are shown separately for FI and SE4 areas and with different C_0 penalty values. Valuations are in Euros for a battery with capacity of 1 MWh and $r = 0.1$ is assumed for all calculations.

Penalty	Market	$P(t) \equiv 0$	$P(t)$ from date:				Average
C_0			20 Jan	16 May	8 Aug	25 Oct	
1	FI	39 345	39 793	42 922	47 770	43 322	43 452
5	FI	33 294	33 807	37 030	42 176	37 469	37 621
10	FI	28 562	29 082	31 616	36 309	32 007	32 253
20	FI	22 436	22 825	24 202	27 404	24 452	24 721
1	SE4	17 543	18 431	21 455	17 782	19 294	19 241
5	SE4	12 729	13 721	16 150	12 975	14 207	14 263
10	SE4	9395	10 123	11 585	9563	10 286	10 389
20	SE4	5967	6190	6664	6016	6236	6277

and analytical techniques to value the option, so practitioners or other interested parties would need to add this in at a later stage. To give an outline of how one might provide an estimate of such a cost to put in the model, we present the following ideas. The penalty should take into account any costs related to the act of trading since this penalty is only incurred when a trade is executed. Typically we might imagine that this would include some accounting for the degradation and maintenance of the battery and any market fees related to trading. For the degradation, we could assume that this should be money set aside from revenues to replace the battery after it gets close to the maximum number of cycles. Therefore, the cost in this model could be roughly estimated as

$$C(x,t) \approx \frac{\text{Capital Cost of Battery Replacement}}{\text{Estimated number of cycles}} + \text{Transaction fees}\ .$$

We show results for $C_0 = 1, 5, 10, 20$, with the idea that those penalties will transition from scenarios where trading against the daily cycle in forward prices is possible, to scenarios where only trading on the regulating prices is possible. Our expectation will be that increasing costs should lower the revenues generated, and the effects might be different when considering that the different regions have different underlying parameters.

First consider the results for Finland in Table 4. As expected, increasing penalties will vastly reduce the revenues generated on each simulated day. Given that there is an optimal balance between the number of trades versus the profit on each trade, we see that the effect of increasing cost does not have a linear effect on the value, as greater costs can be offset by looking for more profitable but infrequent trades. This can be done by submitting higher ask prices when the battery is full and lower bid prices when the battery is empty. You will also notice this effect under inspection of the results for Sweden. Here, the main difference in values generated in different markets will be caused by the parameters of the SDE for the regulating price. This is born out in the fitted parameters as we calibrate Sweden's market with a much lower variance in region II giving much less scope to exploit those profitable but infrequent trades. In particular, on 20th Jan 2019 when the fitted spot price $P(t)$ is the same in both markets, we can see that the revenues in Sweden will be much lower than those we might generate in Finland.

Looking horizontally across Table 4 for fixed penalty and different days, the amplitude of the cycles have an influence on the valuations for that particular day. In Finland, we see that 20th Jan has the smallest valuation and 8th August has the largest corresponding to the days with smaller or larger ranges in $P(t)$ respectively, and similarly in Sweden 16th May has the highest valuation and biggest range in $P(t)$. We would expect this as larger ranges allow the battery operator to take advantage of so called time arbitrage in the market.

In fact in both regions we can see that the fitted parameters in region I and II have a much greater effect on the valuation than the different daily cycles. This means that a method (such as the one outlined in the next section) that can quickly capture the strategy and valuation when P is fixed should be able to replicate most of the important features of the optimal strategy and the resulting valuation. We have included a column headed $P(t) \equiv 0$ where we rerun our calculations based on zero underlying price, and we see that the values generated are just 5% to 10% lower than when there is a variation in the time profile of price. The difference between the results reduces when penalties on charging are increased as the optimal strategy veers away from time arbitrage.

4.2. The effect of cost on the optimal strategy

In Fig. 3 we show how the optimal strategy changes during the day in the two different markets of Finland and Sweden, and with different penalties on charge/discharge. In general, although the magnitude of the spread may change, the features of the optimal strategy look to be quite similar across the different markets and penalty scenarios. For the low penalty scenarios of (a) and (c), the strategy in regime I is to follow the price (light blue line) with a small margin either side, and in regime II that margin is increased. The optimal price we offer to sell at is raised in advance of the peak (at around 15:00) to ensure we do not sell before the peak, and similarly drops lower in advance of the dips to ensure we buy at the lowest price. This has the effect of suggesting that the phase of the sine waves for the bid/ask prices are shifted by one or two hours compared to the price $P(t)$. In Figs. 3(b) and (d) there is a much higher penalty, so the amplitude chasing the peak and dip in the price process is much less pronounced, simply because the profit margin we seek to make back after penalties is much larger.

4.3. The effect of regime switching

First in Fig. 4 we demonstrate the effect regime switching has on the strategy for optimal trading. Inspecting the four figures, we see that the bid/ask strategies in (a) and (b) are almost identical to the strategies obtained in (c) and (d) when regime switching is turned off. In other words, the primary concern when setting the bid/ask spread is what the current underlying stochastic process might do, rather than what it might do if a regime change happens. This is a result of the fact that the regime changes are Markovian and therefore completely random and cannot be predicted. So our optimal strategy is only really concerned with what might happen in the current regime when making a decision over the next short time period. We use this observation to motivate the approximation later on where the solution for each regime is derived independently and then merged in a simplistic fashion to derive the combined result.

