A Risk-Based Bidding Strategy for Continuous Double Auctions

Perukrishnen Vytelingum and **Rajdeep K. Dash** and **Esther David** and **Nicholas R. Jennings** [1]

Abstract. We develop a novel bidding strategy that software agents can use to buy and sell goods in Continuous Double Auctions (CDAs). Our strategy involves the agent forming a bid or ask by assessing the degree of risk involved and making a prediction about the competitive equilibrium that is likely to be reached in the marketplace. We benchmark our strategy against two of the most common strategies for CDAs, namely the Zero-Intelligence and the Zero-Intelligence Plus strategies, and we show that our agents outperform these benchmarks. Specifically, our agents win in 100% of the simulations against the ZI agents and, on average, 75% of the games against the ZIP agents.

1 Introduction

Software agents that can act autonomously and interact in flexible ways are increasingly being used in a variety of electronic marketplaces. One of the most common forms of these marketplaces is the Continuous Double Auction (CDA) in which traders submit offers to buy (bid) and offers to sell (ask) at any time during the trading period and in which the market clears.[2] continuously. Such CDAs have emerged as the dominant financial institution for trading securities and financial instruments and today the major exchanges (like the NASDAQ and the NYSE) use variants of the CDA institution.

In a centralised CDA where an auctioneer has complete knowledge of the market (i.e. is aware of all traders' private information), it is possible to achieve *optimal market efficiency* where the total profit of all the traders is maximised and the transactions are made at the theoretical *competitive equilibrium* (CE) price. However, such complete information is highly unlikely to occur in practice because traders are typically selfish profit maximisers.

Against this background, Smith [5] showed in his seminal study of competitive market behaviour that in experiments with human traders, the market efficiency achieved in a decentralised[3] environment is close that of the centralised one. More recently, Das *et al.* described a set of controlled experiments where human traders interact with software agents. Hence, they found that the agents consistently obtained higher profits than their human counterparts and this led them to speculate that agent-based strategies may improve to the point where they can always outperform human opposition [7].

Given this potential, several researchers have developed trading strategies for software agents. One of the earliest attempts in this direction was the Zero-Intelligence (ZI) strategy [4]. A ZI agent essentially ignores the state of the market when forming a bid or ask and submits a random value drawn from a uniform distribution. Although it is extremely simple, the ensuing market efficiency, if all agents adopt it, is close to optimal. However, Cliff *et al.* argue that this efficiency is a consequence of the intrinsic structure of the CDA, rather than an indicator of the suitability of the ZI as a profitable trading strategy [2]. To illustrate this point, they developed the Zero-Intelligence Plus (ZIP) strategy. In this strategy, the traders maintain a profit margin (the ratio of the trader's profit to its valuation of the good) that determines their bid or ask at any time during the trading process. Importantly, this profit margin adapts to the prevailing market conditions through a learning mechanism, so that the trader can submit bids or asks that remain competitive. Other strategies for the CDA include the GD strategy [8] which uses a belief function for price formation and the FL strategy [6] which uses fuzzy logic to decide on a bid or ask (based on a reference price, which is the median of the previous transaction prices). In empirical studies, it has been shown that the ZIP strategy outperforms the GD strategy [3], but nevertheless, we believe that ZIP is a too simple a bidding strategy. It relies principally on the adaptation of its profit margin to the prevailing market conditions. Moreover, we believe the FL approach of making decisions based on a reference price is sound, but that better results would be obtained if the reference was based on the CE price (since this is where prices are likely to stabilise).

Given these insights, we believe that the key to developing a more effective strategy is to perform some degree of prediction about the CE price (since it cannot be calculated *a priori* in decentralised environments), with bidding decisions based on the notion of risk attitude (adapted to best fit the prevailing market situations). We believe that risk (here defined as expected utility loss resulting from missing out on a transaction) is the most appropriate way to characterise the agent's willingness to trade, rather than the profit margin used in the ZIP strategy. To this end, we develop a profit-maximising adaptive bidding strategy based on risk and CE price prediction. We term our strategy as Risk-Based (RB) and study its behaviour in both homogeneous and heterogeneous agent populations. As the CDA cannot be analysed from a theoretical perspective, we measure its effectiveness by empirically benchmarking RB against the ZI and the ZIP strategies (since these are probably the two most common strategies).

