

How to Make Specialists NOT Specialised in TAC Market Design Competition? Behaviour-Based Mechanism Design*

Dengji Zhao^{1,2}, Dongmo Zhang¹, and Laurent Perrussel²

¹ Intelligent Systems Laboratory, University of Western Sydney, Australia
`{dzhao,dongmo}@scm.uws.edu.au`

² IRIT, University of Toulouse, France
`laurent.perrussel@univ-tlse1.fr`

Abstract. This paper proposes an approach to design behaviour-based double auction mechanisms that are adaptive to market changes under the Trading Agent Competition Market Design platform. Because of the dynamics of the market environment, it is not feasible to test a mechanism in all kinds of environments. Since the strategies adopted by traders are well classified and studied, we will analyse and utilise the behaviour of traders with each kind of strategy, design specific (trader-dependent) mechanisms for attracting them, and finally integrate these trader-dependent mechanisms to achieve adaptive mechanisms.

Keywords: Adaptive Mechanism Design, Double Auction, Behaviour-based Mechanism, E-Commerce, Trading Agent Competition.

1 Introduction

A double auction market allows multiple buyers and sellers to trade commodities simultaneously. Most modern exchange markets, e.g. the New York Stock Exchange, use double auction mechanisms. In a typical double auction market, buyers submit *bids* (buy orders) to the auctioneer (the market maker) offering the highest prices they are willing to pay for a certain commodity, and sellers submit *asks* (sell orders) to set the lowest prices they can accept for selling the commodity. The auctioneer collects the orders and tries to match them using certain market clearing policies in order to make transactions.

An annual Trading Agent Competition (TAC) Market Design Tournament (CAT Tournament) was established in 2007 to foster research in the design of double auction market mechanisms in a dynamic and competitive environment, particularly mechanisms able to adapt to changes in the environment [1,2]. A CAT tournament consists of a series of games, and each game is a simulation of double auction markets including traders (buyers and sellers) and specialists

* This research was supported by the Australian Research Council through Discovery Project DP0988750.

(market makers). Traders are simulated and provided by the tournament organiser, while each specialist is a double auction market set up and operated by a competitor. Traders dynamically swap between specialists to trade, while specialists compete with each other by attracting traders, executing more transactions and gaining more profit. Therefore, the CAT tournament environment simulates not only the dynamics of traders but also competition among specialists, which renders the market design particularly challenging.

Although certain winning market mechanisms under the TAC competition platform have been published [3,4,5,6], they cannot guarantee that a winning mechanism is also competitive when the environment changes. This also explains why a winning specialist could not win all games in the final in past tournaments. This is further demonstrated by Robinson *et al.* through one post-tournament evaluation [7]. They showed that most specialists are susceptible to environmental changes. *This phenomenon raises the question of how to design a competitive double auction market that is adaptive to environmental changes.*

Central to becoming a winning specialist in the CAT tournament is attracting as many good traders as possible in order to receive more good shouts, generate more transactions and therefore create more profit for both traders and the market maker. This is also true for a real exchange market, as people normally choose a market based on market liquidity and the number of traders in the market. Moreover, there often does not exist a uniform mechanism that is attractive to all kinds of traders, which also explains why different exchange markets use different policies to target different traders in the real world. *Therefore, it is very important for a market maker to fully understand the market environment and target the right customers.* A key way to understanding the market environment is analysing historical market information.

Therefore, in this paper we propose an approach based on traders' behaviour to design competitive mechanisms that are also adaptive to environmental changes. By classifying and utilising traders' behaviour, we first design mechanisms that are competitive in environments with one kind of trader, and then integrate these trader-dependent mechanisms to obtain competitive mechanisms for any complex environment that is not known in advance.

This paper is organised as follows. After a brief introduction to the CAT tournament platform in Sect. 2, we show how to classify traders based on their behaviour in Sect. 3. Section 4 presents a way to utilise traders' behaviour in the design process and shows an experimental example. In Section 5 we introduce a more general extension of this approach, and conclude in Sect. 6 with some suggested directions for future work.