Now since one of the main contributions of this paper is to outline how to calibrate and solve for regime switching models, it seems appropriate to outline what would happen if we were to simply assume that there is just a single regime. In Table 5 we repeat the results from Table 4 but now we only model a single regime with parameters defined by Table 3. Looking just at Table 5, the effect of increasing penalty is similar, however there is one big difference between the

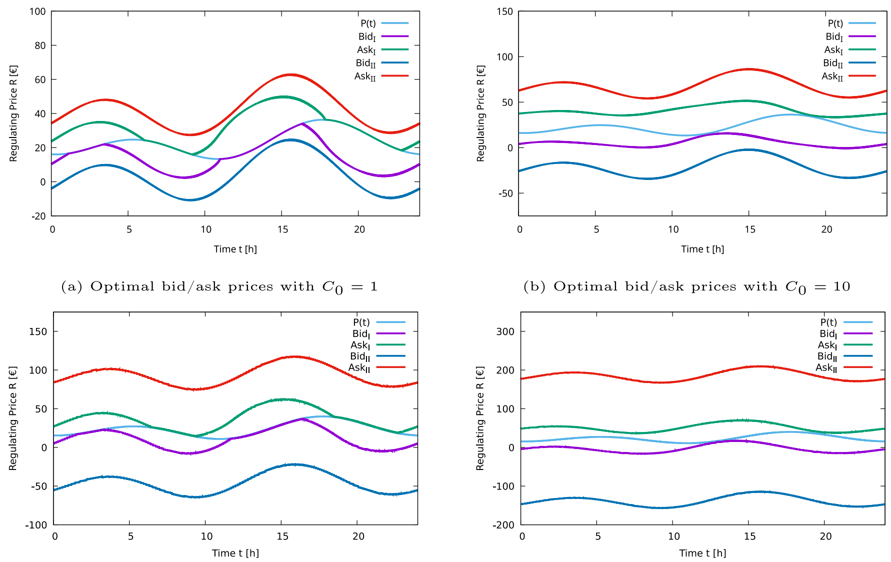

(a) Optimal bid/ask prices with $C_0 = 1$

(b) Optimal bid/ask prices with $C_0 = 10$

(c) Optimal bid/ask prices with $C_0 = 1$

(d) Optimal bid/ask prices with $C_0 = 10$

Fig. 3. In these figures we show the optimal $\text{bid}_S/\text{ask}_S$ prices for regime S as a function of time over the 24 h time period. In (a) and (b) we show the optimal strategy in Sweden on 16th May 2019, and (c) and (d) the optimal strategy in Finland on the same date.

Table 5
In this table we compare the results from a single regime (Table 3) with the regime switching model (Table 2). Estimated A equivalent fixed continuous yearly revenue rate with $P(t) \equiv 0$, and the corresponding average over the four considered days are both displayed. Fixed revenue rates are generated from the periodic solution according to the formula (25). Valuations are in Euros for a battery with capacity of 1 MWh. The interest rate $r = 0.1$ is assumed for all calculations, thus the rate used for calculating the A fixed revenue rate is $r = 0.1$ as well.

Penalty	Market	Single regime yearly revenue rates		Regime switching yearly revenue rates	
C_0		$P(t) \equiv 0$	$P(t)$ seasonal	$P(t) \equiv 0$	$P(t)$ seasonal
1	FI	93 248	94 453	39 345	43 452
5	FI	82 674	83 990	33 294	37 621
10	FI	73 750	75 124	28 562	32 253
20	FI	60 609	61 970	22 436	24 721
1	SE4	25 154	26 017	17 543	19 241
5	SE4	18 773	19 714	12 729	14 263
10	SE4	13 860	14 707	9395	10 389
20	SE4	7651	8211	5967	6277

results here and those in Table 4, and that is that modelling with a single regime leads to much higher valuations across all cases. In the results for Finland we see the values increase from around €40,000 when penalties are $C_0 = 1$ to over €90,000 in the same case for a single regime. This should not be seen as a good result, as we believe it to be much less accurate. Implementing this strategy in a real market would not lead to better outcomes, since the model wrongly attaches a higher variance of $\sigma = 64.994$ at all times when the real data has long periods of low volatility followed by short periods of high volatility. So, for example the strategy it delivers would set the ask prices too high during low volatility periods and too low ask prices when the volatility is high.

4.4. Strategies in the real world

The results from Table 5 leads us nicely onto a comparison of how the strategies perform when applied to the real data. To capture 'real data', we simply run through the hourly recorded data of regulating prices for the last 90 days of 2021, and keep track of a trading account in which:

- if the battery is full and the regulating price $R(t)$ is above our ask price, we sell electricity at the ask price and empty the battery — updating the trading account by adding ask price and minus the penalty cost;
- if the battery is empty and the regulating price $R(t)$ is below our bid price, we buy electricity at the bid price and fill the battery — updating the trading account with minus the bid price and minus the penalty cost.