This work advances the state of the art in the following ways. First, we develop a novel strategy based on the concept of risk. Thus, our agents form a bid or ask depending on how risk-seeking or risk-averse they are and adapt their risk attitude, based on past experience, to be more effective in prevailing market conditions. Second, we demonstrate the effectiveness of this approach to price formation in CDAs and show that our strategy is in experimental equilibrium [1] where there is no incentive for an RB agent to deviate to another strategy in a population of RB agents. The remainder of this paper is

[1] University of Southampton, Southampton SO17 1BJ, UK email:{pv03r, rkd02r, ed, nrj}@ecs.soton.ac.uk

[2] The market clears as soon as a bid exceeds an ask.

[3] Traders' preferences are private information in a decentralised environment.

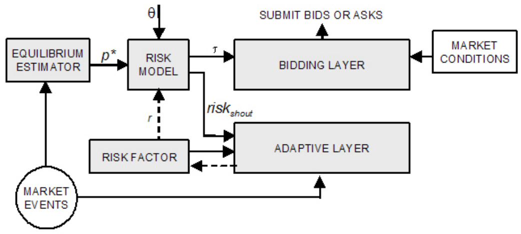

Figure 1. The model of the RB Strategy

structured as follows. We begin, in section 2, by describing our RB strategy and provide an empirical evaluation of its effectiveness in section 3. We conclude in section 4 and highlight future work.

2 The Risk-Based Strategy

The trading agent's preferences are determined by its *limit price*, which is the maximum a buyer is willing to pay and the minimum a seller will to accept. These preferences are then fed into the agent's trading strategy which determines how it responds to the market conditions, with bids and asks. This strategy can be risk-seeking, that is the trader tries to achieve high profit (but has a correspondingly lower probability of transacting) or risk-averse, which trades-off (lower) profit for a higher probability of transacting. The risk-neutral strategy considers a bid or ask that maximises its expected profit. Our RB strategy is flexible in that it can vary its risk attitude depending on the prevailing market conditions, to remain competitive.

In more detail, RB is represented by two distinct layers (see figure 1). The first represents the *reactive* bidding behaviour, where bids or asks are submitted according to a set of bidding rules (see figure 2). These rules are influenced by τ, the target price produced by the risk model. To compute this value, the risk model considers the agent's current estimate of the CE price, p^*, which is provided by the equilibrium estimator. The second layer represents the *adaptive* behaviour where the trader updates its risk factor, when triggered by a market event (such as when a transaction occurs or a new bis/ask is submitted). This change causes the agent to be more risk-seeking if it can transact at higher profit or more risk-averse if it is targeting too high a profit. Each component is now described in turn.

2.1 The Bidding Layer

At the beginning of a period, the trader has *no information* other than its limit price. Thus, a buyer, i, simply bids *towards* the minimum of its limit price, l_{ik}, and the outstanding ask, *oask*, (see equation 1) to maximise its surplus. Similarly, the seller, j, submits an ask towards the maximum of its cost price, c_{jk}, and the outstanding bid, *obid*, (see equation 2). The buyer or the seller will agree to a transaction when its bid or ask improvement (on the outstanding bid or ask respectively) is within an absolute value Δ_{spread} which we set to the minimum indivisible unit of currency (0.01).

$$ask_j = \begin{cases} oask - (oask - max\{c_{jk}, obid\})/\eta & \text{if } first\ round \\ oask - (oask - \tau)/\eta & \text{otherwise} \end{cases} \quad (1)$$

$$bid_i = \begin{cases} obid + (min\{l_{ik}, oask\} - obid)/\eta & \text{if } first\ round \\ obid + (\tau - obid)/\eta & \text{otherwise} \end{cases} \quad (2)$$

where $\eta \in [1, \infty)$ is a constant that determines the decrease rate of profit margin. A low η implies a faster rate of convergence of bids or asks towards a transaction price and, conversely, a high η implies a more conservative bidding approach and a slow convergence.

A buyer with a low limit price, should maintain its bargaining power by being able to bid for as long as possible. In contrast, a buyer with a high valuation, relative to the outstanding ask, should maximise its profit by *exploring* the market and exponentially reducing the *bid-ask spread* [4]. In either case, the behaviour is given in equation 1 and 2. However, if the outstanding bid (ask) is higher (lower) than the buyer's limit price (seller's cost price), the trader should not submit any bid (ask) until the beginning of the next round.