2 Preliminary

This section will introduce the CAT tournament platform, called JCAT [8]. JCAT provides the ability to run CAT games. A CAT game consists of a CAT server and CAT clients including traders (buyers and sellers) and specialists (market makers). The CAT server works as a communication hub between CAT

clients and records all game events and validates requests from traders and specialists. A CAT game lasts a certain number of days, say 500, and each day consists of rounds. Each trading agent is equipped with a specific bidding strategy and can only choose one specialist to trade in each day, while each specialist is a combination of policies. Traders are configured by the competition organiser, and each specialist is set by a competitor.

Each trader is configured with a private value (i.e. its valuation of the goods it will trade), a *market selection strategy* and a *bidding strategy*. The market selection strategy determines a specialist to trade in each day, and the bidding strategy specifies how to make offers. The main market selection strategies used in previous competitions are based on an n-armed bandit problem where daily profits are used as rewards to update the value function. Bidding strategies integrated in JCAT are those that have been extensively studied in the literature, namely ZIC (Zero Intelligence-Constrained [9]), ZIP (Zero Intelligence Plus [10]), GD (Gjerstad Dickhaut [11]), and RE (Roth and Erev [12]).

Each specialist operates one exchange market and designs its own market rules in terms of five components/policies, namely accepting policy, clearing policy, matching policy, pricing policy and charging policy. Accepting policy determines what shouts/orders are acceptable. Clearing policy schedules clearing time during a trading day. Matching policy specifies which ask is matched with which bid for clearing. Pricing policy calculates a transaction price for each match given by matching policy. Charging policy is relatively independent from other policies and determines the charges a specialist imposes on a trading day, e.g. fees for each transaction.

3 Behaviour-Based Trader Classification

Given an unknown environment, the key to understanding it is analysing traders' behaviour. Especially when the strategies adopted by traders can be clearly classified, we want to find out traders' behaviour patterns for different strategies, i.e. the relationship between traders' strategies and their behaviour. Therefore, we can distinguish traders in terms of their behaviour and apply different policies for different traders. In this section, based on JCAT, we introduce how to collect traders' behaviour-related information, define the categories of traders and finally show how to classify traders based on their behaviour.

3.1 Data Acquisition

In JCAT, for each trader i and each specialist s , all specialists can obtain the following trader-related historical information.

- Accepted shouts of i by s .
- Cleared/Matched shouts of i by s .

The above information is also the only information about each trader available for all specialists. The trader of a rejected shout is never revealed to any specialist, even the specialist whom the shout was submitted to. Therefore, the acceptance of a shout cannot depend on the sender's historical information. Given

the above information about each trader, we need to pre-process it depending on what we need for the design process, e.g. the average clearing price for a trader in a specialist during a period of time and a trader's trading time distribution.

3.2 Defining Categories of Trader

Given the perfect equilibrium price p_e^* of a market¹, we classify traders into two different categories, intra-marginal and extra-marginal:

- *Intra-marginal*: A seller (buyer) i with private valuation v_i is intra-marginal if $v_i \leq p_e^*$ ($v_i \geq p_e^*$).
- *Extra-marginal*: Otherwise.

The reason for classifying traders into these two categories is that intra-marginal traders can bring profitable shouts to a market, while extra-marginal traders do not. Therefore, a competitive specialist needs to attract more intra-marginal traders. We can further classify intra-marginal traders in terms of their bidding strategies.

3.3 Category Recognition from Behaviour

We say a trader is attracted by a specialist if the trading time the trader spent in that specialist is much greater than the time it spent in any other specialist. We know that a profit-seeking trader chooses a specialist that has given it the highest profit in some past period. In order to give a trader profit, a specialist has to match its shouts as many as possible with profitable clearing prices. Therefore, intra-marginal traders are more likely to be attracted. Thus, a trader's trading time distribution (i.e. stability) will be the main information to be considered in its category recognition.

Trading Time Distribution. As the main market selection strategy adopted in CAT competitions, *ϵ -greedy selection* determines what is the most profitable specialist for a trader and then selects this specialist with probability of $1 - \epsilon$ and the others randomly with probability ϵ . This selection strategy uses reinforcement learning method based on the profit a trader received from each specialist. ϵ is mostly set to be 0.1 in CAT competitions.

Based on the above market selection strategy, we recognised the following trading time distribution patterns. We say a trader i is more stable if the time (w.r.t. the number of days) that i spent in each market varies significantly, i.e. the standard deviation of the trading time is higher. Generally speaking, intra-marginal traders are much more stable than extra-marginal traders under the same bidding strategy, but the degree of stability varies with bidding strategies.