To give some benchmark values to compare against, as well as taking prices directly from the recorded observations, we also simulate the values of the regulating price $R(t)$ using our regime switching model.

Now, to keep things simple, we apply the bid/ask strategies found by solving (17) with either a single regime or regime switching for the underlying dynamics and $P(t) \equiv 0$. In this case, the optimal regulating bid/ask prices can be expressed in terms of X and are symmetric around zero, so if the optimal ask price is X^*, the optimal bid price is $-X^*$. When using simulated regulating prices with $P \equiv 0$, they correspond directly to

$$\text{Ask}(t) = X^*, \qquad \text{Bid}(t) = -X^*. \tag{26}$$

However, when using 'real data' or the observed values of the regulating price in the trading algorithm, clearly P is not zero and varies with time, so we must convert back X^* back into real pricing strategies by

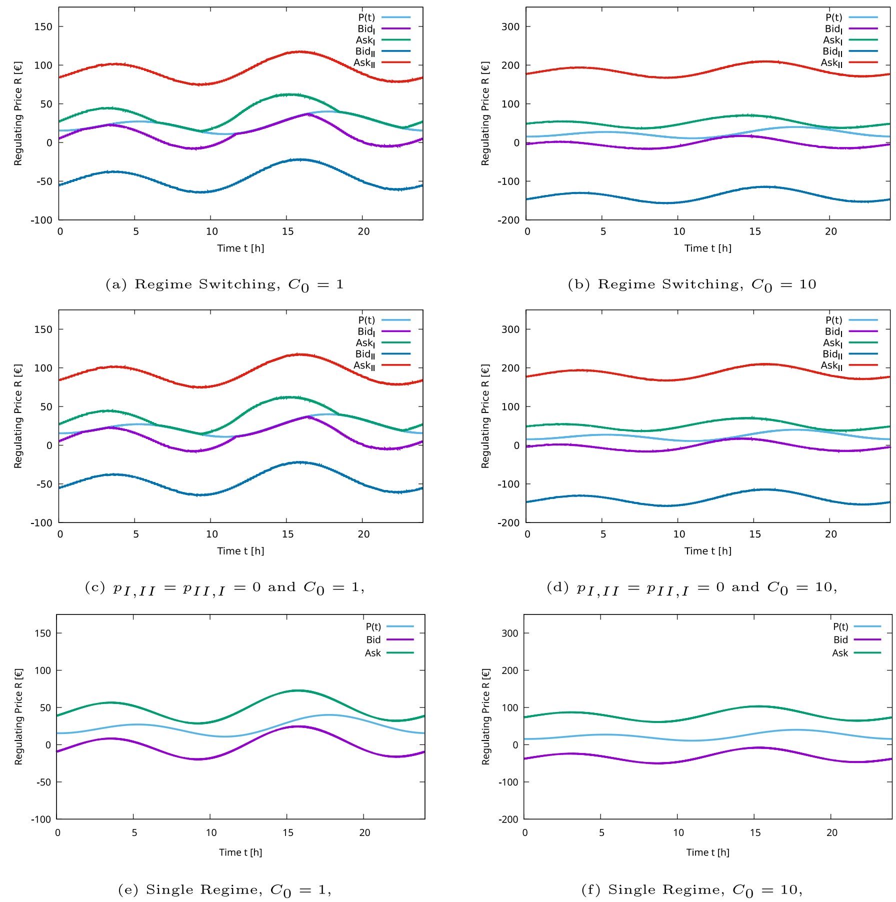

(a) Regime Switching, $C_0 = 1$ (b) Regime Switching, $C_0 = 10$

(c) $p_{I,II} = p_{II,I} = 0$ and $C_0 = 1$, (d) $p_{I,II} = p_{II,I} = 0$ and $C_0 = 10$,

(e) Single Regime, $C_0 = 1$, (f) Single Regime, $C_0 = 10$,

Fig. 4. Comparison of the optimal bid/ask strategy when regime switching enabled as in Table 2 in (a) and (b), and then using the same OU parameters but setting $p_{I,II} = p_{II,I} = 0$ in (c) and (d), and finally running with a single regime from Table 3.

Table 6
Here we present the simulated revenue generated (V) during 90 days of trading in the FI regulating market, when implementing different optimal trading strategies. For the top rows, we assume $P \equiv 0$ and so the regulating price $R(t) = X(t)$ is simulated using SDEs (2)–(4) for a single regime (Table 3) or regime switching (Table 2) as appropriate, whilst the bottom rows use the actual regulating price $R(t)$ observed on that date. The ask/bid thresholds are set in accordance with (26)–(27), where the different X^* values are calculated as part of the numerical solution to (17) with $P \equiv 0$.

