Designing Robust Strategies for Continuous Trading in Contemporary Power Markets

Themistoklis G. Diamantopoulos,
Andreas L. Symeonidis, and Anthony C. Chrysopoulos

Electrical and Computer Engineering Dept., Aristotle University of Thessaloniki,
Informatics and Telematics Institute, CERTH, Thessaloniki, Greece
{thdiaman,asymeon,achryso}@issel.ee.auth.gr

Abstract. In contemporary energy markets participants interact with each other via brokers that are responsible for the proper energy flow to and from their clients (usually in the form of long-term or short-term contracts). Power TAC is a realistic simulation of a real-life energy market, aiming towards providing a better understanding and modeling of modern energy markets, while boosting research on innovative trading strategies. Power TAC models brokers as software agents, competing against each other in Double Auction environments, in order to increase their client base and market share. Current work discusses such a broker agent architecture, striving to maximize his own profit. Within the context of our analysis, Double Auction markets are treated as microeconomic systems and, based on state-of-the-art price formation strategies, the following policies are designed: an adaptive price formation policy, a policy for forecasting energy consumption that employs Time Series Analysis primitives, and two shout update policies, a rule-based policy that acts rather hastily, and one based on Fuzzy Logic. The results are quite encouraging and will certainly call for future research.

Keywords: Double Auctions, Trading Agent Competition (TAC), Energy Market, Fuzzy Logic, Time Series Analysis.

1 Introduction

A *Double Auction (DA)* is an auction where multiple sellers and buyers participate, placing *asks/bids* on the product(s) they want to sell/buy, trying to maximize their profit. DAs are a particularly interesting case of dynamics for Computer and Economics scientists, since product prices and demand/response may fluctuate in an unpredictable manner, thus giving room for research.

Power TAC, initially launched in 2011 [8], provides a powerful competition benchmark of applying DAs' theory to real-life problems. It simulates a modern energy market, where producer, consumer and prosumer needs of electrical power are modeled. Competing agents act as brokers, procuring energy from producers and selling it to consumers in order to acquire the maximum possible profit. The profit, however, does not depend only on monetary units, but also on factors

E. David et al. (Eds.): AMEC/TADA 2012, LNBIP 136, pp. 30–44, 2013.

such as the agent's fees or its bias towards renewable energy sources [8]. In the highly dynamic environment of Power TAC agents face the following challenges:

- Create tariffs to attract customers and create a portfolio of producers and consumers of Electric Energy, according to the current state of the Market.
- Determine the amount of energy to be procured based on the prediction of customers' power load consumption, since any imbalance of the power load is penalized.
- Maximize their profit either by attracting new customers, or improving terms of the existing contracts.

Thus, broker agents have to be equipped with innovative trading and decision making strategies, in order to optimize agent performance. To this end, we have designed agent `Mertacor`. The rest of the paper is organized as follows: Section 2 provides a state-of-the-art review on price formation strategies through the prism of DAs, whereas Section 3 discusses the key modeling elements of Power TAC. The `Mertacor` architecture is presented in Section 4, while Section 5 presents the experiments conducted. Section 6 provides useful conclusive remarks, directions for future research and concludes the paper.

2 Double Auctions

Double Auctions are thoroughly described using *microeconomics* theory, where a product's *price* and *quantity* are determined by the *supply* and *demand* for that product [9,6]. In a given *exchange*, each participant has a *limit price* (aka *private value*) denoting the maximum or minimum price one is willing to offer/accept, being a buyer/seller respectively. In DAs, product price is inversely proportional to the demand and proportional to product supply. The intersection point of the supply and demand curves is called *competitive equilibrium* and is a crucial factor dictating whether shouts are successful or not.

This process is identified as *market clearing* and its timing is actually one of the main criteria that differentiates DAs into various types. Although there are several types, they can typically be divided into two main lines of research: the *Continuous Double Auctions (CDAs)* and the *Clearing House Double Auctions (CH)*. The latter provide an interesting research framework, mainly due to their high market efficiency. Nevertheless, the complexity of CDAs attracts even more researchers interested not only in designing such markets, but also in developing antagonistic price formation policies.

2.1 A Taxonomy of Strategies

Although properly designing a market is an interesting challenge, no market is complete without traders. A trading strategy may have several features depending on the type of the market it is designed for. Strategies may be designed to converge to some equilibrium, mimic human behavior or simply gain profit.