- *Under the same bidding strategy.* All intra-marginal traders have similar trading time distribution, in other words, intra-marginal traders with valuations

¹ The equilibrium of a market where traders truthfully report their demand and valuations.

far from the perfect market equilibrium are not more stable than those with valuations close to the perfect market equilibrium. Extra-marginal traders with valuations close to the perfect equilibrium are less stable than intra-marginal traders, but they still have preferences between markets. When valuations of extra-marginal traders are far from the perfect market equilibrium, they have no strict preference for any market, i.e. the times spent in each specialist are very close to each other.

- *Degree of stability with different strategies.* Given similar valuations, GD, ZIP and ZIC traders are more stable than RE traders. One reason is that an RE trader uses the profit that it was able to obtain in the most recent trading in a market to adjust (increase) its bidding price, so it will keep increasing its bidding price in a market until finally its shouts cannot be successfully matched, which will cause the trader to move to another market.

Stability vs Intra-marginality. As we have mentioned in the above, most intra-marginal traders are very stable. However, some extra-marginal traders with valuations close to the perfect equilibrium can also be very stable if there are some specialists that have very high probability to match their shouts while others cannot do so. Therefore, a stable trader doesn't need to be intra-marginal. To find out whether or not a stable trader is intra-marginal, we need further information about their behaviour, e.g. bidding prices. If a stable seller's (buyer's) average bidding price is above (under) the equilibrium price, then it maybe not intra-marginal. In general, the selected information should be able to efficiently classify traders into the categories you defined.

4 Behaviour-Based Policy Design

A mechanism in a specialist is a combination of different policies and the relationship between these policies are not completely clear, so searching a competitive combination without restriction under this setting will be computationally intractable. In general, we limited the search space for each policy to certain well-known alternatives that are normally trader-independent. Moreover, there often exist many policy combinations that are competitive under the same market environment, which can be seen from the results in [6]. However, in our approach, since we have gained an understanding of traders' behaviour, we are able to further limit the search space by utilising traders' behaviour. More importantly, we want to further utilise traders' behaviour information to design trader-dependent mechanisms that attract one kind of trader, and integrate those trader-dependent mechanisms to achieve adaptive mechanisms that are attractive to all kinds of traders. In the rest of this section we will define the policies of a specialist by using traders' behaviours and propose a two-step method to search adaptive mechanisms.

4.1 A Search Space of Behaviour-Based Policies

Combined with traders' behaviours, the following policies are adapted from the literature.

Accepting Policy. Once a specialist gets a new shout, it has to first decide whether or not to accept it. If too many extra-marginal shouts are accepted, they will not be matched and therefore the transaction rate will be very low. So why does not a specialist only accept shouts from traders that it wants to attract? Unfortunately, a specialist does not know who is the sender of a shout before the shout is accepted in CAT competitions. Instead some other general market information can be used here, e.g. the equilibrium price of historical shouts received in a market. We will use the equilibrium price of historical shouts to set up a maximum (minimum) acceptable ask (bid) price for each day, as historical equilibrium can approximately distinguish between intra-marginal and extra-marginal shouts.

Given current day t , most recent M historical shouts H_t^M , the maximum acceptable ask price A_t^a and minimum acceptable bid price A_t^b are defined as:

$$\begin{aligned} A_t^a &= E(H_t^M) + \theta^a * F_t^a \\ A_t^b &= E(H_t^M) - \theta^b * F_t^b \end{aligned}$$

where $E(H_t^M)$ is the equilibrium price of H_t , $F_t^a, F_t^b \geq 0$ are relaxations, and $\theta^a, \theta^b \in [0, 1]$ are the relaxation rates. F_t^a and F_t^b are calculated for each day, and θ^a, θ^b are dynamically updated during a day, say, updated after each round.

