

AstonCAT-Plus: An Efficient Specialist for the TAC Market Design Tournament

Meng Chang¹, Minghua He¹ and Xudong Luo²

¹Aston University, Birmingham, United Kingdom

²City University of Hong Kong, Hong Kong SAR, China
{changm2, m.he1}@aston.ac.uk

Abstract

This paper describes the strategies used by AstonCAT-Plus, the post-tournament version of the specialist designed for the TAC Market Design Tournament 2010. It details how AstonCAT-Plus accepts shouts, clears market, sets transaction prices and charges fees. Through empirical evaluation, we show that AstonCAT-Plus not only outperforms AstonCAT (tournament version) significantly but also achieves the second best overall score against some top entrants of the competition. In particular, it achieves the highest allocative efficiency, transaction success rate and average trader profit among all the specialists in our controlled experiments.

1 Introduction

Double Auction (DA) market is a market in which multiple buyers compete to purchase many items that are simultaneously offered for sale by multiple sellers [He *et al.*, 2003]. This mechanism has dominated today's financial instruments exchange for its high allocative efficiency and simplicity in implementation. As economy and technologies evolve, the burgeoning online trading system and electronic market-places have offered traders more freedom than ever to trade their securities across the world. Given this, one DA market has to face competitions from other similar markets running concurrently around the world.

In order to explore automated mechanism design, the International Trading Agent Competition (TAC) market design tournament (also called CAT) simulates the competitive environment of multiple double auction markets. Entrants of the competition called specialists need to design their own strategies for the following policies [Cai *et al.*, 2009]: *clearing policy* deciding how to match traders' offers and when to execute transactions; *pricing policy* calculating transactional prices; *accepting policy* judging what offers can be placed in the market; and *charging policy* determining what are appropriate fees. Another principal entity in CAT is the trader who may be either a buyer or seller willing to exchange goods. Traders are provided by the organiser and equipped with a trading strategy selected from four most studied ones: ZI, ZIP, GD, RE [Gode and Sunder, 1993; Cliff and Bruten, 1997;

Gjerstad and Dickhaut, 1998; Nicolaisen *et al.*, 2001] and a market selection strategy which is mainly based on the history of the trader's profit made with each specialist.

A CAT game lasts a number of days (500 days in CAT-2010). Each day consists of a number of trading rounds, which each lasts for a known constant length of time. The daily evaluation of the specialists is based on three metrics: (1) market share, which is the percentage of the total traders' population registered in the market; (2) profit share, which is the ratio of the daily profit a specialist obtains to the profit of all specialists and (3) transaction success rate (TSR), which is the percentage of the shouts accepted that result in transactions. The daily score of each specialist is the mean value of the above three metrics [Cai *et al.*, 2009].

AstonCAT is a specialist designed for the CAT tournament. Inspired by its soaring improvement in performance on Day 3 of the competition (ranked 5th), we developed post-tournament version called AstonCAT-Plus which significantly outperforms its predecessor and achieves the highest TSR, allocative efficiency and average trader profit among all the specialists in controlled experiments. Our main contributions are: (1) We develop a new and effective equilibrium estimation algorithm reflecting both long-term and short-term market conditions. (2) We introduce intra-day shifting threshold shout accepting strategy. (3) We propose for the first time a clearing strategy which clears market based on profit target. (4) Our hierarchical market-adaptive charging strategy successfully solves the trade-off between maintaining a reasonably high market share and generating profit.

The rest of this paper is organised as follows: Section 2 presents the details of AstonCAT-Plus. Section 3 evaluates it through controlled experiments. Section 4 concludes.

2 AstonCAT-Plus

Figure 1 shows the architecture of AstonCAT-Plus. The strategies corresponds to the four policies mentioned in Section 1. *Shout engine* registers, sorts and classifies accepted shouts. It couples tightly with clearing strategy to determine which bids match which asks. *Auctioneer* acts as a coordinator assembling and passing information requested by other components. *Market client* deals with communication issues with the CAT server. Finally, *market equilibrium estimator* generates estimated current equilibrium price $\hat{p}^*$ which is referred to by clearing, accepting and pricing strategies.




