An Analysis of Threshold Policies for Trading in Continuous Intraday Electricity Markets

Gilles Bertrand
CORE, Université catholique de Louvain
Email: gilles.bertrand@uclouvain.be

Anthony Papavasiliou
CORE, Université catholique de Louvain
Email: anthony.papavasiliou@uclouvain.be

***Abstract*—The large-scale integration of renewable energy resources has increased the uncertainty of power production in short-term operations in recent years. This has increased the need for electricity market participants to balance their position closer to real time, in order to hedge against volatile real-time prices, thus increasing the significance o f i ntraday market trading. Trading in the continuous intraday market (CIM) is a difficult p roblem b ecause (i) t rades c an a ppear a nd d isappear at any moment, (ii) the decision of accepting or rejecting a bid is binary, and (iii) acceptance decisions need to be reached quickly in order to lock in interesting trades. In this paper, we show that the problem can be modelled as a one-stage Markov decision process if it is assumed that, by trading, the trader does not influence t he r eal-time p rice. W e f ocus o n t rading p olicies based on price thresholds for arriving rapidly to trading decisions. We analyse the behaviour of price threshold policies, and derive an analytical solution to the problem in particular cases. Finally, we demonstrate the effectiveness of the proposed trading policies by making Cn Cut Cf Cample Cest Cn Che Cerman CIM Chich gives a Crofit C f 2 0 C illions C uros i n 5 00 hours.**

***Index Terms* -- renewable energy sources, power systems, power generation dispatch**

I. Introduction

Following the introduction of the climate and energy package in Europe [1], the penetration of renewable power has strongly increased in the numerous European countries. These renewable energy resources increase the variability of supply, and therefore increase the need to correct the system dispatch closer to real time. An interesting option for market participants to balance their portfolios is to trade in the CIM, which explains the increase of liquidity in this market recently. For instance, the selling volume in the German CIM has evolved from 1005 AWh An 2010 Ao 2461 AWh An 2013 And Anally to A070 AWh An 2016 [2]. Ahe Ancrease Af liquidity An Ahis market Aecessitates A Aefinement A f A rading strategies.

The design of CIM has been addressed in [3]. Models of CIM prices have been proposed in [4] and [5]. A number of papers have focused on CIM bidding strategies, which is also the topic of our paper. In [6], the authors solve the problem As A stochastic Aptimal Aontrol problem Aor A specific continuous time model of price evolution. The authors in [7] propose a trading strategy based on the assumption that a distribution of the price is available for every period. Neither of these approaches accounts for the fact that trades are binary decisions. This aspect of the problem is accounted for in [8], where the authors propose a simple trading strategy which they test against real data. The goal of the proposed trading strategy As Ao select Aids Ahat Are Axpected Ao Ae profitable, Aelative Ao Ahe predicted Ambalance price.

Our contributions in this paper are the following: (i) We cast the CIM trading problem as a single-stage Markov Decision Process (MDP). (ii) In order to solve this Markov decision process, we employ policy function approximation. We focus on a threshold policy, according to which we define a sell threshold above which we accept buy bids, and a buy threshold below which we accept sell bids. These thresholds can be easily optimized in the case of a one-dimensional problem, such as the one presented in this paper. In order to be able to cope with more complicated settings which cannot be cast as single-stage MDPs (e.g. CIM trading with a storage asset, or accounting for impact of agents' exposure on imbalance prices), we present a reinforcement learning algorithm [9] and [10] which can be used for optimizing multiple threshold parameters (e.g. time-varying thresholds). (iii) We analyze the optimal threshold under the specific assumption that the intraday and real-time price are drawn from a bivariate Gaussian distribution. (iv) We introduce risk aversion in our framework by employing a concave utility function, and provide the analytical solution to this problem for a particular class of utility functions.

II. Description of German CIM

The different German short-term electricity markets are presented in figure 1. The CIM commences after the conclusion of the day-ahead market and the intraday auction. It opens for hourly products at 3 PM [2] on the day before (D-1) and for quarterly products at 4PM on D-1. The market closes 30 minutes before delivery [2]. Finally, after delivery, imbalances are settled at the imbalance price, which we refer to in this paper as the real-time price. Our contribution is focused on the CIM.