Matching Policy. The two most used matching policies are *equilibrium matching* and *maximal matching*. Equilibrium matching is used to find the equilibrium price p_e which balances the bids and the asks going to be matched so that all the bids with price $p \geq p_e$ and all the asks with price $p \leq p_e$ are matched [13]. The aim of maximal matching is to maximise the number of transactions/matches by matching high intra-marginal shouts with lower extra-marginal shouts if necessary. The main difference between these two matchings is that maximal matching moves some profit from high intra-marginal traders to lower extra-marginal traders so that lower extra-marginal traders are attracted. Actually maximal matching can also be used for other proposes, e.g. stabilising some high intra-marginal traders, which can be seen in a mechanism for attracting GD traders in Section 4.3. But one disadvantage of maximal matching is that if it moves too much profit from high intra-marginal traders, they will leave the market so that other intra-marginal traders will be affected recursively. At the same time, since equilibrium matching always gives more profit to high intra-marginal traders, some profit seeking traders, like ZIC and RE traders, will keep increasing their profit margin so that their shouts are difficult to match.

Because of the availability of each traders' behaviour information, we will adopt this information for the matching policy. The following are the two additional policies we used in this framework.

1. *Double Equilibrium Matching.* We run two matchings one after another. The first matching is an equilibrium matching based on the bidding price of shouts. The second matching rematches the matched shouts given by the first matching in terms of the average clearing price of each sender's current

Algorithm 4.1. Modified Discriminatory k -pricing Policy

Input: a : ask, b : bid
Output: \hat{p} : clearing price

```

1 begin
2   if  $best(s(a)) = m_i$  and  $best(s(b)) = m_i$  then  $k = 0.5$ ;
3   else if  $best(s(a)) = m_i$  (or  $best(s(b)) = m_i$ ) then
4     | if  $s(b)$  (or  $s(a)$ ) is attractable then  $k = minK$  (or  $k = 1 - minK$ );
5     | else  $k = 0.5$ ;
6   else
7     | if  $s(a)$  is more attractive than  $s(b)$  then  $k = 1 - minK$ ;
8     | else  $k = minK$ ;
9   end
10  if  $p^*(a) \leq p^*(b)$  then  $p_a = \max(p^*(a), p(a)$ ;  $p_b = \min(p^*(b), p(b)$ );
11  else  $p_a = p(a)$ ;  $p_b = p(b)$ ;
12   $\hat{p} = p_a + k * (p_b - p_a)$ ;
13 end
```

best market², called best clearing price. The second matching matches two shouts if the gap between their best clearing prices is very small. This is because their best clearing prices are good enough to attract them and also don't give them too much space to increase their profit margin.

2. *Behaviour-based Maximal Matching.* Maximal matching is guided by traders' behaviour so that extra-marginal shouts are matched only if the senders are those whom we want to attract, i.e. stable traders.

Pricing Policy. Pricing policy will also play a very important role not only in attracting traders but also in stabilising traders. We use a modified *discriminatory k -pricing policy*, where k is dynamically determined for each match according to the two corresponding traders' behaviour. Let $p(x)$ indicate the bidding price of shouts x , $s(x)$ indicate the sender of shout x , $best(t)$ indicate the current best market of trader t , and $p^*(t)$ is the average clearing price for trader t in $best(t)$. Assume the current specialist is m_i , Algorithm 4.1 gives the pseudo-code of the modified pricing policy, where $minK \in [0, 1]$ is what we have to set up for each different goal. The key idea of this policy is stabilising/keeping traders a specialist has already attracted and attracting those that are not attracted yet. The attractability of a trader is dependent on the overall design goal.

Clearing Policy. There are two main clearing policies used in TAC competitions, round-based and continuous. Round-based clearing clears at the end of each round, while continuous clearing clears whenever there is a new match available. Matching policy is sensitive to clearing policy. For instance, maximal matching will be useless with continuous clearing. Moreover, traders will have chances to revise their shouts if the market does not clear for some rounds during a day. We use a modified version of round-based clearing policy in this framework. Instead of clearing in each round, we choose a fixed number of clearing time

² The current best market of a trader is the market where it trades most.

points according to the number of goods each trader has, for example, we clear 5 times a day if each trader requires to exchange 3 items. Then we distribute clearing time points into the 10 rounds of a day by giving greater preference to the first 5 rounds. Thus, we clear more in the beginning of a day while waiting longer near the end of a day, because intra-marginal traders become less and less when it is approaching the end of a day and we want to give unsatisfied traders more chances to improve their shouts.