Figure 1: AstonCAT-Plus architecture.

2.1 AstonCAT-Plus Equilibrium Estimation

According to micro-economy theory [Perloff, 1998], in order to efficiently allocate goods and fairly price transactions in a market, it is indispensable to estimate equilibrium price p^* of the market. We estimate p^* through running sliding windows on two independent streams of market information. One is run over the history of transaction prices to find short-term equilibrium price p_s and the other over the averages of daily maximum transacted asks and the minimum transacted bids to find long-term equilibrium price p_l . p_s uses a higher weight for more recent transactions over a short window (typically 5) since p_s is supposed to be reactive to the instant changes of market conditions. Let the last executed transaction be the k th transaction of the game, p_s is calculated as follows,

$$p_s = \sum_{j=k-W_{short}+1}^k p_t^j \omega_j \quad (1)$$

where

$$\omega_j = \frac{0.9^{k-j}}{\sum_{j=k-W_{short}+1}^k 0.9^{k-j}}, \quad (k - W_{short} < j \leq k) \quad (2)$$

p_t^j denotes the price of j th transaction and W_{short} is the size of the sliding window. To obtain p_l , the history of the maximum transacted ask $\bar{a}$ and minimum transacted bid $\underline{b}$ is maintained and the method described in [Honari *et al.*, 2009] is used. Moreover, p_l uses equal weight on every element over a relatively long window (typically 20) such that it reflects a long-term shifting tendency of our equilibrium price. Assuming day z 's "DAYCLOSED" event has just occurred, p_l can be obtained by,

$$p_l = \frac{1}{2W_{long}} \sum_{i=z-W_{long}+1}^z (\bar{a}_i + \underline{b}_i), \quad (z \geq W_{long}) \quad (3)$$

where W_{long} denotes the sliding window size and when $z < W_{long}$, z itself is used as window size. Once p_l value is established, it will be used for the next trading day.

Hence, we can see that p_s contains only a few transactions' information and gets updated dozens of times a day whereas p_l reflects several days' information and gets updated only once a day. By combining p_s and p_l , we have

$$\hat{p}^* = p_s \omega_s + p_l (1 - \omega_s), \quad (4)$$




Figure 2: General structure of AstonCAT-Plus's accepting strategy. Shaded areas are accepting ranges. The arrows pointing to the equilibrium price line indicate directions by which accepting thresholds shift during a day.

where ω_s is the weight of p_s and the weight of p_l is $(1 - \omega_s)$. According to experiments, the best ω_s is chosen to be 0.3.

As shown by Table 4 (see Section 3.4), our transaction prices, which are normally our estimated equilibrium prices according to our pricing strategy, deviate from the theoretical ones by a small margin of 6.205 (6.28% of p^*). Moreover, AstonCAT-Plus achieves the highest allocative efficiency meaning that our traders retain 95.76% of the maximum profit that they can possibly get. Therefore, Formular 4 is effective for estimating market equilibrium.

2.2 Pricing Strategy

Our pricing policy simply sets transaction prices to $\hat{p}^*$ if $\hat{p}^*$ lies inside the bid-ask spread because p^* represents where demand trades off supply. In the case that $\hat{p}^*$ lies outside the spread, transaction price is set to the bid or ask price whichever is closer to $\hat{p}^*$ to prevent negative transactional profit¹. Comparing with side-biased pricing policy [Vyteilingum *et al.*, 2008], ours effectively rewards the intra-marginal side in a transaction rather than the side with less population because short in number does not change an extra-marginal trader to an intra-marginal one in CAT environment.

2.3 Shout Accepting Strategy

Shout accepting strategy decides whether a shout can be placed in our market. Our accepting thresholds are set around $\hat{p}^*$ (see Figure 2). Firstly, due to the fact that the estimated equilibrium cannot be 100% accurate such that some slack can avoid intra-marginal shouts being wrongly rejected. More importantly, even $\hat{p}^*$ is accurate, tolerance is still necessary because intra-marginal traders tend to attempt extra-marginal shouts in order to gain a higher profit.