In the German CIM, bids appear at random moments in time, and can be selected by any market participant immediately after having been introduced. In this paper we are interested in determining which bids should be selected by a trader. We assume that the trader does not own any physical assets. Thus, if the trader is short (resp. long) in the intraday market, the trader has to buy (resp. sell) back this position at the real-time price. In Germany, the imbalance settlement system is based on a single price, which means that both over- and under-production are settled at the same imbalance price. The bids are characterized by a delivery period (hour or quarter within an hour), a type (buy or sell), a selling or buying price (in €/MWh) and a quantity (in MWh). For the sake of clarity, we neglect bids that are linked between multiple delivery periods and transaction costs.

III. The CIM trading problem as an MDP

We start by presenting our assumptions, and their implications for the problem. We then explain how the CIM trading

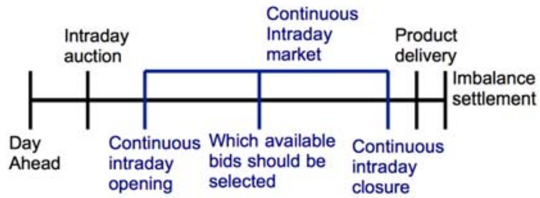

Fig. 1: The sequence of operations in German short-term electricity markets

problem can be cast as a single-stage MDP. After introducing a parametrized policy based on price thresholds, we compute an analytical solution to the problem. We finally examine the properties of the solution for the specific case of a bivariate Gaussian distribution.

A. Assumptions

In what follows, we assume the following:

1) **Traders do not own any physical asset.** This assumption implies that trades taking place in the CIM need to be settled in the real-time market. Thus, the problem is purely financial. Therefore, it can be solved independently for different delivery times. Indeed, there is no reason why the way we trade for bids with delivery time between 3 and 4 PM influences the way we trade for bids with delivery time between 4 and 5 PM.
2) **The cumulative position of a trader does not influence the real-time price.**

These two assumptions imply that each bid can be considered independently from every other bid.

B. MDP Formulation

Given the preceding assumptions, we can cast the CIM trading problem as a one-step MDP model. In this text, we will only present the development for the decision of whether or not to accept buy bids, the reasoning is identical for deciding whether or not to accept sell bids. In order to properly define a one-step MDP, we commence by defining the state variables, the decision variables and the objective function of the problem.

- The state of the problem is the intraday price $S = \{p^{\mathrm{ID}}\}$.
- The set of actions is $A = \{a\}$, where a is a binary variable equal to 1 if a bid is accepted, and 0 otherwise.
- Finally, we can define the objective function by

$$\mathcal{R}(p^{\mathrm{ID}}, a) = \mathbb{E}[p^{\mathrm{ID}} - p^{\mathrm{RT}} | p^{\mathrm{ID}}] a$$

We are interested in a policy $\pi^\star$ which is the solution of

$$\max_{\pi \in \Pi} J(\pi)$$

where $J(\pi) = \mathbb{E}\left[\mathcal{R}(S, A^\pi(S))\right]$ is the likelihood ratio with $A^\pi(S)$ indicating the actions that are selected by following policy π.

Optimizing over the complete space Π is intractable. Instead, we employ policy function approximation. In policy function approximation, we seek a function $\pi(s, a; \theta)$ which is parametrized over a vector of the parameters θ, and assigns a probability for action a if the state is s which is equal to:

$$\pi(s, a; \theta) = \mathbb{P}[A = a | S = s; \theta]$$

We are specifically interested in a threshold policy, which is defined by:

$$a = 1 \text{ if } \theta \leq p_{\mathrm{ID}}.$$

This parametric policy implies that we accept a buy bid if the price is higher than our threshold θ.

C. Analytical solution

If we restrict ourselves to threshold policies and we assume that we have access to the joint distribution of the intraday price and real-time price, it is possible to rewrite the problem as follows:

$$\max_{\theta} \left(\mathbb{E}\left[p^{\mathrm{ID}} | p^{\mathrm{ID}} \geq \theta\right] - \mathbb{E}\left[p^{\mathrm{RT}} | p^{\mathrm{ID}} \geq \theta\right] \right) \cdot (1 - F_{p^{\mathrm{ID}}}(\theta))$$

where p^{ID} is the intraday price, p^{RT} is the real-time price, and $F_{p^{\mathrm{ID}}}$ is the cumulative distribution of the intraday price.