After the first transaction, the trader updates its estimate of the CE price, which it refines after each transaction. Initially, we set the trader's risk factor, r, close to 0 (adopting a risk-neutral attitude); $-1 \leq r < 0$ means that it is risk-averse and $0 < r \leq 1$ that it is risk-seeking. The risk model then uses the CE price estimation and the risk factor to calculate the new target price. Based on the latter and the set of bidding rules, the trader submits a bid or ask towards the target price in a similar way as in the first round.

Bidding Rules for Seller:
if (*bid-ask spread* $\leq \Delta_{spread}$) *accept obid*
else if (*limit price* $\geq$ *oask*) *submit no ask*
else
  if (no information) *submit ask given by Equation 1*
  else
    if (*obid* $\geq \tau$) *accept obid*
    else *submit ask given by Equation 1*

Bidding Rules for Buyer:
if (*bid-ask spread* $\leq \Delta_{spread}$) *accept oask*
else if (*limit price* $\leq$ *obid*) *submit no bid*
else
  if (no information) *submit bid given by Equation 2*
  else
    if (*obid* $\leq \tau$) *accept oask*
    else *submit bid given by Equation 2*

Figure 2. Bidding Rules

If the target price is higher (lower) than the outstanding ask (bid) at any time during the bidding process, the buyer (seller) accepts the outstanding ask (bid). When the outstanding bid (ask) exceeds (falls below) our target price, rather than not submitting a further bid (ask), we move the target price (to still be able to submit a profitable bid or ask) using a learning mechanism which we describe next.

2.2 The Adaptive Layer

In the adaptive layer, the trader uses a set of learning rules, summarised in figure 3, to update its risk factor to better fit prevailing market conditions. Specifically, a learning algorithm is used to increase or decrease the risk factor.

We adapt the trader's risk attitude by gradually changing its risk factor to a *desired risk factor*, $\delta(t)$, which depends on $risk_{shout}$ (the risk factor that corresponds to the last bid, if the agent is a buyer or to the last ask, if a seller). To decrease its risk factor, the trader sets the desired risk factor to slightly lower than $risk_{shout}$ ($\lambda = -0.05$) and when the trader is increasing its risk factor, it sets the desired risk factor to slightly higher than $risk_{shout}$ ($\lambda = 0.05$) when decreasing

[4] The bid-ask spread is the difference between *oask* and *obid*.

the risk factor. The algorithm enacts a continuous-space learning process that backprojects the error between the desired risk factor, $\delta(t)$, and the current risk factor, $r(t)$, onto the current risk factor. λ was chosen, based on simulation results. Specifically,

$$\begin{aligned} r(t{+}1) &= r(t) + \beta(\delta(t) - r(t)) \\ \delta(t) &= (1+\lambda)risk_{shout},\ \lambda = \{-0.05, 0.05\} \end{aligned} \tag{3}$$

where $\beta \in (0,1)$ is the learning rate of the algorithm and influences the moving rate of the target price. Next we consider our risk model from which we determine our target price, τ.

Learning Rules for Seller:
if (transaction at *q*)
    if ($\tau \leq q$) *seller must increase its risk factor*
    else *seller must decrease its risk factor*
else if (ask *a* submitted)
    if ($\tau \geq a$) *seller must decrease its risk factor*

Learning Rules for Buyer:
if (transaction at *q*)
    if ($\tau \geq q$) *buyer must increase its risk factor*
    else *buyer must decrease its risk factor*
else if (bid *b* submitted)
    if ($\tau \leq b$) *buyer must decrease its risk factor*

Figure 3. Learning Rules

2.3 The Risk Model

The role of the risk model is to generate the target price given the trader's risk attitude, which is defined by its risk factor. A target price equal to the estimated CE price implies that the trader is risk-neutral. When a trader adopts a risk-seeking attitude, it considers a target price that is below the CE (for buyer) or above the CE price (for seller), in order to obtain a higher profit margin. Conversely, a risk-averse attitude implies that the trader targets bids above (ask below) the CE, which improves the probability of its bid (ask) being accepted (but decreases its profit margin). Now, because the CE cannot be known *a priori* in the decentralised environment, we need to estimate the CE price based on the history of transactions (which we describe in sub-section 2.4). We model the risk differently for the different types of agents because they may react differently to the market conditions. Generally, agents can be of two types; namely, intra-marginal and extra-marginal.

A buyer (seller) is intra-marginal if its limit price (cost price) is higher (lower) than the CE price and is expected to transact in the market. In contrast, the extra-marginal buyer has too low valuation of the good while the the extra-marginal seller, too high cost, to be likely to transact.