In any case though, the core feature of a strategy is its price formation technique, which is yet an open issue in DAs. This section discusses a taxonomy of price formation techniques, while commenting on the benefits of each approach.

Despite extensive research has been realized on this field, no concrete taxonomy has been established yet, mainly due to the complexity of the techniques involved. The classification proposed within the context of this paper is inspired by the model of Rust et al. [11] and is formed along two axes: *adaptiveness* and *predictivity*.

A strategy is considered *adaptive* if it employs data from previous auction rounds in order to formulate the next shout price. By contrast, a *non-adaptive* strategy is based only on current information, thus achieving simplicity. Strategies employing Reinforcement Learning primitives cannot be classified to either of the above categories, since they usually base decisions upon considering only their own past. Thus, they may be abusively named *self-adaptive*.

An agent may be either *predictive* or *non-predictive*. A non-predictive agent exploits only present and/or past data to form its next shout. In contrast, a predictive agent tries to construct a future state of the market in order to transact in an intelligent manner. To achieve this, it also requires past data, thus predictive strategies are usually adaptive. Finally, non-adaptive and self-adaptive strategies can be regarded as non-predictive, since the former don't make use of past data, while the latter may only construct a model of their own future state, ignoring future states of other agents. Based on the above analysis, the following section discusses the most popular price formation strategies.

2.2 Discussion on State-of-the-Art Price Formation Strategies

Non-adaptive Strategies. The main representative of this category is the Zero *Intelligence (ZI)* strategy, authored by Gode and Sunder [6]. A ZI agent actually submits offers randomly, either freely (*ZI Unconstrained – ZI-U*) or within limit values (*ZI Constrained – ZI-C*). Though ZI may seem as a deprecated strategy, its contribution is substantial in the design of more advanced strategies, since it provides a benchmark against random guessing.

The *KAPLAN* strategy [10], is a "sniping" strategy, in the sense that the agent makes a move towards the end of the current trading period. Although such a simple strategy has had success in the Santa Fe DAs tournament[4], the KAPLAN strategy cannot be applied to competitions like Power TAC, which do not reveal the duration of the rounds.

Adaptive Strategies. Most state-of-the-art strategies are adaptive. Predictive strategies, like *Zero Intelligence Plus (ZIP)*, also model future states [2]. A ZIP agent uses a *learning rule (delta rule)* to update its *profit margin*. It decides whether the latter is increased or decreased based on the type (bid or ask), as well as the success (or not) of the last shout.

An interesting line of research is the one followed by Vytelingum et al. [12], resulting in the *Adaptive Aggressiveness (AA)* strategy. An AA agent initially makes a prediction of the market's equilibrium based on previous shouts, and then computes the price which it should pursue (*target price*). The agent uses a ZIP-like algorithm to determine the trend of its *degree of aggressiveness*, which is used to update the price of the next shout.

On the other hand, non-predictive strategies seem to have a rather straightforward policy on deciding their next move. However, this does not mean that they are less efficient or complex. Strategies like *Gjerstad-Dickhaut (GD)* achieve good results [5]. A GD agent tries to maximize its *expected surplus*, which is defined by the agent's *belief function*. This function is updated according to the success rate of a number of past shouts, i.e. the memory of the agent.

Additionally, He et al. suggest a *Fuzzy Logic based (FL)* strategy [7]. A FL agent computes the *reference price*, i.e. the mean price of the last κ transactions (since the agent has memory), and bases its accepting offer strategy on a comparison to current outstanding shouts. The proximity of the reference price to the shouts is dissolved using fuzzy sets. The agent also has a *learning rate*, which is updated using fuzzy rules, in order to adjust its will to transact according to the frequency of its transactions.

Self-adaptive Strategies. The main advantage of Reinforcement Learning strategies is their utter independence from the other agents' actions, and sometimes even from the market. Although this may seem ineffective, techniques such as *Roth-Erev (RE)* achieve significant performance, especially in sealed DAs, where the amount of information that is given to the agents is limited [3]. The RE agent has been designed to mimic human behavior. The agent's *propensity* of making a move is updated at each round through an *experience function*, which is a reflection of the agent's satisfaction (or disappointment) concerning the considered move. The optimum move is selected using a *choice probability function*.

