IAMwildCAT: The Winning Strategy for the TAC Market Design Competition

Perukrishnen Vytelingum and **Ioannis A. Vetsikas** and **Bing Shi** and **Nicholas R. Jennings**[1]

Abstract. In this paper we describe the IAMwildCAT agent, designed for the TAC Market Design game which is part of the International Trading Agent Competition. The objective of an agent in this competition is to effectively manage and operate a market that attracts traders to compete for resources in it. This market, in turn, competes against markets operated by other competition entrants and the aim is to maximise the market and profit share of the agent, as well as its transaction success rate. To do this, the agent needs to continually monitor and adapt, in response to the competing marketplaces, the rules it uses to accept offers, clear the market, price the transactions and charge the traders. Given this context, this paper details IAMwildCAT's strategic behaviour and describes the wide techniques we developed to operationalise this. Finally, we empirically analyse our agent in different environments, including the 2007 competition where it ranked first.

1 Introduction

Continuous Double Auctions (CDAs) have traditionally been used in stock markets in order to trade securities and other financial commodities. Their attraction lies in the fact that any trader (buyer or seller) can come into the market, at any point, and place a shout for buying (resp. selling) at some desired price, and a trade will take place almost instantly, if there is a matching offer to sell (resp. buy) at that or a better price. Given this, most of the existing work on CDAs addresses ways of designing effective strategies that maximize a trader's profit. However, there is considerably less literature on the design of market protocols for such auctions in order to promote desirable properties (such as improved efficiency or reduced market volatility). Moreover, in today's globalised economy, stocks are often traded simultaneously in different (competing) markets around the world. Thus, the different markets need to differentiate themselves and appeal to traders to conduct their business under their jurisdiction (e.g. by offering attractive prices for participation and trading).

To rectify this shortcoming, the TAC Market Design Competition (CAT) provides a test-bed for exploring the problem of designing competitive and efficient markets (see [3] for the competition rules). Each CAT game lasts a number of days, and each day consists of a number of trading rounds, which each lasts for a known constant length of time. A number of traders and a number of markets participate (the former are determined by the competition organisers, while the latter are the competition entrants). Each trader is given a finite set of goods to trade and is assigned a private value (also referred to as a limit price) for each good. The difference between this price and the transaction price represents the profit of the agent in the transaction; their total profit in the market is the sum of these transaction profits minus any fees that they incur in participating in the market. The traders use various well-known strategies from the CDA literature: ZI, ZIP, GD, RE (Roth-Erev) [5, 1, 4, 6] and are allowed to register with a different market at the beginning of each day. They also have a memory of the profit they achieved historically in each market such that they are more likely to register to the market where they made the highest profit. Thus, the markets must compete for traders by clearing transactions efficiently and not charging excessive fees.

The different competing markets are represented by specialists, each of which is an agent entered by a separate competitor. These specialists set the rules for their respective market; they determine which shouts are accepted in the market (*quote-accepting rule*), which shouts will be matched for transactions (*clearing rule*) and at what price (*pricing rule*), as well as the fees to charge for various services (*charging policy*).

The score of each agent is a combination of three different metrics: the profit obtained as a percentage of the total profit obtained by all specialists, the market share of the agent (i.e. the percentage of traders who register with the specialist), and the transaction success rate (TSR) (i.e. the percentage of shouts accepted by the market that resulted in a transaction). To be successful, therefore, an agent needs to be competitive in making profit, attracting traders and ensuring that shouts placed in the market result in transactions. While these goals are not necessarily contrary to each other, there are a number of trade-offs to be resolved here. For example, charging larger fees will increase the profit but decrease the market share, while improving the TSR by accepting fewer shouts will result in fewer total transactions and thus less profit both for the specialist and the traders.

In order to design an effective specialist agent, we decided to break the agent down to multiple components, where each one deals with a particular trade-off. Then looking at each component, we designed it in such a way as to balance that trade-off. For example, we designed a clearing rule that allows us to maximize the TSR with a minimal drop in the efficiency of the transactions, and a pricing policy, that manages to extract enough profit without compromising the agent's market share. Similar methods, of breaking down a complex problem into multiple parts and then selecting strategies for and optimizing each one separately, have successfully been used in other complex trading domains [9]. Drawing inspiration from this approach, we also started testing the various individual components using experimental comparisons. The goals of these experiments are two-fold: to determine the best possible agent design, and to examine the behaviour of the market and how it is affected by the different strategies.