The trade-off in selecting θ is the following: a large value of θ results in keeping the best-priced offers, however only a few offers are accepted. On the other hand, choosing a low value for θ implies that more offers are accepted, however some are potentially less interesting. We can compute an analytical solution to the problem, by finding the point at which the gradient vanishes. [1] We obtain:

$$\theta^* = \mathbb{E}[p^{\mathrm{RT}} | p^{\mathrm{ID}} = \theta^*].$$

Thus, the threshold corresponding to an extreme point is such that the revenues that we earn from selling to a buyer in the intraday market are exactly breaking even with the costs of balancing our position in real time.

After computing the first derivative of the payoff with respect to θ, we use the second derivative in order to determine if the point that we find is a minimum, a maximum, or a saddle point. The detailed derivation is developed in the extended version of the paper.

D. Bivariate Gaussian case

In order to gain intuition about the optimal treshold policy, in this section we solve the problem for the case where the prices are distributed according to a bivariate Gaussian distribution. Assuming that (p^{RT}, p^{ID}) are bivariate normal, we can write [11]:

$$\mathbb{E}[p^{\mathrm{RT}} | p^{\mathrm{ID}}] = \mu_{p^{\mathrm{RT}}} + \rho \frac{\sigma_{p^{\mathrm{RT}}}}{\sigma_{p^{\mathrm{ID}}}} (p^{\mathrm{ID}} - \mu_{p^{\mathrm{ID}}}).$$

Using this conditional expectation expression, we can compute the extremum $\theta^\star$:

$$\theta^\star = \frac{\mu_{p^{\mathrm{RT}}} - \rho \frac{\sigma_{p^{\mathrm{RT}}}}{\sigma_{p^{\mathrm{ID}}}} \mu_{p^{\mathrm{ID}}}}{1 - \rho \frac{\sigma_{p^{\mathrm{RT}}}}{\sigma_{p^{\mathrm{ID}}}}}$$

[1] The proof is available in the extended version of the paper, which is available in the following link: https://sites.google.com/site/gillesbertrandresearch/publications/eem-paper-2018-extended.

In order to verify whether we have a minimum, a maximum, or a saddle point, we evaluate the second derivative at the extremum:

$$\frac{\partial^2 J(\theta^\star)}{\partial \theta^2} = f_{p^{\text{ID}}}(\theta^\star)\left(\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} - 1\right)$$

We can distinguish 3 cases:

- If $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} > 1$, the extremum is a local minimum.
- If $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} < 1$, the extremum is a local maximum.
- If $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} = 1$, the extremum is a saddle point (see the extended version of the paper).

Figure 2 illustrates these results graphically. In the figure, we fix the parameters of the bivariate distribution, and compute the payoff as a function of θ for different values of $\sigma_{p^{\text{ID}}}$. Several observations can be made:

- The payoff for $\theta = -\infty$ is equal to $\mu_{p^{\text{ID}}} - \mu_{p^{\text{RT}}} = 1$. Indeed, at $\theta = -\infty$, $\mathbb{E}[p^{\text{RT}} | p^{\text{ID}} > \theta]$ is equal to $\mathbb{E}[p^{\text{RT}}]$.
- For $\theta = +\infty$, the payoff is 0. If $\theta = +\infty$, all trades are refused, which explains the zero value of the payoff.
- For $\sigma_{p^{\text{ID}}} = 0.125$ and $\sigma_{p^{\text{ID}}} = 0.25$, $\theta^\star$ is a local minimum, because $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} > 1$. For $\sigma_{p^{\text{ID}}} = 0.5$, $\theta^\star$ is a saddle point because $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} = 1$. For $\sigma_{p^{\text{ID}}} = 1$ and $\sigma_{p^{\text{ID}}} = 2$, $\theta^\star$ is a local maximum, because $\rho \frac{\sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}} < 1$.