Finally, the *Q* strategy is an interesting expansion of the Q-learning technique that employs an e-greedy policy [1]. This way, the Q agent either explores the environment or decides using its knowledge up to that point. The agent updates its pricing policy based on which action has been the most profitable in its recent history.

3 The Power TAC Environment

The broker agents' main challenge is the *Tariff Market*. It contains all households, low energy consumers, and small producers. Agents submit their tariffs (asks or bids) to the market, trying to acquire as large market share as possible, while keeping their prices within an affordable level. As far as the customers' contract choices are concerned, they are based on the concept of *tariff utility*. A customer's tariff utility for a given tariff i is given in equation (1):

$$u_i = -(c_u + c_f) \cdot a_{cost} - e_i \cdot a_{energy} - r_i \cdot a_{risk} - I_i \cdot a_{inertia} \tag{1}$$

where the parameters a_{cost}, a_{energy}, a_{risk} and $a_{inertia}$ define the weights given by the customer to costs (either variable (c_u) or fixed (c_f)), energy sources e_i

(e.g. renewable sources), any risks r_i (e.g. dynamic contracts), and the customer's will to remain idle I_i. This way the customer computes the tariff utility of a subset of tariffs offered. Instead of choosing the tariff with the highest utility, the customer selects a tariff in a probabilistic way (see [8]). Agents may also trade energy amongst each other in the *Wholesale Market*, or with large-scale consumers directly, through *requests for quotes (RFQ)*, in order to balance their *portfolio*.

Other entities of Power TAC include the *Distribution Utility (DU)*, which imposes penalty fees when there is an imbalance between broker procured and consumed energy, the *Weather Service*, which provides weather forecast data to brokers, and the *Accounting Service*, which keeps track of all agents' transactions and provides them with portfolio information.

According to Power TAC specifications, all customers are initially bound to a contract with a *default agent*, which is not meant to be competent. Upon initialization, the number of timeslots is determined but not reported to the agents. Brokers can submit tariffs during each timeslot, nevertheless there is a *tariff publication fee* and a maximum number of shouts per agent to avoid market "spamming".

Given that focus of the current work is on the price policy, only variable rate tariffs are taken into account. Thus, (1) is dealt with as (2):

$$u_i = -c_u \cdot a_{cost} \tag{2}$$

given c_u is computed as the mean energy cost for the consumption during k randomly selected past days and a_{cost} is a parameter defining the weight given to that cost.

Thereby, a comprehensive DA environment is created, in the sense that brokers have to buy energy from the producers and sell it to the consumers, aiming to make profit. However, the environment is still too complex to be treated as a simple price formation problem. With respect to the Power TAC challenges discussed, our broker's objectives in the Tariff Market are to:

– Form the price of the next bid (or ask)
– Update (or not) its current shouts or decide to submit new shouts
– Determine the amount of energy to be requested by its producers.

4 Agent Mertacor Design

In accordance with the above mentioned objectives, *Mertacor*, our agent-broker comprises of three modules defining three policies: a price formation policy, a tariff update policy, and an energy prediction policy. These policies provide the necessary input to the core agent mechanism that integrates decisions and defines the final agent strategy, communicated to the Power TAC server.

4.1 Price Formation Policy

Vast amount of information is shared to the agents participating in the Tariff market, thus adaptive strategies should typically be advantageous compared to non-adaptive and self-adaptive ones. The policy developed aspires to exploit any available information as optimally as possible.

Initially, the agent uses data from the last transactions to calculate the *successful price*, which is defined as the mean of all successful transactions of the market and is computed using a *moving average* method. Considering that the agent's memory holds the latest N transactions, the successful price is given by equation (3):

$$SP = \frac{\sum_{i=T-N+1}^{T} w_i s_i p_i}{\sum_{i=T-N+1}^{T} w_i s_i} \tag{3}$$

where w_i is the weight of shout i with price p_i, and s_i denotes whether the offer was successful ($s_i = 1$) or not ($s_i = 0$). Considering transaction T is the most recent, then $w_T = 1$ and all other weights are updated based on equation (4):

$$w_{i-1} = r \cdot w_i \qquad r \in [0, 1] \tag{4}$$

where r determines the importance given to former shouts by the agent. Equations (3) and (4) are valid for both bids and asks.