Penalty	Market	Single regime 90 days trading		Regime switching 90 days trading	
C_0		Strategy	V	Strategy	V
10	FI (simulated)	$X^* = 55.5$	11 565	$X_I^* = 25.6$ $X_{II}^* = 162.2$	3626
20	FI (simulated)	$X^* = 72.2$	9932	$X_I^* = 35.7$ $X_{II}^* = 208.4$	3449
10	FI (real data)	$X^* = 55.5$	2363	$X_I^* = 25.6$ $X_{II}^* = 162.2$	2551
20	FI (real data)	$X^* = 72.2$	1669	$X_I^* = 35.7$ $X_{II}^* = 208.4$	2410

adding in the spot price. For our purposes, we assume

$$\mathrm{Ask}(t) = \bar{P}(t) + X^*, \quad \mathrm{Bid}(t) = \bar{P}(t) - X^* \tag{27}$$

where $\bar{P}(t)$ is the daily moving average of the last 24 h of prices.

The results for this trading account are shown in Table 6, using the real data from Finland and calibrated parameters from Tables 2 and 3. In the top half of the table, we run a 90 day simulation of $X(t)$ using (4) to generate the regulating prices $R(t)$. In this case we see the single regime simulations appear to give greater V values, due to the high value of variance used as input to (4). This is similar to the results described in Section 4.3, where the single regime gives higher values. However, given the results in the bottom half of the table when the real regulating prices are used in the trading account, we can see that the values from the simulation with a single regime model significantly over estimates the real V by an order of magnitude. The strategy from the regime switching model does better, although even then our simplified model is still missing some important features of the real world, such as periods when the regulating market is not called upon (and we do not allow trading during these periods in the real world simulation). When the trading account uses the real regulating prices as input, and compares the bid/ask strategy derived from a single regime or regime switching model, we see that the regime switching strategy is the best strategy, with an 8% increase compared to a single regime strategy when the penalty is $C_0 = 10$ and over 40% when $C_0 = 20$.

So clearly there are two reasons to favour the regime switching model over the single regime model:

i The strategies perform better in the real world;
ii The estimated values inside the model are closer to the performance in the real world.

4.5. Summary

We demonstrate in this section how the solution to the PDE switching problem delivers an optimal bid and ask price strategy if we know which regime we are in (high or low volatility). These outputs could be used by storage operators on the market to build their own strategy, but importantly it also allows them to generate the value of trading on a particular day. With a sufficient variety of trading days, and a model for the penalties on a charge/discharge cycle, the value of investing in more storage or new technology could be calculated.

The key results from this section are that:

- the optimal strategies are affected more by time varying prices when the penalties are low, and less when penalties are high,
- the bid/ask spread in the optimal strategy appears to be primarily driven by the stochastic dynamics of the current regime,
- single regime dynamics over estimate the value and perform poorly on real world data,
- regime switching dynamics do better at estimating values and the strategies perform better on real data.

In particular, the way the optimal strategy appears to be independent of the regime switching leads us to believe that a quick analytical approximation could be used to generate annuities when trading on the regulating market. In the next section we show how by making the simplifying assumption of fixed prices we can generate approximate valuations for the full time varying problem that match values within 5%–10% given realistic scenarios.

5. Quasi-analytic approximations

In this section we provide a quasi-analytic approximation by applying simplifying assumptions. We note that the results of Section 4 indicate that optimal strategy is closely associated with the first hitting time of the underlying process at some fixed levels and the time varying feature of the optimal strategy only appears due to time dependent $P(t)$ values. Therefore, we assume that $P(t) \equiv 0$. Besides, as indicated in Fig. 4, turning off regime switching has negligible effect on the reigning optimal bid/ask strategy. Thus, in order to successfully apply results corresponding to the theory of first hitting time, we also assume $p_{I,II} = p_{II,I} = 0$, and so we initially consider the problem separately by being in regime I or in regime II.

Hence, we first consider a single regime Ornstein–Uhlenbeck process with zero reversion:

$$dX(t) = -\kappa X(t)dt + \sigma dW(t), \quad X(0) = u. \tag{28}$$

The discounted first hitting time problem is naturally related to the problem of expressing corresponding Laplace transforms and so we recall results of Alili et al. (2005) who considered the problem with a one unit volatility process as:

$$dU(t) = -\kappa U(t)dt + dW(t), \quad U(0) = u \tag{29}$$

and accordingly we introduce $\tau_a = \inf\{s > 0 : U(s) = a\}$. Based on Alili et al. (2005), we know that for $r > 0$ and $u < a$, the Laplace transform $B_a^{r,\kappa}(u) = \mathbb{E}_u^{\kappa}[e^{-r\tau_a}]$ is given by the following formula:

$$B_a^{r,\kappa}(u) = \frac{e^{\kappa u^2/2} D_{-r/\kappa}(-u\sqrt{2\kappa})}{e^{\kappa a^2/2} D_{-r/\kappa}(-a\sqrt{2\kappa})} \tag{30}$$

where $D_\nu(\cdot)$ is the parabolic cylinder function and $\mathbb{E}_u^{\kappa}$ corresponds to $\mathbb{P}_u^{\kappa}$ that denotes the law of $U(t)$ given that $U(0) = u$.