Against this background, in this paper we make the following contributions. First, we describe, for the first time, the various policies of our agent. We explain how the various trade-offs guided the design of the agent and how each one was addressed, in order to generate the most competent and successful agent that participated in the

[1] University of Southampton, UK, email: {pv,iv,bs07r,nrj}@ecs.soton.ac.uk

competition. Thus, we designed a number of novel strategies, e.g. clearing shouts, in some rounds, to maximise the number of transactions cleared rather than the profit. Second, we experimentally evaluate the performance of our agent. We compare the efficiency and performance of our agent against that of the other competitors in the competition. Here we show that our agent achieved the best and most stable performance, both in the score and across other metrics (i.e. attracting "good" traders and maintaining a high market efficiency).

This paper is organized as follows. In Section 2, we give a complete description of our agent and all its components. In Section 3, we present the experiments we conducted. Then, we conclude.

2 The IAMwildCAT Strategy

Given this background on the CAT game and its goals, our objective is to design an agent that maximises the scoring function. Specifically, our strategy consists of a set of different market rules and the charging policy (see Figure 1). Each of these policies involves a particular trade-off; in the rest of this section, we detail how we designed the agent in order to resolve each trade-off.

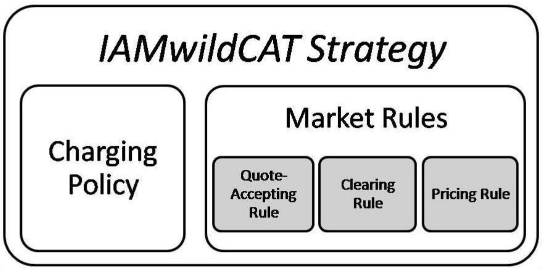

Figure 1. Structure of the IAMwildCAT Strategy

2.1 The Quote-Accepting Rule

We first consider the quote-accepting rule which selects the bids and asks that are accepted into the market (i.e. not all bids submitted by the traders will necessarily be accepted into the marketplace). Such rules are typically employed to speed up the bidding process (e.g. the NYSE quote-accepting rule [10] specifies that any new quote must improve upon the currently outstanding quote), as well as to improve the properties of the auction (e.g. reducing price fluctuations [8]). In the CAT platform, because TSR is a measure of success, it is important to reject the "poor" bids and asks that the market does not expect to clear. Now, we could maximise the TSR by accepting only a few really "good" shouts. However, the fewer shouts that are accepted, the smaller the number of transactions and thus the smaller the profit of both agents and traders; it also makes the market less attractive to traders, which impacts the market share. Thus, we need to select just the right shouts in order to balance this trade-off.

The micro-economic theory of competitive equilibrium states that transaction prices are expected to converge to the competitive equilibrium price p^* where demand meets supply [2]. Thus, we expect the bids (resp. asks) that will be cleared in the market to be at least as high (resp. low) as the competitive equilibrium price. The aim, then, is to accept these bids and asks, rejecting those bids below and those asks above this price. Now, because we can only estimate the equilibrium based on the convergence of transaction prices, we assume some error in our estimation and provide some slack, α_r and α_a, when deciding the minimum bid, $bid_{min} = (1 - \alpha_r)p^* - \alpha_a$ or maximum ask, $ask_{max} = (1 + \alpha_r)p^* + \alpha_a$ to accept. We estimate the competitive equilibrium price using a weighted moving average (giving more weight to more recent transactions because of the expected convergence) that is reset at the end of each trading day.

Furthermore, because traders register with a specialist at the beginning of each trading day, it is impossible to ensure that the same set of traders will remain in the market. Thus we do not expect the equilibrium price to remain constant across trading days. Given this, we reset the moving average of the equilibrium price at the beginning of each day. On the first round, because of the high variance of transaction prices [10] and, hence, the poor estimate of the equilibrium price, we set α_r to be proportional to the variance of the transaction prices (for more slack at the beginning of the trading day).