Mertacor also behaves in a predictive manner. A *risk factor* R determines the agent's eagerness to take risks in order to pursue greater profit. R values range in the interval $[-c, c]$, where c defines the maximum deviation of the final shout price from the successful price. The success rate of the agent's M latest shouts is computed as follows:

$$k = \frac{\sum_{i=1}^{M} acceptedShouts(i)}{M} \tag{5}$$

where $acceptedShouts(i)$ returns 1 or 0 if the shout was successful or not respectively. Equation (5) is then normalized to the interval $[-c, c]$:

$$\hat{k} = c \cdot (2k - 1) \tag{6}$$

Note that the risk factor could be assigned the value of $\hat{k}$. For example, if the success rate of the agent's latest M asks is higher than 50%, then the normalized $\hat{k}$ is positive, meaning that the agent should probably take risks by increasing its shout pricing policy. However, for the sake of experimentation, Mertacor's risk factor follows a Gaussian distribution.

In order to avoid extreme adjustments to the final shout price, the distribution is restricted between intervals, thus it is given by equation (7):

$$g(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{\dfrac{-(x-\mu^2)}{2\sigma^2}} \qquad g(x) \in [-l, l] \tag{7}$$

where the distribution limit l is defined as follows:

$$l = \min\left\{|\hat{k} - c|, |\hat{k} + c|\right\} \tag{8}$$

The distribution's mean is $\hat{k}$ and its standard deviation is given by equation (9):

$$\sigma = \frac{e - l}{3} \qquad e \in (0, \infty) \tag{9}$$

where e is an experimentation parameter, defining the relationship between σ and l. Figure 1 illustrates the influence of e to the height of the distribution as well as its limits. The risk factor is given randomly from the distribution part denoted from the shaded area of Figure 1. The smaller the parameter e, the more likely the agent chooses a value closer to the mean of the distribution, while as the parameter increases, the agent may choose a more risky value.

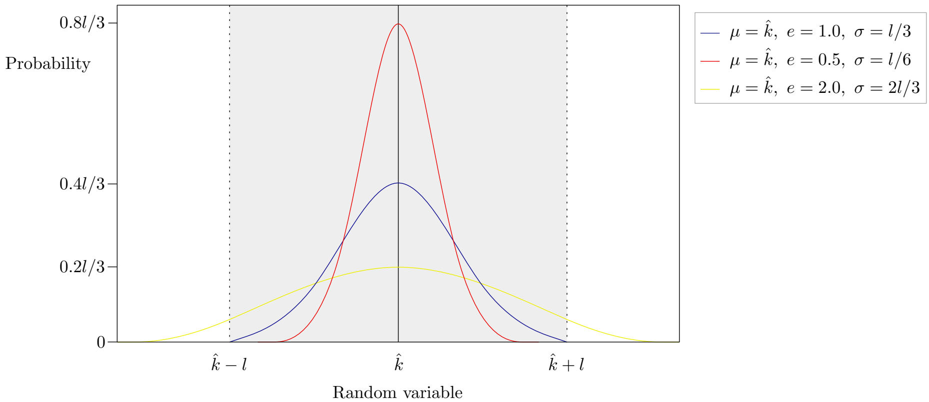

Fig. 1. Experimentation function of the agent

Finally, if s_{ask} is defined to be the agent's ask limit price, the next ask is given by equation (10):

$$ask = \max\left\{SP \cdot (1 + R), s_{ask}\right\} \tag{10}$$

In accordance, let s_{bid} be the agent's bid limit price, the next bid is calculated using equation (11):

$$bid = \min\left\{SP \cdot (1 - R), s_{bid}\right\} \tag{11}$$

The agent's limit values may be considered fixed. In fact, they only change barely to avoid stiff situations, where all agents stick to their limit values.

4.2 Tariff Update Policy

A tariff update policy should deal successfully with submitting tariffs to the Tariff Market, as well as updating already existing ones. Two policies are proposed, both considering the balance of the broker's portfolio and the maximum number of offers permitted.

Basic Update Policy. Striving for an aggressive policy, the agent submits at first a new tariff each round it is given the right to. Thus, it could quickly conquer the market, since the more the tariffs the better the chance of more customers accepting them. However, the policy's spontaneity may lead to ineffective tariffs, since during the first rounds the agent's price formation policy is not likely to have converged to optimal shout values.