By virtue of change of variables we express the solution to the problem with arbitrary volatility. This means that we also introduce

$$\tau_a^{\sigma} = \inf\{s > 0 : X(s) = a\}$$

and a transformed process as $\hat{X}(t) = X(\frac{t}{\sigma^2})$ for $t \geq 0$. Note that $d\frac{t}{\sigma^2} = \frac{dt}{\sigma^2}$ and $dW(\frac{t}{\sigma^2}) = \frac{d\hat{W}(t)}{\sigma}$, for a given $\hat{W}$ Wiener process, and so we obtain:

$$\begin{aligned} &d\hat{X}(t) = -\hat{\kappa}\hat{X}(t)dt + d\hat{W}(t), \quad \hat{X}(0) = u \\ &\hat{\kappa} = \frac{\kappa}{\sigma^2}, \quad \inf\{t > 0 : \hat{X}(t) = a\} = \tau_a^{\sigma}\sigma^2 \end{aligned} \tag{31}$$

Having established the relations in (31), we can also derive the expected discounted value of the first passage time problem for the process given by (28) with the help of parabolic cylinder functions as:

$$\mathbb{E}_u^{\kappa}[e^{-r\tau_a^{\sigma}}] = \mathbb{E}_u^{\frac{\kappa}{\sigma^2}}[e^{\frac{-r}{\sigma^2}\tau_a}] = B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(u).$$

Theorem 1. *Assume that the underlying process is given by (28) and assume that $x < a$. Consider a battery storage operator that perpetually trades by repeatedly making transactions when the underlying process first hits level a from below and subsequently when it first hits $-a$ from above, and at each transaction they receive $M = a - C_0$ in cash. Their total expected discounted profit can be expressed as:*

$$M \cdot e_x \cdot (1 + e_{-a} + (e_{-a})^2 + \dots) = \frac{(a - C_0)e_x}{1 - e_{-a}},$$

where $e_x^a \stackrel{\Delta}{=} B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(x)$.

The perpetual value $V_F(x;a)$ is given by

$$V_F(x;a) = \begin{cases} \frac{(a-C_0)e_x^a}{1-e_{-a}^a} & \text{if } x < a \\ \frac{(a-C_0)e_{-x}^a}{1-e_{-a}^a} + x - C_0 & \text{if } x \geq a. \end{cases} \tag{32}$$

Proof. See Appendix. □

Given that $V_F(0;a) = V_E(0;a)$ we can now express the value of operating the battery according to a strategy of bids and asks at $x = -a$ and $x = a$ as a function of a. Assuming the parameters $C_0 > 0$, $r > 0$, $\kappa > 0$ and $\sigma > 0$, it is trivial to show that the optimal strategy can be derived by solving the convex optimization problem

$$\max_a \frac{(a - C_0)e_0^a}{1 - e_{-a}^a} \text{ s.t. } a \geq 0. \tag{33}$$

Since the parabolic cylinder functions are already implemented in the majority of programming languages (for example Python), the values of e_x^a in (33) can be quickly computed and the optimization problem can be easily solved in a matter of seconds using standard routines.

5.1. Multiple regime case

Having outlined the valuation problem for the single regime case, we now provide an approximation method for the first hitting time problem with two different regimes. We do so by seeking the $\phi = (\phi_I, \phi_{II})$ stationary distribution of the Markov chain, and calculating $V = V_I \cdot \phi_I + V_{II} \cdot \phi_{II}$, where V_I and V_{II} are the separate lifetime values obtained from Theorem 1 with strictly regime I and II, respectively.

The Markov chain introduced in (3) is irreducible and aperiodic, therefore due to the Fundamental Theorem of Markov chains a unique stationary distribution exists. Given the transition matrix P, this unique stationary distribution satisfies the following matrix equation: $\phi = \phi P$. This means that ϕ is a left eigenvector of P with an eigenvalue of 1, or equivalently ϕ^T is a right eigenvector of P^T with an eigenvalue of 1. This ϕ can be readily computed as P is given with the help of transition probabilities in Table 2.

Table 7
This table compares the equivalent A fixed continuous yearly revenues of numerical solutions of the variational inequalities with quasi-analytic approximations, according to the formula (25), and assuming zero Elspot price values; $P \equiv 0$. For both numerical and analytical approaches, values are presented with the underlying process following strictly Regime I; Regime II and allowing regime switches between the two of them. As for the quasi-analytic approximation, we maximized A over potential values of a hitting levels. Regarding the regime switching case, numerical values are obtained as solutions of variational inequalities, whereas quasi-analytic approximations are obtained with the help of stationary distribution as discussed in Section 5.1. Valuations are in Euros for a battery with capacity of 1 MWh and $r = 0.1$ is assumed for all calculations.