In more detail, intra-marginal[2] traders are expected to trade earlier than extra-marginal traders (as their better shouts are more likely to be cleared first) such that most of the profitable transactions occur earlier during the trading day, with the extra-marginal and marginal traders left to bid at the end of the day. To avoid marginal bids and asks (that are slightly lower and higher than the equilibrium price respectively) being submitted and risking that they remain uncleared at the end of the trading day, we further constrain our quote-accepting rule. In particular, on the last few trading rounds, our strategy only accepts bids and asks that can be currently cleared. Thus, we minimise the number of uncleared bids and asks, improving our TSR.

2.2 The Clearing Rule

The clearing rule defines when and how to clear the market. There are two parts to this rule. The first is when to clear. One approach is to collect all bids and asks and clear the market at the end of the trading day to maximise profits. However, because traders bid for single units at a time, this approach would imply traders have the opportunity to trade only a single unit (unable to trade the rest of their multi-unit endowment). An alternative approach is to maximise the number of transactions (instead of profits) by a continuous clearing rule whenever a bid or an ask is accepted in the market (e.g. the Continuous Double Auction clearing [10]). Given this, our strategy adopts a rule in between these two approaches, with the market clearing at the end of each round. In this way, we can be almost as efficient as clearing at the end of the day, while allowing the traders to still trade multiple times. By so doing, we get most of the benefits from both approaches without the drawbacks.

The second part is how to match bids. At the end of each round, our agent has a list of shouts to clear. It can try to maximize the number of transactions, by matching "bad" shouts with "good" shouts, but in so doing, it will reduce the efficiency of the market and give less average profit to the traders (which will have an impact on the market share primarily). On the other hand, it can match the shouts efficiently, and maximise profits to the traders, but it will generate less transactions (and TSR). As mentioned earlier, intra-marginal traders are expected to trade earlier than marginal (and extra-marginal) traders such that the amount of profit to be extracted in the market is higher earlier during the trading day, with less profit to be made at the end of the trading day. Thus we chose the following strategy to deal with this trade-off: our agent clears the market for maximum profits at the end of the earlier rounds of the trading day, while, on the following rounds, with less profits to be made in the market, our agent clears to maximise the number of transactions. By so doing, some extra-marginal traders are allowed to transact while increasing the number

[2] An intra-marginal buyer (resp. seller) is expected to trade in the market because of its limit price is higher (resp. lower) than the equilibrium price. The remaining traders are extra-marginal.

of transactions and hence the TSR, at the expense of some profits (though these are generally low at this point).

2.3 The Pricing Rule

The pricing rule determines the price at which a transaction occurs when a buyer and a seller are chosen to transact (by the clearing rule). This price can have any value between the ask and bid prices. Initially, we used primarily *discriminatory k-pricing*[3] with $k = 0.5$; this means that the mean of the ask and bid prices is chosen as the transaction price. In the competition, we used a variation of this policy, called *side-biased pricing*, which varies k, depending on the number of buyers and sellers participating in the market. Specifically, we looked at a window of the latest 10 trading days for the average number of buyers and sellers our agent attracted, and if the difference between the number of buyers and sellers is bigger than 10% of the total number of traders, we adjust k (proportionately to this difference) in order to give more profit to the side which is under represented. We do this in an attempt to attract more of them. However, as we wanted to be somewhat conservative[4], we only allow k to vary in $k \in [0.3, 0.7]$. In Section 3.3, we discuss this issue in more detail and examine the preformance of the two policies.

2.4 The Charging Policy

The charging policy determines the specific charges that are levied from the traders in the system. A *registration fee* is paid by traders in order to register with the market agent at the beginning of the day, irrespective of whether they transact or not. An *information fee* is paid if transaction history information is obtained. A *shout fee* and a *transaction fee* are the amount payed respectively when a shout is placed and when a transaction occurs. The *profit fee* is the percentage of the difference between the accepted shout and the transaction price that is paid by the traders to the market.[5] Before we describe our policy in detail, it is necessary to note the ways in which an agent's charging policy changes its score:

- the score is increased each day by the percentage of the profit that the agent achieved compared to all agents; this means that extracting profit is most efficient for small absolute values of the profit (compared to the total profit extracted by everyone else).
- the market share is decreased by an amount which is relatively proportional to the absolute value of the profit that any agent extracts in total.