First, we consider the intra-marginal trader. We identified the following constraints that our risk model should satisfy, given the trader's different risk attitudes (e.g. when the buyer is completely risk-averse, it targets a bid at its limit price while when it is completely risk-seeking, it targets a bid at 0). Therefore, the risk function must be continuous and pass through three specified values at risk factor -1, 0 and 1. It should also give symmetric behaviour for risk-aversion and risk-seeking and finally, the ask must be within an arbitrary maximum (*MAX*). Given these constraints, there is an infinite solution space for such a function and so, we chose a

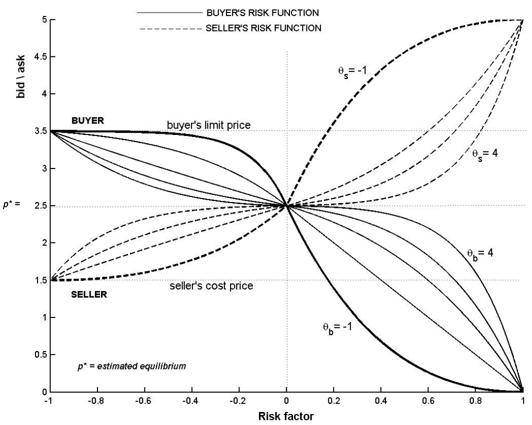

Figure 4. Risk for the intra-marginal trader for different θ

parameterised function (see figure 4) within the solution space with θ determining the behaviour of the function (i.e. its rate of change with respect to risk factor). When θ is high, the absolute gradient tends to 0 at the CE and increases as θ tends to -1. Experimental results suggest that the behaviour of our bidding strategy depends on the value θ and we study its implications in our empirical evaluation. Our function imposes a constraint on $\theta \in [-1, \infty)$ to limit the maximum absolute gradient and avoid the target price updating excessively with respect to change in the risk factor. In more detail, the intra-marginal buyer's and seller's model is as follows:

For a buyer,

$$\tau = \begin{cases} p^* \times (1 - r\,e^{\theta(r-1)}) & \text{if} \quad r \in (0,1) \\ (l_{ik} - p^*)\ (1 - (r+1)\,e^{r\theta_-}) + p^* & \text{if} \quad r \in (-1,0) \end{cases}$$

$$where\ \theta_- \ =\ (p^* \times e^{-\theta})/\,(l_{ik} - p^*) - 1 \tag{4}$$

For a seller,

$$\tau = \begin{cases} p^* + (MAX - p^*)\ re^{(r-1)\theta} & \text{if} \quad r \in (0,1) \\ p^* + (p^* - c_{jk})\ r\,e^{(r+1)\theta_-} & \text{if} \quad r \in (-1,0) \end{cases}$$

$$where\ \theta_- = log\,[(MAX - p^*)/(p^* - c_{jk})] - \theta \tag{5}$$

We next consider the case where the trader is extra-marginal. In this case, the trader modifies its risk function (shown in figure 4) to that of figure 5. This reflects the fact that the extra-marginal trader cannot be risk-averse and its risk factor has to be greater than 0 if it is to have any chance of transacting with a profit. Note that equations 4 and 5 also apply to the extra-marginal trader, with p^* replaced by the limit price.

2.4 The Equilibrium Estimator

We use the *moving average* method for estimating the CE price, p^*, based on the history of transactions. We make this choice because it is an objective analytical tool that gives the average value over a time frame spanning from the last transaction. Moreover, it is sensitive to price changes over a short time frame, but over a longer time span, is less sensitive and filters out the high-frequency components of the signal within the frame. Moving averages are commonly used to emphasize the direction of a trend and smooth out price fluctuations that can misinform the trader. Based on our assumption that the transaction prices converge to equilibrium, we introduce the notion

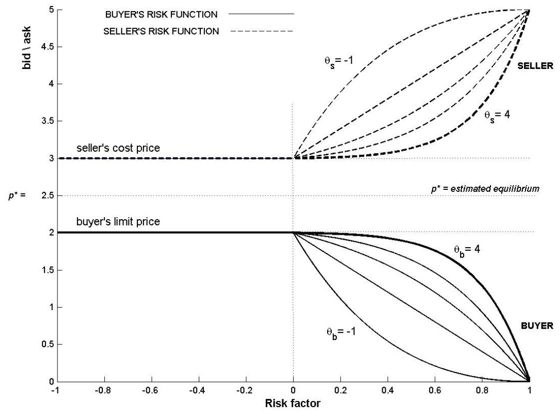

Figure 5. Risk for the extra-marginal traders for different θ

of *recency* in the Moving Average by giving more weight to the more recent transaction prices. In so doing, it emphasises any convergence pattern in the history, improving our estimation.