Penalty	Market	Numerical			Quasi-analytic		
C_0		Regime I	Regime II	Switching	Regime I	Regime II	Switching
1	FI	21 803	570 067	39 345	21 806	570 248	38 812
5	FI	16 217	549 086	33 294	16 219	549 257	32 747
10	FI	11 921	530 666	28 562	11 926	530 832	28 016
20	FI	6518	502 038	22 436	6521	502 197	21 891
1	SE4	9051	60 195	17 543	9054	60 211	17 486
5	SE4	5005	51 473	12 729	5007	51 488	12 669
10	SE4	2440	44 246	9395	2441	44 259	9334
20	SE4	390	33 887	5967	390	33 899	5913

5.2. Comparison of results

In Table 7 we compare the fixed continuous yearly A revenues obtained by the numerical solution of the variational inequalities against the quasi-analytic approach outlined previously in this section. We can make a one-to-one comparison between the corresponding results of the two approaches, to obtain that for regimes I and II all the differences are less than 0.1% of the actual values. This is reassuring about the implementation of the two methods, and the difference can be attributed solely to the error in the numerical computations. Regarding the regime switching case, the numerical solution is still obtained as solution of the variational inequalities, whereas the quasi-analytic approach used the approximation technique by calculating the weighted average of the single regime values multiplied by the stationary probabilities. Indeed, this approximation technique does not examine a process in which regime switches do happen, it rather averages out the values of two optimal stopping problems. Therefore, one would intuitively expect that the corresponding value of the optimal switching problem, in which the user is aware of regime switches is higher. This is because they could exploit the information that future regime switches are to happen, and the dynamics of the process is to change. Comparing the last column of both approaches in Table 7 reveals that the numerical value is a few percent higher for the numerical solution for all different cases. Nonetheless, the quasi-analytic approach is computationally less time consuming and provides a rather accurate estimation.

6. Conclusions

We find that the regime switching model captures more realistic valuations and trading strategies that can take advantage of the nature of the real data, without over-predicting what might be possible. The single regime model gives valuations over twice the size of the two regime switching model, indicating a large overshoot. We are also able to observe that the way spot price P varies during the day only has second order effects on the valuation, with fixed price model $P(t) = 0$ valuations coming to within 5 to 10% of the full model.

After observing the symmetries and behaviour of the optimal strategy, we believe that analytic approximations might deliver good benchmark results in a fraction of the time that full numerical schemes take. When comparing single regime models, the full numerical solution and quasi analytic agree to 0.01%, which is comparable to the grid errors in the numerical scheme. Once the simplified model to generate the full regime switching value is used, the difference is around 1% to 2%. Combining this with the difference between fixed price and the full model, we should expect the quasi-analytic solution to deliver approximations within 10% of the full numerical solutions for the majority of realistic scenarios. This can be delivered in a few seconds compared to 10 s of seconds to solve the full numerical finite difference scheme with sufficient accuracy. This technique could prove useful in a variety of scenarios where solving the full model is too complex or too time consuming, and a fairly accurate analytic approximation is good enough. The analytic approximation should work on any Markovian regime switching model with optimal control, as long as the rate of switching happens on a slower timescale than the optimal controls. In a world where investment in renewable energies is vital, being able to quickly assess the potential benefits of installing battery capacity on the market is extremely important. Our model assumes that the battery is small enough to not impact prices on the market, so it cannot directly be used to value large projects, but as it takes costs and prices as input to produce an optimal strategy for the operator, it could be used as part of an agent based model to simulate larger market behaviours.

There are numerous assumptions in this model, but it is worth mentioning that implementing such strategies in the real world is complicated by the fact that one cannot simply observe whether you are in regime I or regime II. If the operator spots high volatility periods too late, they may already have executed the trade and lost an opportunity. How running volatility is calculated or how reliable change point analysis can be on live data will influence how well this method works. Trying to bring this uncertainty into the model using Bayesian learning, or simple testing strategies on real data, are both interesting topics that could motivate further research. It is also worth considering different stochastic processes, especially non Markovian ones, to assess the accuracy of the analytic approximation.

CRediT authorship contribution statement

Paul Johnson: Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Dávid Zoltán Szabó:** Writing – review & editing, Writing – original draft, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Peter Duck:** Writing – review & editing, Writing – original draft, Validation, Investigation, Formal analysis, Conceptualization.

Appendix. Analytical approximation

We now present the proof to Theorem 1.

Proof. Define the following stopping times iteratively:

$$\mathrm{T}_1 = \inf\{s > 0 \, : \, X(s) \geq a\} \tag{A.1}$$

$$\mathrm{T}_2 = \inf\{s > \mathrm{T}_1 \, : \, X(s) \leq -a\}$$

$$\mathrm{T}_3 = \inf\{s > \mathrm{T}_2 \, : \, X(s) \geq a\}$$

$$\dots$$

for n odd

$$\mathrm{T}_n = \inf\{s > \mathrm{T}_{n-1} : X(s) \geq a\} \tag{A.2}$$
$$\mathrm{T}_{n+1} = \inf\{s > \mathrm{T}_n : X(s) \leq -a\}$$

for n even

$$\mathrm{T}_n = \inf\{s > \mathrm{T}_{n-1} : X(s) \leq -a\} \tag{A.3}$$
$$\mathrm{T}_{n+1} = \inf\{s > \mathrm{T}_n : X(s) \geq a\}.$$