Charging Policy. Charging is a trade-off between traders' profits and a specialist's profit. It is not closely related to the above policies, but it affects traders' market selection. Therefore, most specialists in previous competitions do not charge in the beginning of a TAC game in order to attract traders. However, for most high intra-marginal traders, charging does not affect their profit too much, because they already reserved a large profit margin by bidding a very low (high) price to buy (sell). This framework will only focus on profit fee, as other fees, i.e. registration fee, transaction fee and information fee, could lead to 0 profit even for a trader who has successfully traded in the market.

4.2 Searching Adaptive Mechanisms

We know the main challenge for stabilising/attracting traders is stabilising their bidding prices, which depends on their bidding strategies. In other words, we might not be able to find a uniform mechanism that is attractive to traders with any kind of bidding strategy. Therefore, instead of searching for competitive mechanisms in a mixed environment from the very beginning, we propose a two-step approach. We first identify trader-dependent mechanisms that are competitive in an environment with only one kind of trader. Then we combine trader-dependent mechanisms together to achieve mechanisms that are competitive in any environment.

Trader-dependent Mechanism Design. Given the goal of a trader-dependent mechanism that we want to achieve (or a function of trader-dependent mechanism to maximise), we first set up the testing environment according to the goal and an initial mechanism as the current best mechanism, and then monotonically modify only one of the parameters in the search space to compete with the current best to find the next best one that increases the goal function the most, until we cannot find any modification that has any significant improvement of the function. Note that we require the modification of each parameter to be monotonic, i.e. update/change in one direction. Algorithm 4.2 describes the searching process for trader-dependent mechanisms. This algorithm will return mechanisms that locally maximise the goal function. In order to get an overall optimal mechanism, we can repeat this process with different initialisations.

Adaptive Mechanisms with Trader-dependent Mechanisms. Once we get trader-dependent mechanisms for each kind of trader offline, we will adapt

Algorithm 4.2. Searching Trader-dependent Mechanism

Input: m_0 : initial mechanism, f_m : a function of mechanism to maximise, δ : the minimum improvement
Output: m^* : the local best mechanism

```

1 begin
2    $CurrBest \leftarrow m_0$ ;
3   repeat
4      $m^* \leftarrow CurrBest$ ;
5     foreach policy parameter  $r$  do
6        $m' \leftarrow$  monotonically update  $r$  in  $m^*$ ;
7       if  $f_m(m') > f_m(CurrBest)$  then  $CurrBest \leftarrow m'$ ;
8     end
9   until  $f_m(CurrBest) < f_m(m^*) + \delta$ ;
10 end

```

them online for any market environment. The main idea is to use the classification learned in Section 3 to determine each trader’s category and apply the corresponding trader-dependent mechanism. However, we might end up with many inconsistent trader-dependent mechanisms that are required to run together for some environments. In such a case, we have to either apply only one of them or mix them by giving different priority to apply each of them. In order to make such a discrimination, we need to ascertain which trader-dependent mechanism will attract more good traders, which can be done, for example, by statistical analysing traders’ behaviour.

4.3 Experiments

In this section, we show a trader-dependent mechanism that is attractive to intra-marginal traders with GD bidding strategy, which is also the most attractive bidding strategy adopted by traders [14].

GD traders use the market history of submissions and transactions to form their beliefs over the likelihood of a bid or ask being accepted, and use this belief to guide their bidding [11]. Then the bidding strategy is to submit the shout that maximises a trader’s expected profit, i.e. the product of its belief function and its linear utility function.

Based on the search space given in Section 4.1 and our specialist agent *jackaroo*³, we identified a trader-dependent mechanism that is very good at attracting intra-marginal GD traders. The value of each parameter of the mechanism is given in Table 1, where A_r and B_r are respectively the accepted asks and bids until round r in one day. We have tested this trader-dependent mechanism (JaGD) with other competitive agents available from the TAC agents repository⁴, *CUNY.CS.V1* (Cu09.1), *CUNY.CS.V2* (Cu09.2), *Mertacor* (Me09), *cestlavie* (Ce09), *jacakroo* (Ja09) from CAT 2009 final, and *PoleCat* (Po10), *Mertacor* (Me10) from CAT 2010 final. Tables 2 and 3 show the average trading time

³ Achieved 3rd, 1st, and 2nd in CAT Tournament 2008, 2009, and 2010, respectively.