To this end, we set $\hat{p}^*(1 + \alpha)$ and $\hat{p}^*(1 - \alpha)$ as ask and bid thresholds respectively, where α , named slack rate, determines the degree of openness of the accepting policy. A small α will result in fewer accepted shouts and consequently less transactions than it should be. On the other hand, a large α will result in excess extra-marginal shouts and unfair matching between extra-marginal shouts and intra-marginal ones. Moreover, a too-open policy would reduce TSR due to lots of unmatched shouts [Vyteilingum *et al.*, 2008]. To solve this trade-off, we decrease α with transactions such that the more the transactions, the tighter the thresholds become. At

¹ Transactional profit is defined as profit generated by the difference between bid and ask prices.

the beginning of each day, a large α can encourage shout submissions, which is important for maintaining market share. As transactions are executed, intra-marginal shouts (goods) are consumed which reduces the probability of new shouts being submitted by intra-marginal traders and extra-marginal shouts being matched. Therefore, a decreasing α can effectively block unmatchable shouts from extra-marginal traders and improve TSR. As a result, AstonCAT-Plus is not only attractive to traders (2nd best in market share) but also achieves the highest TSR in heterogeneous games of the controlled experiments (see Section 3.2).

An appropriate initial α is found via experiments between 0.15 and 0.35. Because seller and buyer quantities are usually not exactly symmetric, a bias to the side that is inadequately represented would give the fewer side more freedom and result in more balanced ask and bid profile. Let $\alpha_{0,a}$ denote initial slack rate for sellers and $\alpha_{0,b}$ for buyers respectively. $\alpha_{0,a}$ and $\alpha_{0,b}$ are updated daily as follows:

$$\alpha_{0,a} = \alpha_0 \frac{\frac{(n_b+n_s)}{2} + (\beta-1)n_s}{\beta n_s} \quad (5)$$

$$\alpha_{0,b} = \alpha_0 \frac{\frac{(n_b+n_s)}{2} + (\beta-1)n_b}{\beta n_b} \quad (6)$$

Where $\frac{(n_b+n_s)}{2} + (\beta-1)n_s$ and $\frac{(n_b+n_s)}{2} + (\beta-1)n_b$ are bias ratios. n_s and n_b are the average number of sellers and buyers over last 5 days respectively. $\beta \in [2, 5]$ is used to flatten the result such that the output will not be absurdly far from 1 even if there is a large difference in quantity between buyer and seller. During a day, $\alpha_{0,a}$ and $\alpha_{0,b}$ are deducted by a small value ϵ on every transaction until ask ratio r_a and bid ratio r_b reach their pre-defined limit $l_a = 1.05$ and $l_b = 0.95$. Finally the ask and bid accepting thresholds τ_a and τ_b are,

$$\tau_a = \hat{p}^*(\max(1 + (\alpha_{0,a} - n_t\epsilon), l_a)) \quad (7)$$

$$\tau_b = \hat{p}^*(\min(1 - (\alpha_{0,b} - n_t\epsilon), l_b)) \quad (8)$$

where n_t is the number of transactions happened in a trading day by the time of the calculation.

2.4 Market Clearing Strategy

While shout accepting strategy decides which shouts to be accepted, market clearing strategy decides how to match the accepted shouts and when to convert matches into transactions. AstonCAT-Plus's clearing strategy is a combination of Continuous Double Auction² (CDA) and our innovative profit per transaction (PPT) based clearing mechanism. Our PPT-based clearing scheme is designed to promote traders' profit. The idea behind it is: under certain conditions, we postpone transactions such that matched pairs within the same set of shouts can be reselected to prevent low-profit or extra-marginal transactions. Statistically, more than 70% of new shouts are submitted during first 3 rounds. Therefore, we use PPT-based clearing mechanism during this period to offer intra-marginal traders sufficient opportunities to make their deserved profit. Responding to the change of market

²Trade takes place as soon as there is a matchable pair of bid and ask offers in the market.