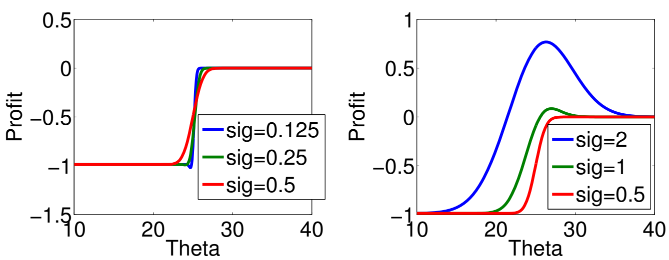

Fig. 2: Profit for different $\sigma_{p^{\text{ID}}}$ with respect to theta for $\mu_{p^{\text{ID}}} = 25$, $\mu_{p^{\text{RT}}} = 26$, $\sigma_{p^{\text{RT}}} = 2$ and $\rho = 0.5$.

IV. Risk aversion

The assumption that traders do not influence the real-time price can lead to a very large imbalance position at the optimum. Such an exposure is risky, especially when one factors in the high volatility of the real-time price. One approach towards controlling the risk of the threshold policy is to apply a concave utility function to the payoff. The introduction of risk aversion implies a coupling between bids, since risk can be mitigated by accepting multiple bids. In order not to obscure the analysis by introducing multiple bids, we will continue our analysis under the assumption that only a single bid is being traded. In this analytically tractable case, we will demonstrate that bid-ask spread can be induced by risk aversion. However, we acknowledge that the consideration of a single bid under risk aversion is a worst-case analysis in terms of diversification, and that a complete analysis would require the simultaneous consideration of multiple bids. We leave this extension to future work.

In this paper, we will investigate the exponential utility function which can be expressed by:

$$U(x) = -e^{(-bx)},$$

where b is a positive parameter. Larger values of b correspond to greater risk aversion of the agent. In what follows, we compute an analytical solution to the MDP.

A. Solution in the general case

The expected utility of the payoff is given by:

$$J(\theta) = \mathbb{E}[-e^{-b(p^{\text{ID}} - p^{\text{RT}})} | p^{\text{ID}} \geq \theta](1 - F_{p^{\text{ID}}}(\theta)) + (-1)F_{p^{\text{ID}}}(\theta)$$

The second term is due to the fact that, if the trader refuses to trade, the utility is equal to -1. The point at which this gradient vanishes is given by the following expression (see the extended version of the paper):

$$\theta^\star = \frac{1}{b}\log(\mathbb{E}[e^{bp^{\text{RT}}} | p^{\text{ID}} = \theta^\star])$$

The consequence of the risk aversion of the producer will be the increase of $\theta^\star$, relative to the risk-neutral case. This follows from Jensen's inequality, $\exp(\mathbb{E}[X]) \leq \mathbb{E}[\exp(X)]$.

B. Solution in the Gaussian case

We compute the extremum in the case of bivariate Gaussian distribution (see the extended version):

$$\theta^\star = \frac{\mu_{p^{\text{RT}}} - \frac{\rho \sigma_{p^{\text{RT}}} \mu_{p^{\text{ID}}}}{\sigma_{p^{\text{ID}}}} + \frac{(1-\rho^2) b \sigma^2_{p^{\text{RT}}}}{2}}{1 - (\frac{\rho \sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}})}$$

Relative to the case without risk aversion, an additional term appears in the numerator, which increases the optimal threshold in proportion to the variance of the real-time price. Note that if the correlation ρ is equal to 1, then this additional term is equal to 0. The intuition behind this is that, when $\rho = 1$, there is no uncertainty regarding the real-time price when the intraday price has been revealed.

In the extended version, we compute the second derivative at the optimum threshold. We get:

$$\frac{\partial^2 J(\theta^*)}{\partial \theta^2} = f_{p^{\text{ID}}}(\theta^\star)\left(-b + \frac{b \rho \sigma_{p^{\text{RT}}}}{\sigma_{p^{\text{ID}}}}\right) \cdot \exp\Big(-b\theta^\star + b\mu_{p^{\text{RT}}} + \frac{b\rho\sigma_{p^{\text{RT}}}(\theta^\star - \mu_{p^{\text{ID}}})}{\sigma_{p^{\text{ID}}}} + \frac{b^2(1-\rho^2)\sigma^2_{p^{\text{RT}}}}{2}\Big)$$

The different cases can be defined as in the case without risk aversion, since the same expression as before is multiplied by a positive function.

C. Influence of b

In figure 3 we present the payoff for different values of the parameter b. As expected, the greater the risk aversion parameter b, the greater the optimal threshold θ^* that we obtain. It can be seen that, if $b = 0$, then we arrive to the same result as in the case without risk aversion.