⁴ <http://www.sics.se/tac/>

Table 1. GD Attractive Mechanism

Policy	Parameter	Value
Accepting	F_t^a, F_t^b	6
	θ^a	$1 - \max(0, \frac{A_r - B_r}{A_r})$
	θ^b	$1 - \max(0, \frac{B_r - A_r}{B_r})$
Matching	Behaviour-based Maximal Matching	
Pricing	$minK$	0.15
Clearing	Modified Round-based	
Charging	12% profit fee	

Table 2. Average Trading Time Distribution of Each Type of Trader

	Specialists								Standard Deviation
	Cu09.1	Cu09.2	Me09	Me10	Po10	Ce09	Ja09	JaGD	
ZIC Sellers	39.60	40.40	54.20	136.27	56.83	55.07	42.87	74.77	31.95
ZIC Buyers	41.77	36.13	46.23	125.53	67.13	53.93	46.17	83.10	29.64
ZIP Sellers	15.43	16.77	50.00	179.20	62.50	50.83	<u>59.40</u>	65.87	51.04
ZIP Buyers	18.30	21.07	49.00	197.83	64.50	40.90	45.37	63.03	57.24
GD Sellers	20.73	22.46	49.29	77.80	<u>87.43</u>	62.37	37.84	142.09	40.21
GD Buyers	22.91	19.57	51.23	69.50	79.84	<u>69.66</u>	41.34	145.94	40.26
RE Sellers	53.10	47.31	53.59	89.76	69.46	67.90	55.91	62.97	13.43
RE Buyers	<u>55.19</u>	<u>51.56</u>	<u>55.61</u>	86.56	73.07	64.94	55.07	58.00	11.91

distribution of one CAT game (500 days), where the bold value in each row shows which market the traders in this row selected most and the underlined value in each column indicates which kind of traders were attracted most by the specialist in that column. The environment is mixed with 70 GD, 70 RE, 30 ZIC, and 30 ZIP buyers and sellers respectively, with valuations uniformly distributed in $[60,160]$, i.e. the perfect market equilibrium is 110. From Table 2 we can see that *JaGD* attracted about 30% of GD traders' trading time (the average for each market is 12.5%). Table 3 further shows that most traders attracted by *JaGD* are intra-marginal GD traders, and some lower extra-marginal traders are also attracted because of the use of maximal matching. It is worth mentioning that, except GD traders, this trader-dependent mechanism is not appealing to other traders, and it is also the case vice versa which can be seen from *Me10*.

5 A Framework for Behaviour-Based Mechanism Design

In this section, we want to summarise our behaviour-based design approach to a more general adaptive mechanism design framework based on traders' behaviour. This framework consists of *data acquisition*, *behaviour-based classification of traders*, *defining behaviour-based policies*, *trader-dependent mechanism design* and *integrating trader-dependent mechanisms*.

1. *Data acquisition* collects and aggregates market information, especially trader related information, which will be the foundation of the other components. Some statistical and data mining methods can be adapted here.

Table 3. Average Trading Time Distribution of Buyers

	Specialists								Standard Deviation
	Cu09.1	Cu09.2	Me09	Me10	Po10	Ce09	Ja09	JaGD	
<i>intra-marginal buyers (with valuations between 160 and 110)</i>									
ZIC	27.60	17.27	34.13	181.27	65.27	46.20	36.73	91.53	53.39
ZIP	18.70	20.85	42.10	256.15	58.85	27.40	40.90	35.05	79.31
GD	23.97	18.64	36.72	70.92	75.42	64.81	18.53	191.00	57.02
RE	42.53	39.88	48.84	113.66	82.91	68.22	47.84	56.13	25.12
<i>lower extra-marginal buyers (with valuations between 110 and 90)</i>									
ZIC	48.25	49.38	56.13	79.38	71.88	58.75	50.38	85.88	14.62
ZIP	16.60	23.40	46.00	89.80	71.80	60.20	34.20	158.00	45.77
GD	28.44	21.44	38.89	47.67	108.89	65.22	33.56	155.89	46.81
RE	68.86	57.86	55.14	66.00	70.29	62.71	59.14	60.00	5.43
<i>other extra-marginal buyers (with valuations between 90 and 60)</i>									
ZIC	64.71	61.43	60.86	58.86	65.71	65.00	61.57	61.86	2.39
ZIP	18.40	19.60	79.60	72.60	79.80	75.60	74.40	80.00	26.99
GD	19.40	20.24	76.56	75.32	75.76	78.24	77.00	77.48	26.36
RE	65.16	62.19	62.71	63.23	63.55	62.06	61.61	59.48	1.64