1. IF "SHOUTPLACED" event occurs THEN
2. shout engine sorts matched bid-ask pairs
3. IF $round < 3$ THEN /* clear market using PPT-based clearing scheme */
4. $flag = true$
5. calculate average PPT for matched bid-ask pairs
6. IF matched shouts contain extra-marginal ones THEN $flag = false$
7. IF $flag = true$ AND average PPT $> \theta_s$ THEN trigger clearing
8. ELSE IF $flag = false$ AND average PPT $> \theta_l$ THEN trigger clearing
9. ELSE IF $matching\ volume > \frac{n_{trader}}{10n_{market}}$ THEN trigger clearing
10. ELSE trigger clearing /* clear market using CDA */

Figure 3: Pseudocode for AstonCAT-Plus market clearing policy.

conditions after 3 rounds such as reduced trading opportunities, our clearing mechanism switches to CDA in order to focus on boosting transaction quantity as CDA gives intra-marginal traders the greatest chances to exchange their remaining goods within the time left.

Figure 3 illustrates our clearing mechanism. The shout engine implements four-heap algorithm [Wurman *et al.*, 1998]. n_{trader} is the total number of traders and n_{market} is the total number of specialists. Two different minimum PPT limits θ_l and θ_s are employed. Extra-marginal transactions³ do not get executed unless their average transactional profit is larger than θ_l . θ_s is considerably smaller than θ_l such that intra-marginal transactions have the priority to be cleared unless their profit is too small. However, as matching volume increases, we should not hold matches for too long even either PPT target can be reached because quantity of transaction must also be considered in order to maximise traders' total profit in our market. Therefore, statement on line 9 in Figure 3 sets a point where matches are cleared regardless of PPT. As for the values of θ_s and θ_l , we made them adaptive to perceptible traders' private value distribution. Assume that seller (buyer) will not attempt an ask (bid) under (over) his private value, highest attempted bid ($\bar{b}_t$) and lowest attempted ask ($\underline{a}_t$) over a number of days give an indication of trader's private distribution which confine the maximum of transaction profit. Accordingly, θ_s and θ_l are set 2% and 16% of $\bar{b}_t - \underline{a}_t$. In our evaluation, specialists' actual PPTs are compared to show the effects of our clearing strategy (see Section 3.3).

2.5 Charging Strategy

The charging policy selects the type and the amount of the fees that registered traders should pay to obtain market services. AstonCAT-Plus only charges profit fee so that our market entry is free.

Our charging policy adapts fees according to our market status evaluated from several dimensions. As Figure 4 shows, our charging strategy consists of three hierarchical levels of rules. Upper level rules dominate lower ones such that fee updates fired by lower level rules are constraint within a rational range and direction defined by upper levels. Level-1 rules are set based on our current *market share target completion* which is a ratio between AstonCAT-Plus's current trader n_{cur} and trader target $n_{tar} = \frac{\text{number of traders}}{\text{number of markets}}$. Average trader quantity of last 15 days is used as n_{cur} . Level-1 func-

³An ask (bid) over (under) estimated equilibrium is identified as extra-marginal shout and transactions involving extra-marginal shouts are extra-marginal transactions.




Figure 4: Hierarchical market-adaptive stabilized charging strategy.

tions to confine fees to a rational range instead of updating fees directly. Level-2 rules determines the direction of fee modification based on *market trend* which is identified using market trend ratio $r_t = \frac{n_{cur}}{\text{overall mean of daily traders}}$. If $r_t > 1.16$, “up” market trend is identified. If $r_t < 0.92$, “down” market trend is identified. In the case that *market trend* cannot be detected, decision will be made by Level-3 which also determines step size of fee updates. The benchmark for setting Level-3 rules is called *moving trend* identified by moving trend ratio r_v which is weighted n_{cur} with higher weight for more recent days over n_{cur} . $r_v > 1.06$ indicates “up” moving trend, $r_v < 0.97$ indicates “down” moving trend, and $0.97 \leq r_v \leq 1.006$ indicates no moving trend. Combining with decisions made by Level-2 rules, appropriate updating step sizes are selected (see Figure 4). Based on this charging strategy, our fees are rationally-confined, market-adaptive but stabilised against high volatility of market activities. Thus maintaining high market share and making high profit are well balanced.