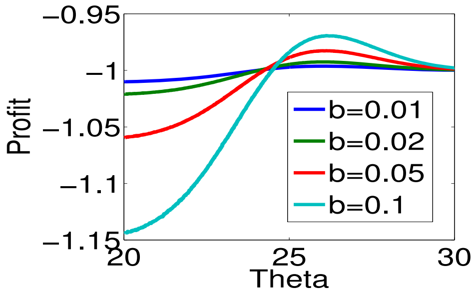

Fig. 3: Payoff for different values of b for the bivariate normal case with $\mu_{p^{\mathrm{ID}}} = 25$, $\mu_{p^{\mathrm{RT}}} = 26$, $\sigma_{p^{\mathrm{ID}}} = 2$, $\sigma_{p^{\mathrm{RT}}} = 2$ and $\rho = 0$.

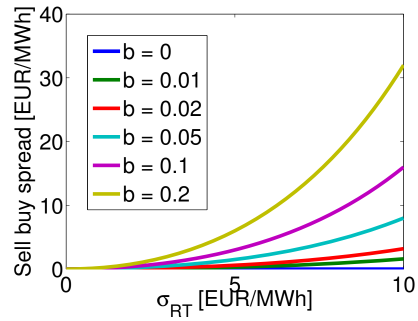

Fig. 4: Spread between the threshold for buy and sell bids for different choices of b, with $\mu_{p^{\mathrm{ID}}} = 25$, $\mu_{p^{\mathrm{RT}}} = 26$, $\sigma_{p^{\mathrm{ID}}} = 2$.

V. Relative value of buy and sell thresholds

In this section we analyze the relationship between the optimal thresholds for buy and sell bids.

In the case without risk aversion, the optimal threshold for accepting sell orders is given in the general case by the condition $\theta^\star = \mathbb{E}[p^{\mathrm{RT}}|p^{\mathrm{ID}} = \theta^\star]$. We observe that this is identical to the optimal buy threshold[2].

For the case with risk aversion, the extremum of the payoff function is given by

$$\theta^* = -\frac{1}{b}\log\left(\mathbb{E}[e^{-bp^{\mathrm{RT}}}|p^{\mathrm{ID}} = \theta^*]\right)$$

For the buy threshold, the consequence of the risk aversion of the producer will be the decrease of $\theta^\star$, relative to the risk-neutral case. This follows again from Jensen's inequality. In the particular case of a bivariate Gaussian, we get

$$\theta^\star = \frac{\mu_{p^{\mathrm{RT}}} - \frac{\rho\sigma_{p^{\mathrm{RT}}}\mu_{p^{\mathrm{ID}}}}{\sigma_{p^{\mathrm{ID}}}} - \frac{(1-\rho^2)b\sigma^2_{p^{\mathrm{RT}}}}{2}}{1 - (\frac{\rho\sigma_{p^{\mathrm{RT}}}}{\sigma_{p^{\mathrm{ID}}}})}$$

The only term that changes compared to the case where we are determining thresholds for accepting buy orders is that the third term in the numerator has the opposite sign. As the intuition suggests, the introduction of risk aversion results in a bid-ask spread. In order to illustrate this result, we present in figure 4 the evolution of the spread between the sell and buy threshold with respect to the standard deviation of the real-time price for different values of b. We observe that, for small variance, the spread is negligible. When the standard deviation increases, the gap between the thresholds also increases. One further observes that the spread increases as b increases.

VI. Learning the optimal threshold

In order to determine the optimal threshold θ in more complicated settings, it is possible to employ the REINFORCE algorithm, as defined in [9] and [10]. The goal of the algorithm is to determine the threshold value θ which maximizes the payoff on the basis of episodes, or epochs, of learning. An episode in the context of our problem is simply the arrival of a bid in the CIM and the realization of the real-time price. The learning algorithm is implemented as follows:

- Initialize θ
- for each episode $\{p^{\mathrm{ID}}, a, r\} \sim \pi(p^{\mathrm{ID}}, a; \theta)$

$$\theta = \theta + \alpha\nabla_\theta \log(\pi(p^{\mathrm{ID}}, a; \theta))r$$

end for

This algorithm requires the derivative of the policy π with respect to the parameter vector θ. The problem is that if one employs a deterministic policy, such as the one described before, we get a non-differentiable policy with respect to the parameter θ. In order to solve this problem, we employ a randomized policy as shown in figure 5. The idea is that our parameter θ is the mean of a Gaussian distribution. We accept a buy bid with price p^{ID} with probability $F_\theta(p^{\mathrm{ID}})$, it is the green part of the distribution. We refuse a buy bid with price p^{ID} with probability $1 - F_\theta(p^{\mathrm{ID}})$, it is the red part of the distribution. This way of modelling the problem matches the intuition that the higher the intraday price is, the bigger chance we have to accept it.