2. *Behaviour-based classification of traders* distinguishes traders in terms of their behaviour. This step heavily depends on the information obtained in the first step. Some machine learning methods, e.g. decision tree leaning, might be useful here.
3. *Defining behaviour-based policies* determines how to utilise behaviour in specialist policies. The main contribution of traders' behaviour in this stage is connecting the loosely coupled policies to reduce the search space.
4. *Trader-dependent mechanism design* identifies mechanisms that are competitive in environments with only one of kind of trader.
5. *Integrating trader-dependent mechanisms* combines all trader-dependent mechanisms to achieve mechanisms that are competitive under an environment containing a mixture of any kinds of traders.

6 Conclusion

We have introduced a behaviour-based adaptive mechanism design approach under the Trading Agent Competition Market Design platform. This approach consists of behaviour-based trader classification, mechanism design for specific environments (called trader-dependent mechanism design) and integrating trader-dependent mechanisms for any complex environments that are not known in advance. To the best of our knowledge, this is the first market design framework heavily depending on traders' behaviour (i.e. market history). By integrating traders' behaviour into market policies, we are able to constrain the search space of double auction mechanisms. More importantly, because of gaining an understanding of the market environment from traders' behaviour, the resulting mechanisms will apply differential policies for attracting different traders and therefore be more focused, more competitive and adaptive.

However, how to use traders' behaviour information more efficiently in trader classification and specialist policy design is worth further investigation. For instance, we might be able to use certain well studied methods from data mining,

e.g. decision tree learning, in behaviour-based trader classification, and even build a clearer relationship between the loosely coupled policies of specialist by using traders' behaviour information to further improve the design quality.

References

1. Cai, K., Gerding, E., Mcburney, P., Niu, J., Parsons, S., Phelps, S.: Overview of cat: A market design competition version 2.0. Technical report, University of Liverpool (2009)
2. Parkes, D.C.: Online mechanisms. In: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
3. Vytelingum, P., Vetsikas, I.A., Shi, B., Jennings, N.R.: Iamwildcat: The winning strategy for the tac market design competition. In: *Proceeding of the 18th European Conference on Artificial Intelligence*, pp. 428–432 (2008)
4. Stavrogiannis, L.C., Mitkas, P.A.: Cat 2008 post-tournament evaluation: The mercatocors perspective. In: *IJCAI Workshop on Trading Agent Design and Analysis (2009)*
5. Honari, S., Ebadi, M., Foshati, A., Gomrokchi, M., Benatahr, J., Khosravifar, B.: Price estimation of persiancat market equilibrium. In: *IJCAI Workshop on Trading Agent Design and Analysis, TADA (2009)*
6. Niu, J., Cai, K., Parsons, S.: A grey-box approach to automated mechanism design. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1. AAMAS 2010*, pp. 1473–1474 (2010)
7. Robinson, E., McBurney, P., Yao, X.: How specialised are specialists? Generalisation properties of entries from the 2008 and 2009 TAC market design competitions. In: David, E., Gerding, E., Sarne, D., Shehory, O. (eds.) *AMEC 2009. LNBIP*, vol. 59, pp. 178–194. Springer, Heidelberg (2010)
8. Niu, J., Cai, K., Parsons, S., Gerding, E., McBurney, P., Moyaux, T., Phelps, S., Shield, D.: Jcat: a platform for the tac market design competition. In: *AAMAS 2008 (2008)*
9. Gode, D.K., Sunder, S.: 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)
10. Cliff, D.: Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report (1997)
11. Gjerstad, S., Dickhaut, J.: Price formation in double auctions. In: *E-Commerce Agents, Marketplace Solutions, Security Issues, and Supply and Demand*, pp. 106–134. Springer, London (2001)
12. Erev, I., Roth, A.E.: Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88(4), 848–881 (1998)
13. Friedman, D., Rust, J.: *The Double Auction Market: Institutions, Theories, And Evidence*. Westview Press, Boulder (1993)
14. Phelps, S., Mcburney, P., Parsons, S.: Evolutionary mechanism design: a review. *Autonomous Agents and Multi-Agent Systems* 21, 237–264 (2010)