3 Evaluation

This section analyses the performance of AstonCAT-Plus through controlled experiments in a similar way to [Niu *et al.*, 2010] by which we attempt to relate market dynamics to our adaptive auction rules. Seven specialist agents⁴ (see Table 1) are included in our experiments. Two types of experiments are conducted: heterogeneous games (240 traders) and head-to-head games (120 traders). Traders are uniformly distributed on the four provided trading strategies (ZIP, RE, ZI-C and GD) for both games. In heterogeneous games, AstonCAT-Plus competes with five opponent specialists developed by other institutes. In head-to-head games, AstonCAT-Plus competes with only one specialist each time which includes download agents and AstonCAT.

Our experiment setting is similar to that of CAT-2010 tournament: 500 days and 10 rounds on each day. However, 10 iterations were run instead of 3 in order to obtain statisti-

⁴We mainly include CAT-2010 agents and the latest version of known successful agents. All agents including AstonCAT and AstonCAT-Plus are available for download at <http://www.sics.se/tac/showagents.php>




Figure 5: Score comparison for heterogeneous games.

Specialist Name	Year	Description
Mertacor	2010	Winner of CAT-2010
Jackaroo	2010	Runner-up of CAT-2010
TWBB	2010	5th in CAT-2010
PersianCAT	2008	Winner of CAT-2008
IAMwildCAT	2008	CAT-2008 Final
AstonCAT	2010	CAT-2010 Day 3 version, ranked the 5th
AstonCAT-Plus	2010	CAT-2010 Post-tournament version

Table 1: Specialists used in controlled experiments.

cally significant results. For intra-game analysis, data of a randomly selected representative game iteration are used.

3.1 Overall Performance

In heterogeneous games, AstonCAT-Plus achieves high performance - ranked 2nd according to the overall scores (see Figure 5). AstonCAT-Plus’s average overall score of 10 iterations is only 1.71% lower than Mertacor (winner of CAT-2010) but 29.96% higher than the next best agent Jackaroo (runner-up of CAT-2010). Statistically, AstonCAT-Plus’s lead over other entrants is significant due to a very small p value ($p\text{-value} \ll 0.0001$) in one tail paired t-test against each of them. Specifically, from Figure 5 we can also see that market share is a vital factor for the overall performance as the rank of overall score is monotonically associated with market share. Hence, maximising market share should undoubtedly be the primary target for CAT specialist design. Mertacor and AstonCAT-Plus together dominate the global market as sum of their market share (56.63%) is considerably more than the total of the other four specialists. Large market share also helps AstonCAT-Plus and Mertacor make the most profit although their average profit fee rate is just 4.15% which is far lower than some other agents like PersianCAT (20%) which make much less profit. Apparently, a small amount of fee on a large number of possible transactions after adequate market is gained is the best way to maximise a specialist’s profit.

3.2 Transaction Success Rate

AstonCAT-Plus achieves the highest TSR among all the specialists. As Table 2 shows, AstonCAT-Plus is the only specialist that gained a more than 90% average TSR throughout heterogeneous games. From Table 2, we can also see that AstonCAT-Plus and Jackaroo outperforms Mertacor by 4.1% and 6.2% respectively in terms of TSR. However, AstonCAT-Plus’s TSR is far more stable than Jackaroo as its standard de-




Figure 6: Average trader profit per transaction.




Figure 7: Average transaction (goods traded) per trader.

Specialist	AstonCAT-Plus	IAMwildCAT	Mertacor	PersianCAT	TWBB	Jackaroo
Mean	0.910	0.737	0.857	0.739	0.718	0.892
Stdev	0.0108	0.038	0.0109	0.059	0.098	0.028
Max	0.920	0.816	0.876	0.797	0.882	0.921
Min	0.891	0.693	0.841	0.609	0.597	0.824

Table 2: TSR summary for heterogeneous games

viation (0.0108) is less than half of that of Jackaroo (0.028). We attribute the success to our shifting threshold accepting strategy which ensures that unmatched extra-marginal shouts submitted by extra-marginal traders can be blocked outside our market effectively as trading progresses.

3.3 Average Trader Profit

This subsection analyses the effects of our clearing strategy by comparing average trader profit of each specialist. Since a trader’s average profit is determined by the average profit per transaction (PPT) and the average number of transactions per trader (TPT), we will look into these two factors in turn.