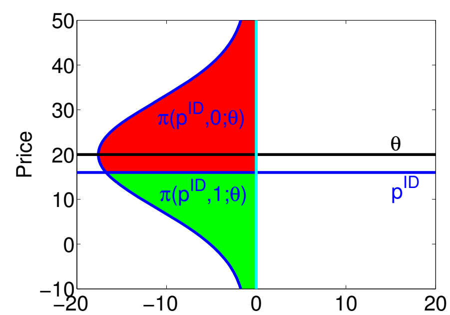

Fig. 5: Stochastic threshold illustration for a buy bid.

Let us develop the mathematical expression of the threshold defined previously (Recall that $\pi(s, a; \theta) = \Pr\{a_t = a|s_t = s, \theta\}$).

$$\pi(p^{\mathrm{ID}}, 0; \theta) = 1 - F_\theta(p^{\mathrm{ID}})$$
$$\pi(p^{\mathrm{ID}}, 1; \theta) = F_\theta(p^{\mathrm{ID}})$$

[2]Nevertheless, when we apply the REINFORCE algorithm against data which is, by construction, characterized by a bid-ask spread (because all remaining orders in the CIM are orders which have not been matched), we can expect a spread between the thresholds for selling and buying power.

In the algorithm, we also need the derivative of the policy with respect to the parameter, we get

$$\frac{\partial \pi(p^{\mathrm{ID}}, 0; \theta)}{\partial \theta} = f_\theta(p^{\mathrm{ID}})$$
$$\frac{\partial \pi(p^{\mathrm{ID}}, 1; \theta)}{\partial \theta} = -f_\theta(p^{\mathrm{ID}})$$

VII. Case study

In this section, we present results for learning the optimal threshold for accepting buy bids. We first demonstrate the efficiency of the REINFORCE algorithm for the case where the price data is drawn from a bivariate Gaussian distribution. Then we apply this algorithm on real German data.

Figure 6 demonstrates the convergence of the optimal threshold according to the REINFORCE algorithm. The optimal threshold converges to its theoretically optimal value for the case where the price data follows a bivariate normal distribution.

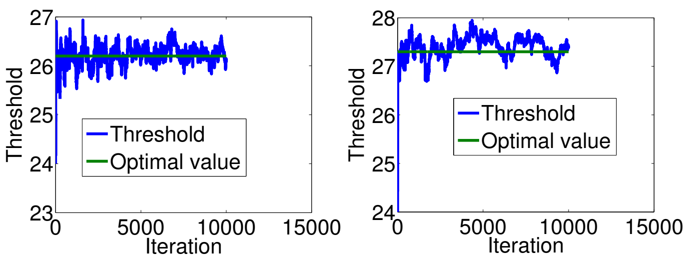

Fig. 6: Convergence of the threshold according to the REINFORCE algorithm for the case of bivariate normal price data, with $\rho = 0$ (left) and $\rho = 0.5$ (right) and $\mu_{p^{\mathrm{ID}}} = 25$, $\mu_{p^{\mathrm{RT}}} = 26$, $\sigma_{p^{\mathrm{ID}}} = 2$, $\sigma_{p^{\mathrm{RT}}} = 2$, $b = 0.1$.