These two facts led us to design a charging policy that is mainly trying to maintain a minimum amount of target market share, while at the same time extracting the best possible score from the profit, without compromising the market share. More specifically, we use a *target profit percentage* charging policy, that during each day aims to extract a predetermined profit score. This target score depends on the agent's current market share MS. Specifically, our agent aims to maintain a target market share MS_{target} which takes a value in: $MS_{target} \in [\frac{1}{M}, \frac{1.25}{M}]$, where M is the total number of competing markets. Thus it tries to obtain a market share slightly higher than the average market share that all markets have. We regulate our market share by getting more profit than our opponents when our market share is high, and less when our market share is below our target. We thus distinguish between two states in this strategy:

- If $MS < MS_{target}$, then the market is in *trader attraction mode*[6] and we aim to extract a small profit percentage equal to $P\% = \frac{50\%}{M}$; as this percentage is about half that of the average profit made by other agents, it will lead (all other things being equal) to an increase of market share within some trading days.
- If $MS > MS_{target}$, then the market is in *trader exploitation mode*[7] and we aim to extract a larger profit percentage equal to $P\% = \frac{200\%}{M}$; as this percentage is about twice that of the average profit made by other agents it will lead (all other things being equal) to a reasonable score, but at a cost of some market share loss within the next trading days.

The target share MS_{target} is gradually decreased if trader attraction mode lasts for more than 10 days and is increased for every day that the agent is in trader exploitation mode.

In more detail, let Π, σ, τ and ϕ be, respectively, the total opponent profit, the number of traders in our market, the number of transactions and the average transaction profit (measured as the difference between the ask and bid prices in each transaction), averaged over the last few days. These average values are reasonable expectations for these variables during the following day. Our agent target profit π_{target} is set to $\pi_{target} = P\% \cdot \Pi$. Therefore the average fee paid by each trader must be $\frac{\pi_{target}}{\sigma}$. In trader attraction mode, we set the registration fee equal to 75% of this value, while, in exploitation mode, this is set to 50%. The remaining profit is extracted through the profit fee by dividing the remainder by ϕ. If this value is more than 100%, then we set the profit fee to 100% and gain the remaining profit by additionally setting a transaction fee equal to the remaining profit divided by τ. We don't set an information nor a shout fee. The reason for chosing to extract most of the profit through the registration fee is because all traders, whether intra or extra-marginal, pay this, while only successful (i.e. intra-marginal) traders pay the other two. In this way, we also achieve the effect of attracting the desirable, intra-marginal traders and driving away the undesirable, extra-marginal traders.

A final adjustment to this strategy is made to account for the beginning and end of the game. As market share is more important at the beginning and becomes progressively less so towards the end, we try to build market share at the beginning, by not extracting any profit for a set number of days (set to 80), and increasing the target percentage during the last 100 days of the game, and in particular during the last 40, when the increase becomes quite pronounced.[8]

3 Evaluation

In this section, we analyse the performance of our specialist against other competitors entered in the CAT competition. To this end, we adopt a similar experimental setup[9] as in the competition, with a

[3] The value of k determines the difference of the transaction price from the ask and bid prices.

[4] In the CAT game, because traders consider their entire history of profits and some randomness introduced in the trader's selection of the market to trade in, the effect of giving more profit to one side could be delayed; if we are too aggressive, it might lead us to overshoot our goal of balancing the populations of buyers and sellers and thus cause the behaviour to oscillate.

[5] Note that if this is 100%, then the pricing rule does not matter at all, since all the difference between the ask and bid prices is levied by the market.

[6] To avoid thrashing, we also count the number of trading days since we last switched modes in the strategy; there is a minimum number of days since the last switch before the next switch is allowed.

[7] In fact we use this additional rule before we switch to this mode: we aim to exploit when the total profit made by the opponents drops below its historical average (by a certain discount), as this will allow us to get more score with less penalty to the market share. This discount is being adjusted depending on the number of times that this rule succeeds or fails.

[8] It should be noted that the length of the CAT game, during the competition, was set to 500 days and each day had 10 rounds; these facts were common knowledge and this allowed us to use this end game strategy.