Figure 6 shows AstonCAT-Plus maintains a prominent advantage in terms of PPT. Although AstonCAT-Plus’s average market share falls behind Mertacor’s from day 210, our full-game mean of PPT exceeds Mertacor by 9.94 (21.2% of its mean). According to daily PPTs without averaging, only Mertacor and AstonCAT-Plus’s daily PPT are stable with low standard deviation (3.72 and 4.30 respectively). The others’ swing violently by a least standard deviation of 9.50. According to relative standard deviation to PPT mean (AstonCAT: 7.6%, Mertacor: 7.9%), AstonCAT-Plus achieves the highest PPT with the lowest variance.

High PPT alone does not mean traders make good profit in a market if their average number of transactions per trader (TPT) is small. At the start, everyone’s TPT is around 1.5. But within a very short time, Mertacor and AstonCAT-Plus establish their lead almost simultaneously. Figure 7 shows, at the end of the game, traders with AstonCAT-Plus has traded averagely 1.77 goods out of three total entitlements which is 33.5%, 187.2%, 172.3%, 78.4% more than IAMwildCAT, PersianCAT, TWBB, Jackaroo respectively, but only 1.6% less than Mertacor. Our 2nd best TPT is mainly due to the clearing strategy that encourages intra-marginal transactions using smaller PPT threshold and switches to CDA at the right time.

Ultimately, $\frac{PPT \times TPT}{2}$ gives us average trader profit.

AstonCAT-Plus’s average trader profit (50.38) in this representative game significantly exceeds that of Mertacor, IAMwildCAT, Jackaroo, PersianCAT and TWBB by 19%, 75%, 130%, 303% and 335%, respectively.

3.4 Efficiency and Convergence

Allocative efficiency and convergence coefficient are two essential metrics to identify whether a market is efficient and stable. According to [Cai *et al.*, 2009], the allocative efficiency is defined as the ratio of the trades’ actual profit to the theoretical maximum profit (obtained had all traders traded at the theoretical equilibrium according to microeconomic theory), and the convergence coefficient is defined as standard deviation of transaction prices around daily theoretical equilibrium. According to the heterogeneous game results shown in Table 4, Mertacor and AstonCAT-Plus’s efficiencies are significantly higher with significantly smaller standard deviation than other specialists, which means they are much more efficient and stable markets. AstonCAT-Plus is arguably more efficient than Mertacor not only because both our efficiency mean and maximum are better than Mertacor, but also that we keep our market in high efficiency considerably longer than Mertacor. AstonCAT-Plus achieves “> 95%” allocative efficiency for 395 days (79% of the game length) versus 308 days (62% of the game length) by Mertacor. Mainly we attribute our high efficiency to our accepting and clearing strategies.

3.5 Head-to-Head Games

We have also run head-to-head games between AstonCAT-Plus and each specialist from Table 1 and the results are generally consistent with heterogeneous games (see Table 3). Mertacor is still the only specialist that scores better than AstonCAT-Plus. Rather surprisingly, AstonCAT-Plus’s average overall score is a massive 240% of that of AstonCAT which indicates the newer version successfully overcomes certain vulnerable points of the original one.

In this situation, we are also interested in how traders migrate and finally converge to one of the two markets. Figure 8 shows that market quickly converges to AstonCAT-Plus in the games against PersianCAT, TWBB and IAMwildCAT, gradually converges toward AstonCAT-Plus in the games against AstonCAT and Jackaroo. In the game against Mertacor, AstonCAT-Plus managed to hold an equilibrium of market share where traders do not converge to either market.

Opponent	Overall Score	Market Share	TSR	Efficiency %
IAMwildCAT	0.842 vs 0.402	0.774 vs 0.226	0.862 vs 0.869	93.02 vs 93.41
Mertacor	0.558 vs 0.659	0.462 vs 0.538	0.839 vs 0.824	92.08 vs 94.57
PersianCAT	0.757 vs 0.475	0.693 vs 0.307	0.908 vs 0.789	94.08 vs 70.29
TWBB	0.789 vs 0.470	0.750 vs 0.250	0.884 vs 0.893	94.05 vs 82.77
Jackaroo	0.681 vs 0.574	0.574 vs 0.426	0.836 vs 0.930	94.32 vs 92.85
AstonCAT	0.853 vs 0.356	0.663 vs 0.337	0.898 vs 0.730	94.11 vs 81.39

Table 3: Results of head-to-head games. Each repeated 10 times. First values in each column refer to mean of AstonCAT-Plus and second ones refer to means of the corresponding opponents.