Determining whether to update an already existing tariff or not is a matter of the agent's portfolio. Mertacor compares its portfolio balance with what is called the *balance limit*. If Mertacor's portfolio balance surpasses this limit, then it is penalized by the DU. Taking the absolute balance limit into account, the Mertacor's main strategy for asks is depicted in Figure 2a and for bids in Figure 2b. *PBalance* and *BLimit* stand for portfolio balance and balance limit, respectively.

```
if (PBalance > BLimit)
   if (unsuccessful asks exist)
      Update min unsuccessful ask
   else
      Update min ask
else_if (PBalance < - BLimit)
   if (unsuccessful asks exist)
      Update max unsuccessful ask
   else
      Update max ask
```

(a)

```
if (PBalance > BLimit)
   if (unsuccessful bids exist)
      Update min unsuccessful bid
   else
      Update min bid
else_if (PBalance < - BLimit)
   if (unsuccessful bids exist)
      Update max unsuccessful bid
   else
      Update max bids
```

(b)

Fig. 2. Basic update policy update algorithms for (a) asks and (b) bids

This way Mertacor adapts and actually understands any imbalances between the number of producers and the number of consumers contracted to it. For example, when the portfolio balance is greater than the positive balance limit, the agent has an energy deficit, thus alters its minimum shouts to reduce the assigned consumers' and increase the respective producers' market share.

Changing minimum or maximum shouts provides a neat way to adjust Mertacor's contracts to producers and consumers and, as a result, its portfolio balance. However, since the Tariff Market is highly competitive, breaking a contract is rather risky. Thus, Mertacor first attempts to modify any unsuccessful shouts and, if all offers are accepted, only then does it decide to break successful contracts (the less profitable ones).

Fuzzy Logic Update Policy. In order to provide a more comprehensible way of defining the values of the metrics that affect the agent's state, Fuzzy logic primitives have been employed. In fact, two fuzzy sets are defined: one with respect to the number of consumers (producers) under contract and one with respect to the agent's portfolio balance.

The first fuzzy set is defined as the number of consumers (producers) who have a contract with Mertacor. Let B be the number of brokers and C the number of the current consumers (producers) of the agent, the fuzzy set is shown in Figure 3.

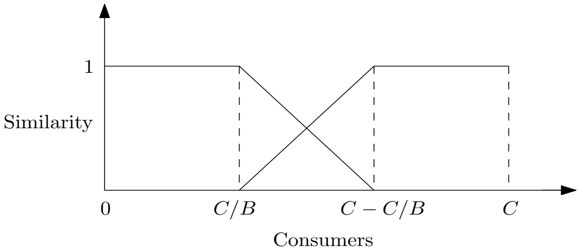

Fig. 3. Fuzzy set for the consumers of the agent. Producers set is defined in an analogous manner.

Dividing the number of brokers by the number of e.g. the agent's consumers provides the number of consumers that the agent should have if they were shared equally among brokers.

The second fuzzy set, describing the agent's portfolio balance, is depicted in Figure 4, where L is the balance limit.

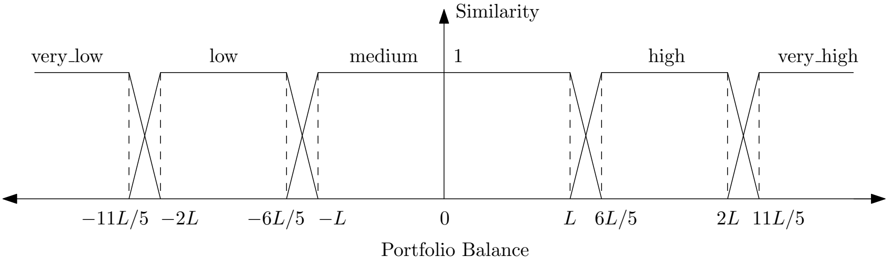

Fig. 4. Fuzzy set for the portfolio balance of the agent

Observing Figure 3 and Figure 4, both fuzzy sets' values are summed to 1 regardless of the independent variable. Thus, the fuzzy sets are regarded in a probabilistic manner, using a uniform random distribution from 0 to 1. For example, as far as the fuzzy set of Figure 3 is concerned, if Mertacor's customers are $C/2B$, they are considered too few. However, if they are $C/2$ then Mertacor has a 50% probability for considering them few and the same probability for believing they are many.