[9] Note that, in our experiments, we used all the available binaries of competition entrants, with the exception of Havana (because of the unavailability of the CPLEX optimisation library they employ) and PSUCAT (because of their unstable implementation).

game[10] running over 500 trading days each lasting 10 trading rounds. The trader population comprises of 180 ZIP traders, 180 RE traders, 20 ZI traders and 20 GD traders, equally split as buyers and sellers. Each trader is endowed with 10 goods to buy or sell at a limit price that is independently drawn from a uniform distribution between 50 and 150 such that the theoretical equilibrium price[11] is 100.

In particular, we first analyse the competition results reported by Nui *et al.* [7] in Subsection 3.1. Then, we analyse in detail the performance of IAMwildCAT. Specifically, we consider the following aspects that Nui *et al.* do not analyse. First, we look at how the number of globally intra-marginal[12] buyers and sellers compares over the trading days (to analyse its effectiveness in attracting "good" traders) in Subsection 3.2. Second, we look at our policy for side-biased pricing in Subsection 3.3 and how it improved our market share and, finally, we look at some more general experiments on the efficiency of our strategy in a homogeneous environment in Subsection 3.4. The purpose of this exercise is to observe its efficiency if all the agents adopt the IAMwildCAT strategy. Note that in figures 2, 3 and 4 we plot only the 5 best strategies for clarity.

3.1 The CAT Competition

Nui *et al.* reported the results of the 2007 CAT competition which was won by IAMwildCAT, with the highest score (at 240.2) outperforming the second placed one by 13% and the third one by 25% [7]. They also empirically evaluated all strategies to identify how they perform in difference cases. They showed that IAMwildCAT had the lowest standard deviation (at 2.8), which suggests consistent behaviour over all the runs. Furthermore, they showed that IAMwildCAT had the highest market share and the highest TSR throughout most of the games. We attribute the former to our strategic choice of maximising market share at the beginning, sacrificing all profits. After 80 trading days, our agent starts charging the traders (see Subsection 2.4) which gradually increases our profit share. We typically expect its market share to decrease (as traders are less profitable in its market). However, by adapting its charging policy effectively, IAMwildCAT does not compromise its market share and, indeed, it is able to increase its profit share while sustaining its market share. Furthermore, our quote-accepting and clearing strategies (see Subsections 2.1 and 2.2) are proved to be very effective, with the TSR increasing from 0.92 at the beginning of the game to over 0.99 after 150 days, outperforming that of all the other agents.

3.2 Intra-Marginal and Extra-marginal Traders

We observe in Figures 2 and 3 that the ratio of intra-marginal traders registered with IAMwildCAT converges to 0.9 (which is considerably higher than that of the other agents). This suggests that our agent successfully incentivises intra-marginal traders to join its market, driving away extra-marginal ones. This is done through setting the fees appropriately (see the charging policy in Subsection 2.4) such that extra-marginal traders, which are not expected to trade, would make negative profit by being charged a registration fee. A market with more intra-marginal traders would imply better bids and asks that can be cleared, which improves our TSR in the process. Now,

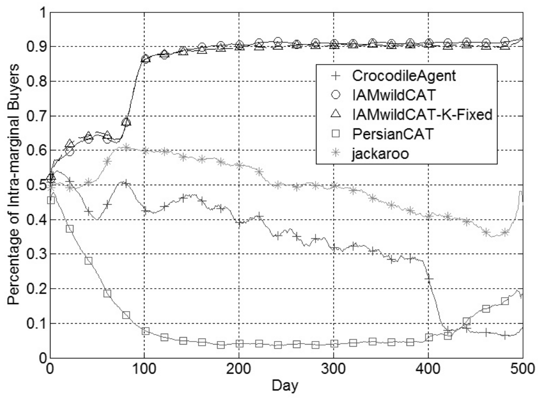

Figure 2. Percentage of intra-marginal buyers.

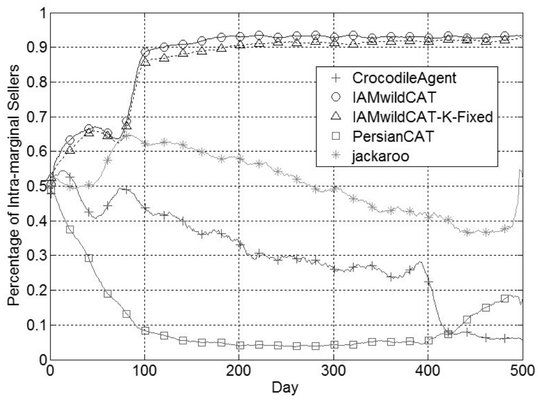

Figure 3. Percentage of intra-marginal sellers.

the intuition behind this ratio capping at around 0.9 is that given the trader's selection strategy, there is a probability of 0.1 that a trader, whether it is intra-marginal or extra-marginal, randomly selects a specialist. Thus, there is always a chance that extra-marginal traders will register with a specialist, such that the ratio can never be 1.