The expected discounted profit after reaching two hitting times is given as: $M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_1}] + M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_2}]$. We further express the second term as:

$$M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_2}] = M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r\mathrm{T}_2}|\mathrm{T}_1]] \tag{A.4}$$
$$= M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r(\mathrm{T}_1+\hat{\mathrm{T}}_a^{-a})}|\mathrm{T}_1]] = M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_1} \cdot \mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_a^{-a}}|\mathrm{T}_1]]$$

where $\hat{\mathrm{T}}_a^{-a} = \inf\{s > 0 : X_1(s) \leq -a; X_1(0) = a\}$ and X_1 follows the same dynamic as X. As the underlying process is symmetric, $\hat{\mathrm{T}}_a^{-a}$ and $\hat{\mathrm{T}}_{-a}^{a}$ have the same distribution. Thus $\mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_a^{-a}}|\mathrm{T}_1] = \mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_a^{-a}}] = B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(-a)$ and so

$$M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_1}] + M \cdot \mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_a^{-a}}] = M \cdot B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(x) + M \cdot B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(x) \cdot B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(-a).$$

In an analogous way we can assess for n even:

$$M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_n}] = M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r\mathrm{T}_n}|\mathrm{T}_{n-1}]] \tag{A.5}$$
$$= M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r(\mathrm{T}_{n-1}+\hat{\mathrm{T}}_a^{-a})}|\mathrm{T}_{n-1}]] = M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_{n-1}}\, \mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_a^{-a}}]]$$

and for n odd:

$$M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_n}] = M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r\mathrm{T}_n}|\mathrm{T}_{n-1}]] \tag{A.6}$$
$$= M \cdot \mathbb{E}_u^\kappa[\mathbb{E}_u^\kappa[e^{-r(\mathrm{T}_{n-1}+\hat{\mathrm{T}}_{-a}^{a})}|\mathrm{T}_{n-1}]] = M \cdot \mathbb{E}_u^\kappa[e^{-r\mathrm{T}_{n-1}}\, \mathbb{E}_u^\kappa[e^{-r\hat{\mathrm{T}}_{-a}^{a}}]].$$

We solve the problem recursively, and so the expected discounted value by assuming that the trader executed n transactions following their strategy, then stops:

$$M \cdot e_u^a + M \cdot e_u^a \cdot e_{-a}^a + M \cdot e_u^a \cdot (e_{-a}^a)^2 + \cdots + M \cdot e_u^a \cdot (e_{-a}^a)^{n-1}$$

where $e_x^a \triangleq B_a^{\frac{r}{\sigma^2},\frac{\kappa}{\sigma^2}}(x)$.

Finally, the lifetime value as a limit is given as:

$$M \cdot e_u^a(1 + e_{-a}^a + (e_{-a}^a)^2 + \dots) = \frac{M \cdot e_u^a}{1 - e_{-a}^a}.$$

Now if we write the perpetual value function of this single regime problem when using a strategy to sell at $x = a$ and buy at $x = -a$, as $V_S(x; a)$. Consider that initially the battery is full, if $x \geq a$, the operator sells immediately at $p = x$ to derive the value

$$V_F(x; a) = V_E(x; a) + x - C_0 \text{ for } x \geq a.$$

Using the symmetry in x we can say

$$V_E(x; a) = V_F(-x; a)$$

and therefore

$$V_F(x; a) = \begin{cases} \frac{(a-C_0)e_x^a}{1-e_{-a}^a} & \text{if } x < a \\ \frac{(a-C_0)e_{-x}^a}{1-e_{-a}^a} + x - C_0 & \text{if } x \geq a. \end{cases} \qquad \square$$

References

Alili, L., Patie, P., & Pedersen, J. L. (2005). Representations of the first hitting time density of an ornstein-uhlenbeck process. *Stochastic Models*, *21*(4), 967–980.

Alizadeh, A. H., Nomikos, N. K., & Pouliasis, P. K. (2008). A Markov regime switching approach for hedging energy commodities. *Journal of Banking & Finance*, *32*(9), 1970–1983.

Bierbrauer, M., Trück, S., & Weron, R. (2004). Modeling electricity prices with regime switching models. In *Computational science-ICCS 2004: 4th international conference, KrakóW, Poland, June 6-9, 2004, proceedings, Part IV 4* (pp. 859–867). Springer.

Chattopadhyay, D. (2014). Modelling renewable energy impact on the electricity market in India. *Renewable and Sustainable Energy Reviews*, *31*, 9–22.

Dai, R., Esmaeilbeigi, R., & Charkhgard, H. (2021). The utilization of shared energy storage in energy systems: A comprehensive review. *IEEE Transactions on Smart Grid*, *12*(4), 3163–3174.

De Grauwe, P., & Vansteenkiste, I. (2007). Exchange rates and fundamentals: A non-linear relationship? *International Journal of Finance & Economics*, *12*(1), 37–54.

Escribano, A., Ignacio Peña, J., & Villaplana, P. (2011). Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics*, *73*(5), 622–650.

European Commission (2023). Electricity Market Design. https://energy.ec.europa.eu/topics/markets-and-consumers/market-legislation/electricity-market-design_en. Accessed April 2023.

Frömmel, M., MacDonald, R., & Menkhoff, L. (2005). Markov switching regimes in a monetary exchange rate model. *Economic Modelling*, *22*(3), 485–502.

Gudkov, N., & Ignatieva, K. (2021). Electricity price modelling with stochastic volatility and jumps: An empirical investigation. *Energy Economics*, *98*, Article 105260.