Upon defining the fuzzy sets as well as their use, the policy is analyzed along two main decisions: submitting new tariffs, and updating existing ones. Considering the former, Mertacor submits new offers if customers are few, as long as current offers don't exceed the maximum allowed number of offers for each agent. If consumers are few, Mertacor submits a new ask and if producers are few, Mertacor submits a new bid. The agent submits new shouts only if its market share is not satisfactory, in order to avoid any loss due to tariff submission fees.

As far as tariff updates are concerned, Mertacor updates its maximum ask if the balance is low or very low or if consumers are few. If Mertacor's balance is high or very high, or if producers are few, then Mertacor updates the minimum bid. So far, the policy seems rational, yet not optimal. In order to optimize it, Mertacor also updates the minimum ask if its portfolio balance is very high and it also updates the maximum bid if it is very low. Figures 5a and 5b depict Mertacor's strategy with respect to asks and bids update policy.

```
if ((PBalance is low or very_low)
      or (consumers are few))
   if (unsuccessful asks exist)
      Update max unsuccessful ask
   else
      Update max ask
else_if (PBalance is very_high)
   if (unsuccessful asks exist)
      Update min unsuccessful ask
   else
      Update min ask
```

(a)

```
if ((PBalance is high or very_high)
      or (producers are few))
   if (unsuccessful bids exist)
      Update min unsuccessful bid
   else
      Update min bid
else_if (PBalance is very_low)
   if (unsuccessful bids exist)
      Update max unsuccessful bid
   else
      Update max bid
```

(b)

Fig. 5. Fuzzy Logic update policy update algorithms for (a) asks and (b) bids

Interpreting Metacor's update strategy leads to identifying the relation between the agent's market share and its respective portfolio balance. When that balance is low, Mertacor needs to balance the energy deficit by updating the maximum ask. Thus, Mertacor decreases ask prices in order to increase his consumers' market share. Mertacor acts in a similar manner when consumers are few. However, if the portfolio balance is very low then Mertacor not only considers increasing consumers (Figure 5a) but also decreasing producers (Figure 5b), in order to fix the portfolio imbalance more efficiently. When Mertacor's portfolio balance is high or very high, respective decisions are made.

4.3 Energy Prediction Policy

Through the energy prediction policy Mertacor estimates the amount of energy needed to cover contracted consumers' needs during the forthcoming timeslot. Mertacor receives energy consumption measurements for each timeslot and constructs a time series $\{x_1, x_2, \ldots, x_n\}$, where x_i is the total energy consumption in the market for timeslot i and n is the number of previous timeslots. This way the problem is transformed to predicting the future value of the series. The latter is found using *exponential smoothing*, according to equation (12):

$$\hat{x}_{n+1} = a \cdot x_n + (1 - a) \cdot \hat{x}_n \tag{12}$$

where a defines the importance given to every previous value. Equation (12) is initialized as $\hat{x}_1 = x_1$.

The total energy that the agent is going to need for the next timeslot is computed using equation (13):

$$E_{n+1} = \frac{C_{agent}}{C} \cdot \hat{x}_{n+1} \tag{13}$$

given C_{agent} is the number of the agent's consumers and C is the number of all consumers of the market. Upon checking its portfolio, Mertacor asks for the following amount of energy from every one of its contracted producers:

$$e_{n+1} = \frac{E_{n+1}}{P_{agent}} \tag{14}$$

where P_{agent} is the number of producers in the portfolio.

5 Experiments

Three sets of experiments were conducted in order to identify the parameters that would lead to the most profitable policies for Mertacor. Policies were evaluated against their mean market share of producers/consumers, and the total profit of the agents participating. Each of the experiments was conducted for different configurations of producers and consumers (discussed below) and each experiment was conducted ten times, so the mean values of the derived metrics are presented.

It should also be denoted that all agents tested were assigned the same energy prediction policy. Results are quite intriguing and discussed later on.

5.1 Price Formation Parameters Experiments

From the various parameters of the price formation policy proposed in 4.1, the experimentation parameter is by far the most interesting, since it is crucial for the agent's competing tactics. Other parameters may be given some determined pseudo-optimal values as their impact is not optimum for a certain value. Take agent memory for instance: the moving average method utilized ensures proper weights as long as the memory is not too small. To this end, variables c and r are assigned the values 0.20 and 0.95 respectively.

Several experimentation parameter values were tested (omitted due to space limitations). Table 1 discusses market share and total profit for three different consumer/producer compilations, for experimentation parameter values 0.5, 1.0, and 2.0.