Specialist	AstonCAT-Plus	IAMwildCAT	Mertacor	PersianCAT	TWBB	Jackaroo
Allocative Efficiency %	Mean	95.762	82.744	95.474	71.625	68.301
	Stdev	0.541	2.316	0.239	4.120	8.598
	Max	96.539	87.175	95.800	75.944	83.540
	Min	94.480	79.396	95.116	63.580	57.780
Convergence Coefficient	Mean	6.205	12.502	5.163	8.013	10.308
	Stdev	0.816	1.211	0.722	1.221	0.749
	Max	7.217	14.114	6.061	9.506	11.592
	Min	4.896	10.702	3.943	5.844	8.889

Table 4: Summary of allocative efficiency and convergence coefficient in heterogeneous games.

4 Conclusion

This paper details the strategy used by market specialist AstonCAT-Plus (post-tournament version for the TAC CAT). Specifically, we provide the novel strategies for accepting shouts, pricing transactions, clearing market and setting fees. AstonCAT-Plus is evaluated empirically with AstonCAT (tournament version for CAT-2010) and other top entrants of the competition, in which it performs stably and highly efficiently. Moreover, our mechanism has shown many adaptive features. For example, our clearing threshold adapts to the number of transactions and PPT limits adapt to perceptible traders' private value distributions. We believe some of our original ideas described in the paper can be borrowed by real world DA market designers.

As for future work, we will further improve our shout engine algorithm to enable clearing decision to be made on each individual bid-ask pair rather than the matched shouts bunch. Furthermore, most of our parameters are determined through experiments at the moment. We are going to design an evolutionary approach to learn optimal values for selected parameters through repeated games. By doing so, we believe the reusability of our mechanism can be improved too.

References

- [Cai *et al.*, 2009] K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps. Overview of CAT: A market design competition version 2.0. *Technical Report ULCS-09-005, University of Liverpool, UK*, 2009.
- [Cliff and Bruten, 1997] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. *Tech Report HPL-97-91*, 1997.
- [Gjerstad and Dickhaut, 1998] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22, pages 1–29, 1998.
- [Gode and Sunder, 1993] D. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Mar-




Figure 8: Daily market shares in head-to-head games.

ket as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), pages 119–137, 1993.

- [He *et al.*, 2003] M. He, H. Leung, and N. R. Jennings. A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Transactions on knowledge and data engineering*, 15(6), pages 1345–1363, 2003.
- [Honari *et al.*, 2009] S. Honari, M. Ebadi, A. Foshati, M. Gomrokchi, J. Benatahr, and B. Khosravifar. Price estimation of PersianCAT market equilibrium. *Association for the Advancement of Artificial Intelligence*, 2009.
- [Nicolaisen *et al.*, 2001] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation* 5(5), pages 504–523, 2001.
- [Niu *et al.*, 2010] J. Niu, K. Cai, S. Parsons, P. McBurney, and E. Gerding. What the 2007 tac market design game tells us about effective auction mechanisms. *Autonomous Agents and Multi-Agent Systems. Special Issue on Market-Based Control of Complex Computational Systems*, 2010.
- [Perloff, 1998] J.M. Perloff. *Microeconomics*, chapter 2, pages 29–32. 1998.
- [Vytelingum *et al.*, 2008] P. Vytelingum, I. A. Vetsikas, B. Shi, and N. R. Jennings. IAMwildCAT: The Winning Strategy for the TAC Market Design Competition. *In Proceedings of 18th ECAI*, pages 428–432, 2008.
- [Wurman *et al.*, 1998] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: theory and implementation. *Decision Support Systems* 24(1), pages 17–27, 1998.