In order to test the performance of the REINFORCE algorithm against German CIM, we use the first 4000 hours data of 2015. We build seven blocks of data to learn and test. The i^{th} block of the learning set contains data between hour $500\cdot(i-1)+1$ and $500\cdot i$. Its corresponding test block contains data between hours $500 \cdot i + 1$ and $500 \cdot (i + 1)$. In order to make a fair analysis, we present the results for the 4^{th} most profitable block (median scenario). In figure 7, we show the cumulative profits. At the end of the period, we get a profit of 20 million euros. In table I, we present the optimal threshold $\theta^\star$ for different values of the risk aversion parameter b ($b = 0$ implies risk neutrality).

b	0	0.0002	0.0005	0.001	0.002
$\theta^\star$ (€/MWh)	-27.4	-27.4	-27.4	-14	121.2

TABLE I: Optimal threshold $\theta^\star$ for different values of the risk aversion parameter b.

VIII. Conclusions and perspectives

In this paper we analyze the problem of CIM trading using a threshold policy. We model the problem as a one-stage MDP if we assume that traders' positions in the intraday market do not influence the real-time price. We use a policy function approximation which is parametized by a price threshold

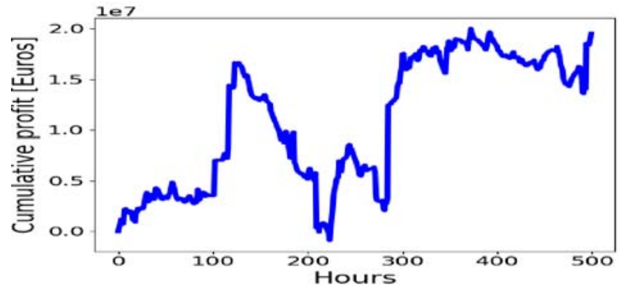

Fig. 7: Cumulative profit achieved by the REINFORCE algorithm against German intraday data.

for accepting bids, and we derive the analytical solution to the problem. We then analyze the optimal solution for the specific case of a bivariate Gaussian distribution. We introduce risk aversion through a concave utility function, and we prove analytically that higher risk aversion results in an increasing bid-ask spread for intraday traders, as intuition suggests. Finally, we apply reinforcement learning in order to compute the optimal trading threshold, and we demonstrate the effectiveness of our approach on the German CIM by making an out of sample test, we obtain a profit of 20 million euros in 500 hours. In future work, we are interested in considering the case in which the trader influences the real-time price through its intraday position. In this framework, we face a multistage MDP, for which we can resort to reinforcement learning in order to optimize time-varying trading thresholds.

Acknowledgment

This research was funded by the ENGIE Chair on Energy Economics and Energy Risk Management, by an Electrabel grant on "Modeling the Value of Flexibility at Sub-Hourly Operating Time Scales", and by the Belgian National Science Foundation (FSR-FNRS) through a FRIA grant. The authors are grateful to Gauthier de Maere, Johan Paduart and Guillaume Erbs for their comments and thoughtful discussions.

References

[1] EU. Retrieve from https://ec.europa.eu/clima/policies/strategies/2020_en

[2] EPEXSPOT Retrieve from https://www.epexspot.com/fr/donnees_de_marche/intradaycontinuous

[3] C.Weber, " *Adequate intraday market design to enable the integration of wind energy into the European power systems*" Energy Policy, Volume 38, 3155-3163, 2010.

[4] R. Kiesel and F. Paraschiv, " *Econometric analysis of 15-minute intraday electricity prices*" Energy Economics, Volume 64, 77-90, May 2017.

[5] R. Kiesel. "*Modeling market order arrivals on the intraday power market for deliveries in Germany with Hawkes processes with parametric kernels*", Energy Finance Christmas Workshop, 2017.

[6] R. Aid et al., "*An optimal trading problem in intraday electricity markets*", arXiv:1501.04575, 2015.

[7] S. Braun and R. Hoffmann. "*Intraday Optimization of Pumped Hydro Power Plants in the German Electricity Market*" Energy Procedia vol. 87, pp. 45-52, 2016.

[8] A. Skajaa et al., "*Intraday Trading of Wind Energy*" IEEE Transactions on Power Systems Volume 30, Issue 6, Nov. 2015.

[9] R. Williams, "*Simple statistical gradient-following algorithms for connectionist reinforcement learning*", Machine Learning, vol. 8, no. 23, 1992.

[10] D. Silver, Retrieve from http://www0.cs.ucl.ac.uk/staff/D.Silver/web/Teaching_files/pg.pdf

[11] D. P. Bertsekas and J. N. Tsitsiklis, "Introduction to Probability" 2002. Retrieve from http://athenasc.com/Bivariate-Normal.pdf