Referring to Table 1, deviations are small, something expected since the other policies encapsulated in the agent are identical, resulting in partial absorption of the influence of the experimentation parameter. However, when the latter is set to 0.5, one may notice that Mertacor makes more total profit than its

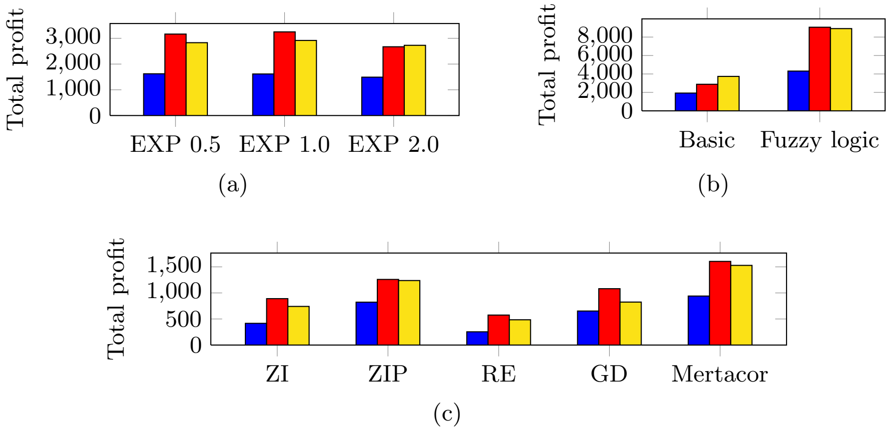

Fig. 6. Graphs showing the total profits for experiments competing (a) different values of the experimentation parameter (b) the two tariff update policies and (c) different price formation policies. The experiments are shown for different settings of producers and consumers, 4 producers & 8 consumers ( ), 4 producers & 16 consumers ( ), 8 producers & 16 consumers ( ).

competitors for small numbers of customers (4 consumers and 8 producers). On the contrary, setting the parameter to 2.0 is more effective with respect to market share, something expected since the agent is more willing to explore its space (see Section 4.1). Taking the golden rule, the value 1.0 is selected as the optimal one for Mertacor, in order to perform well for both profit and market share.

Table 1. Results of price formation parameters experiments

	EXP 0.5			EXP 1.0			EXP 2.0		
	Market share		Total profit	Market share		Total profit	Market share		Total profit
4 cons 8 prod	31,11	32,25	1621,9	31,19	32,65	1614,6	36,89	31,95	1490,7
4 cons 16 prod	29,88	32,32	3164,4	31,98	30,02	3249,8	37,33	34,70	2671,4
8 cons 16 prod	29,31	26,29	2830,3	31,42	31,91	2917,1	37,69	36,26	2728,1

5.2 Tariff Update Experiments

The two tariff update policies introduced in Section 4.2 were tested against each other. Mertacor's experimentation parameter was set to 1.0. The results are depicted in Table 2, where it is evident that the fuzzy logic tariff update policy clearly outperforms the basic scheme, both with respect to profit, as well as market share.

5.3 Price Formation Experiments

The third set of experiments is a comparison of Mertacor's strategy against the four strategies analyzed in Section 2.2: ZI, ZIP, RE, GD. The strategies were adjusted in the Mertacor model, in order to preserve experiment consistency. All five competing agents employ the fuzzy logic tariff update policy and the energy prediction policy discussed in Sections 4.2 and 4.3, respectively. Mertacor's experimentation parameter is set to its optimal value (1.0). Experiment results are shown in Table 3.

Table 2. Results of tariff update experiments

	Basic			Fuzzy logic		
	Market share		Total profit	Market share		Total profit
4 cons 8 prod	25,25	33,95	1926,9	72,94	57,40	4310,1
4 cons 16 prod	19,36	31,55	2880,5	76,79	58,33	9065,0
8 cons 16 prod	23,96	39,6	3734,3	72,88	41,79	8917,2