3.3 Discriminatory Versus Side-Biased Pricing

We next evaluate our side-biased pricing policy (where we vary the k parameter); we considered an experiment with 7 different agents, including IAMwildCAT (with this policy) and a modified version of IAMwildCAT, which used the fixed discriminatory k-pricing policy. We believe it is necessary to vary k because intra-marginal traders in a specialist's market might not necessarily be globally intra-marginal. Thus, given our aim to incentivise only intra-marginal traders to join our market, we vary k to give more profit to globally intra-marginal traders than to globally extra-marginal ones. Here, we analyse the effect of side-biased pricing on our strategy.

Now, from Figure 4, we observe that our side-biased pricing policy does increase our ratio of intra-marginal sellers to intra-marginal buyers in the market. However, it introduces a small bias for sellers with more intra-marginal sellers than intra-marginal buyers. It is also interesting to note that IAMwildCAT has a ratio of globally intra-marginal sellers to buyers stable around 1 compared to the hugely varying one of the other agents. This is indeed effective behaviour as a ratio that deviates from 1 implies an equilibrium price that is higher or lower than the theoretical equilibrium in the global market such that some of the profits are distributed to globally extra-marginal

[10] We repeat each game for 15 runs to improve our estimate of performance.

[11] Because the limit prices are drawn from a uniform distribution, the demand and supply curves are expected to be linear, intersecting at 100.

[12] A trader is globally intra-marginal if it is intra-marginal when we consider *all* traders in the system. In our experiments, buyers (resp. sellers) are expected to be intra-marginal if their limit prices are at least higher (resp. lower) than the theoretical equilibrium price at 100.

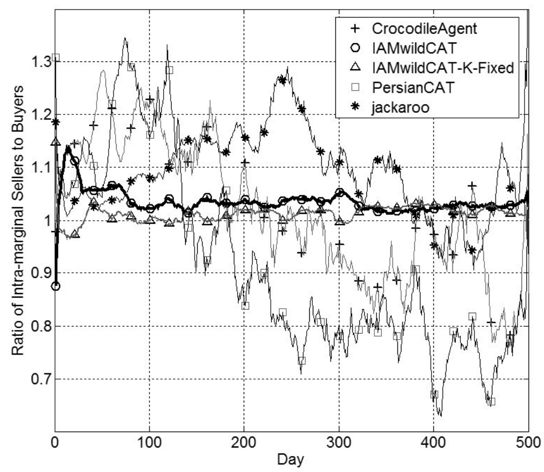

Figure 4. Ratio of intra-marginal buyers to sellers.

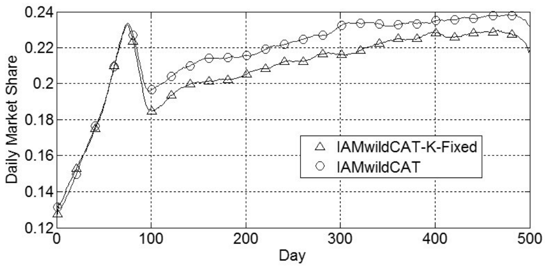

Figure 5. Market share with discriminatory and side-biased pricing.

traders at the expense of globally intra-marginal ones. While the pricing does not affect the specialist's profit share (but rather the distribution of profits among buyers and sellers) or its TSR, we can see from Figure 5 that our side-biased pricing is an improvement over the fixed discriminatory pricing, since it does increase the market share.