Harikae, S., Dyer, J. S., & Wang, T. (2021). Valuing real options in the volatile real world. *Production and Operations Management*, *30*(1), 171–189.

Janczura, J., & Weron, R. (2012). Efficient estimation of Markov regime-switching models: An application to electricity spot prices. *AStA. Advances in Statistical Analysis*, *96*, 385–407.

Karakatsani, N. V., & Bunn, D. W. (2008). Intra-day and regime-switching dynamics in electricity price formation. *Energy Economics*, *30*(4), 1776–1797.

Killick, R., & Eckley, I. (2014). Changepoint: A comprehensive changepoint analysis package for R. *Journal of Statistical Software*, *58*, 1–17.

Korn, R., Melnyk, Y., & Seifried, F. T. (2017). Stochastic impulse control with regime-switching dynamics. *European Journal of Operational Research*, *260*(3), 1024–1042.

Kristiansen, T. (2012). Forecasting nord pool day-ahead prices with an autoregressive model. *Energy Policy*, *49*, 328–332.

Lenardon, M. J., & Amirdjanova, A. (2006). Interaction between stock indices via changepoint analysis. *Applied Stochastic Models in Business and Industry*, *22*(5–6), 573–586.

Mao, X., & Yuan, C. (2006). *Stochastic differential equations with Markovian switching*. Imperial College Press.

Moriarty, J., & Palczewski, J. (2017). Real option valuation for reserve capacity. *European Journal of Operational Research*, *257*(1), 251–260.

Mount, T. D., Ning, Y., & Cai, X. (2006). Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, *28*(1), 62–80.

Nadarajah, S., & Secomandi, N. (2023). A review of the operations literature on real options in energy. *European Journal of Operational Research*, *309*(2), 469–487.

Nord Pool AS (2022). Market data | nord pool. https://www.nordpoolgroup.com/en/Market-data1/#/nordic/table. 14 November 2022.

Nordic Balancing Model (2021). Questions and Answers Nordic webinar - mFRR energy activation market - implementation guide. https://nordicbalancingmodel.net/wp-content/uploads/2021/01/mFRR-EAM-webinar-QA.pdf. 02 February 2024.

Øksendal, B., & Sulem, A. (2005). *Stochastic control of jump diffusions*. Springer.

Rahman, M. M., Oni, A. O., Gemechu, E., & Kumar, A. (2020). Assessment of energy storage technologies: A review. *Energy Conversion and Management*, *223*, Article 113295.

Scarcioffolo, A. R., & Etienne, X. L. (2021). Regime-switching energy price volatility: The role of economic policy uncertainty. *International Review of Economics & Finance*, *76*, 336–356.

Schaller, H., & Norden, S. V. (1997). Regime switching in stock market returns. *Applied Financial Economics*, *7*(2), 177–191.

Shah, M. I., Kirikkaleli, D., & Adedoyin, F. F. (2021). Regime switching effect of COVID-19 pandemic on renewable electricity generation in Denmark. *Renewable Energy*, *175*, 797–806.

Shao, J. (2015). Strong solutions and strong feller properties for regime-switching diffusion processes in an infinite state space. *SIAM Journal on Control and Optimization*, *53*(4), 2462–2479.

Shao, J., & Xi, F. (2013). Strong ergodicity of the regime-switching diffusion processes. *Stochastic Processes and their Applications*, *123*(11), 3903–3918.

Shao, J., & Xi, F. (2014). Stability and recurrence of regime-switching diffusion processes. *SIAM Journal on Control and Optimization*, *52*(6), 3496–3516.

Szabó, D. Z., Duck, P., & Johnson, P. (2020). Optimal trading of imbalance options for power systems using an energy storage device. *European Journal of Operational Research*, *285*(1), 3–22.

Thies, S., & Molnár, P. (2018). Bayesian change point analysis of bitcoin returns. *Finance Research Letters*, *27*, 223–227.

Thompson, M., Davison, M., & Rasmussen, H. (2009). Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics*, *56*(3), 226–238.

Tiwari, A. K., & Menegaki, A. N. (2019). A time varying approach on the price elasticity of electricity in India during 1975–2013. *Energy*, *183*, 385–397.

Wahab, M., & Lee, C.-G. (2011). Pricing swing options with regime switching. *Annals of Operations Research*, *185*, 139–160.

Weron, R. (2007). *Modeling and forecasting electricity loads and prices: A statistical approach*. John Wiley & Sons.

Weron, R., Bierbrauer, M., & Trück, S. (2004). Modeling electricity prices: jump diffusion and regime switching. *Physica A. Statistical Mechanics and its Applications*, *336*(1–2), 39–48.

Weron, R., & Misiorek, A. (2008). Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models. *International Journal of Forecasting*, *24*(4), 744–763.

Ye, W., Liu, X., & Miao, B. (2012). Measuring the subprime crisis contagion: Evidence of change point analysis of copula functions. *European Journal of Operational Research*, *222*(1), 96–103.

Zhou, Y., Scheller-Wolf, A., Secomandi, N., & Smith, S. (2016). Electricity trading and negative prices: Storage vs. disposal. *Management Science*, *62*(3), 880–898.