Table 3. Results of price formation experiments

	ZI			ZIP			RE			GD			Mertacor		
	Market share		Total profit	Market share		Total profit	Market share		Total profit	Market share		Total profit	Market share		Total profit
4 cons 8 prod	11,58	8,60	416,4	20,61	20,15	822,7	4,13	5,85	254,1	18,36	10,55	652,7	44,66	51,35	940,1
4 cons 16 prod	14,39	5,80	890,1	19,36	19,50	1259,0	6,46	7,20	575,1	18,00	11,94	1080,9	41,55	51,23	1604,7
8 cons 16 prod	10,28	12,96	741,0	20,56	19,39	1237,5	6,11	11,11	485,0	15,26	16,15	824,4	46,91	34,15	1528,5

Results are quite encouraging for the Mertacor approach. Additionally, useful conclusions are drawn for the potential of the various approaches on Power TAC in general: ZIP outperforms all other agents, GD outperforms ZI and RE, and RE has disappointing results since its total profit is comparable to the random ZI agent. Thus, one may argue that adaptive strategies (ZIP, GD, Mertacor) seem to achieve better results than non-adaptive (ZI) or self-adaptive (RE) ones. In addition, predictive agents (ZIP, Mertacor) perform better than all others. These conclusions are rather expected since Power TAC's tariff market is an open auction market that provides agents with vast amount of exploitable information.

6 Conclusion

The research challenges provided by competitions such as Power TAC are manyfold. Even when the problem of a free decentralized market is delimited within such a competition, the field is certainly productive when it comes down to designing DA strategies. To do so, one has to explore various research areas, such as Fuzzy Logic or Time Series Analysis.

Based on the experiments conducted, adaptive predictive price formation policies prove efficient in such open-type markets. Learning techniques perform worse, since their sealed-type advantage is depressed. In general, strategies that exploit optimally information seem to be advantageous, like in the case of the fuzzy logic tariff update strategy. In addition, the agent's ability to explore is crucial. This is proven not only by the optimization of the experimentation parameter, but also by the success of predictive strategies.

Further research is encouraged along several aspects of the agent's strategy. As far as the price formation policy is concerned, it would be possible to design different experimentation procedures (e.g. use another probability density function). The strategy could prove even more effective if certain values, such as the experimentation parameter or the limits of the fuzzy logic update policy, were dynamically adjusted to the market, instead of being pre-specified. Finally, with respect to the energy prediction policy, exponential smoothing could be replaced by several time series forecasting models (e.g. an ARMA model), in order to further explore for the optimal production scheme.

References

1. Borissov, N., Wirström, N.: Q-Strategy: A Bidding Strategy for Market-Based Allocation of Grid Services. In: Meersman, R., Tari, Z. (eds.) OTM 2008, Part I. LNCS, vol. 5331, pp. 744–761. Springer, Heidelberg (2008)
2. Cliff, D.: Minimal-intelligence agents for bargaining behaviors in market-based environments (1997)
3. Erev, I., Roth, A.E.: Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. The American Economic Review 88(4), 848–881 (1998)
4. Friedman, D.: The Double Auction Market: Institutions, Theories and Evidence (Santa Fe Institute Studies in the Sciences of Complexity Proceedings). Westview Press (March 1993)
5. Gjerstad, S., Dickhaut, J.: Price Formation in Double Auctions. Games and Economic Behavior 22(1), 1–29 (1998)
6. Gode, D.K., Sunder, S.: Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. The Journal of Political Economy 101(1), 119–137 (1993)
7. He, M., Leung, H.F., Jennings, N.R.: A Fuzzy-Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions. IEEE Transactions on Knowledge and Data Engineering 15(6), 1345–1363 (2003)
8. Ketter, W., Collins, J., Reddy, P.P., Flath, C.M.: The Power Trading Agent Competition. Report Series ERS-2011-011-LIS, Erasmus Research Institute of Management (2011)
9. Mas-Colell, A., Whinston, M.D., Green, J.R.: Microeconomic Theory. Oxford University Press, USA (1995)

10. Rust, J., Miller, J.H., Palmer, R.: Behavior of trading automata in a computerized double auction market. In: The Double Auction Market: Institutions, Theories, and Evidence (Santa Fe Institute Studies in the Sciences of Complexity Proceedings), pp. 155–198 (March 1993)
11. Rust, J., Miller, J.H., Palmer, R.: Characterizing effective trading strategies: Insights from a computerized double auction tournament. Journal of Economic Dynamics and Control 18(1), 61–96 (1994)
12. Vytelingum, P., Cliff, D., Jennings, N.R.: Strategic bidding in continuous double auctions. Artificial Intelligence 172(14), 1700–1729 (2008)