3.4 Homogeneous and Heterogeneous Markets

Finally, as per previous evaluation methodologies of double auctions [10, 2], we analyse the global efficiency (and the convergence of the daily market efficiency) of the strategies in both homogeneous and heterogeneous settings. Now, if agents were allowed to select their strategy, they would all chose the most efficient one, i.e. IAMwildCAT, and it would then be very insightful to see how the market efficiency changes if all agents use the same strategy. In particular, in a homogeneous setting, IAMwildCAT does better than in the heterogeneous setting, with a global efficiency of 90.6% (see Figure 6). While PersianCAT has the highest global efficiency (slightly higher than IAMwildCAT at 90.9%), it does poorly in the heterogeneous environment where it scores 128.8, i.e. 47% less than IAMwildCAT. PersianCAT performs well in the homogeneous case because its strategy favours profit-maximisation (sacrificing its TSR) that contributes to the high efficiency. Thus, overall, IAMwildCAT performs well in both a homogeneous (with a high global market efficiency) and a heterogeneous environment (with a high score).

4 Conclusions

This paper details the IAMwildCAT agent, winner of the 2007 TAC Market Design Competition. In particular, we presented the trade-offs present in the design of the agent and gave our strategic rules for quote-accepting, clearing, pricing and charging. We analysed the competition results and, in particular, the IAMwildCAT agent's market share, profit share and transaction success rate compared to the other agents. We then looked at how IAMwildCAT is very successful at incentivising intra-marginal traders to join its market, driving away extra-marginal ones. Furthermore, we examined experimentally the advantage of our side-biased pricing over the standard fixed discriminatory pricing and showed that our agent is able to balance the number of globally intra-marginal buyers and sellers which avoids distributing profits to undesirable, extra-marginal traders. Finally, we analysed the strategies outside the scope of the competition by looking at the market efficiency in homogeneous and heterogeneous environments. As discussed in Subsection 3.4, such insights are particularly important if agents are allowed to change strategies and they all choose the most efficient one. We empirically demonstrated that a market with only IAMwildCAT agents does reasonably well at only 0.3% less than the most efficient one, PersianCAT, while outperforming the heterogeneous market in terms of market efficiency.

As future work, we intend to improve on all the policies we currently have. For example, we intend to improve our charging policy, by better understanding how the different fees individually affect the market share and profit share. This would allow us to experiment with various combinations of strategies (like in [9]) and select the best combination, so as to improve our agent even more. As such strategies are designed to be more and more effective, they will be the foundations for automating real markets in a global economy.

Experiment	Global Efficiency	Convergence Coefficient
6 PersianCATS	90.9%	8.1
6 IAMwildCATS	90.6%	6.2
6 Heterogeneous CATS	88.7%	6.4
6 CorcodileAgents	79.8%	6.1

Figure 6. Efficiency of homogeneous and heterogeneous markets.

ACKNOWLEDGEMENTS

We would like to thank Rajdeep K. Dash who participated in the initial design of IAMwildCAT. Part of this research was undertaken under ALADDIN (joint EPSRC and BAE project EP/C548051/1).

REFERENCES

[1] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Tech Report HPL-97-91, 1997.

[2] D. Friedman and J. Rust, *The Double Auction Market: Institutions, Theories and Evidence*, Addison-Wesley, New York, 1992.

[3] E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps, 'Overview of CAT: A market design competition', Tech Report ULCS-07-006, Dept. of Computer Science, University of Liverpool, Liverpool, UK, (2007).

[4] S. Gjerstad and J. Dickhaut, 'Price formation in double auctions', *Games and Economic Behavior*, **22**, 1–29, (1998).

[5] D. K. Gode and S. Sunder, 'Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality', *Journal of Political Economy*, 101(1):119–137, 1993.

[6] J. Nicolaisen, V. Petrov, and L. Tesfatsion, 'Market power and efficiency in a computational electricity market with discriminatory double-auction pricing', *IEEE Trans: Evolutionary Computation*, **5(5)**, 504–523, (2001).

[7] J. Niu, K. Cai, E. Gerding, P. McBurney, and S. Parsons, 'Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition', in *AAMAS-08*, 1079-1086, (2008).

[8] S. Parsons, J. Niu, K. Cai, and E. Sklar, 'Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement', *AAMAS-06*, 1143–1150, (2006).

[9] I. A. Vetsikas and B. Selman, 'A principled study of the design tradeoffs for autonomous trading agents', in *AAMAS-07*, pp. 473–480, (2003).

[10] P. Vytelingum, *The structure and behaviour of the Continuous Double Auction*, Ph.D. dissertation, School of ECS, Univ. of Southampton, 2006.