3 Empirical Evaluation

We describe the experimental setup of our CDA market and evaluate the RB strategy in both homogeneous and heterogeneous environments (that is with similar and different strategies respectively).

3.1 Experimental Setup

Our CDA market consists of a set I of buyers and a set J of sellers (we use a population of 10 buyers and 10 sellers). At the beginning of each trading period, each buyer $i \in I$ is given a set of M limit prices, $\vec{L_i}$ (ordered from highest to lowest), for the goods it wants to buy. Similarly, each seller $j \in J$ is given a set of N cost prices, $\vec{C_j}$ (ordered from lowest to highest) for goods to sell:

$$\vec{L_i} = \{l_{i1}, ..., l_{iM}\} \;\; \forall i \in I \quad (6)$$

$$\vec{C_j} = \{c_{j1}, ..., c_{jN}\} \;\; \forall j \in J \quad (7)$$

We assume that the trading agents have fixed roles (that is agents are either buyers or sellers) to conform to previous studies on CDAs [2, 5] and that agents do not have any information about the competition. A trading period is deemed finished after a specified period of inactivity which typically implies that the highest limit price of the buy side is lower than the lowest cost price of the sell side. We simulate a real-time CDA through the random activation of agents from a pool of active traders (traders that are still willing to trade) at each time step. The limit prices and cost prices are drawn from the same normal distribution and we set the constant η at 8.

Our simulations were conducted over 100 periods, with buyers and sellers receiving an allocation of limit and cost prices at the beginning of each period. We enforce a no-order queuing market rule where there is a unique outstanding bid and a unique outstanding ask at any time and the NYSE spread-improvement rule which states that any new bid and ask must improve upon the outstanding bid and ask respectively. Whenever the outstanding bid and outstanding ask match (bid is equal or greater than ask), a transaction is executed.

We first consider the performance of a strategy in a homogeneous environment where all agents use the same strategy. Here the metric is allocative efficiency which is the ratio of total profit from all trades to the maximum surplus (i.e. the total profit if allocation were optimal in a centralised environment). We then consider heterogeneous populations in which agents have varying strategies. Specifically, we consider a balanced population in which there are two different strategies and where each trader has a counterpart (using the other strategy). Each trader and its counterpart receive the same allocation of limit prices (for an unbiased comparison) and we compare the total profit from all traders using one strategy to that of traders using the other strategy.

3.2 Homogeneous Populations

Each buyer, i was given a set, $\vec{L_i}$ of 10 limit prices $l_{ik}, \forall k \in \{1, 2, ..., 10\}$ and each seller, j, a set $\vec{C_j}$ of 10 goods at different cost prices, $c_{jk}, \forall k \in \{1, 2, ..., 10\}$. All allocations were normally distributed between 1.50 and 4.50. The risk factor of the trader is initially arbitrarily distributed between -0.2 and 0.2 in order to ensure that agents start trading with a risk neutral attitude.

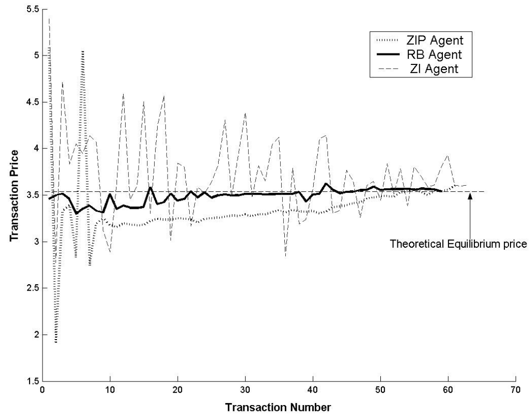

Figure 6. Transactions in a homogeneous ZIP, ZI and RB environment

In a typical experiment, using the above configuration, the theoretical CE price was 3.56 and the optimal allocation was 60. The result, in figure 6, clearly shows the transaction prices converging to the theoretical equilibrium within several rounds for the RB ($\theta = 1$) and ZIP strategies. We observe a non-converging fluctuating pattern (about p^*) with the ZI strategy. It can be seen that our strategy avoids the initial fluctuations of transaction prices which are observed with the ZIP strategy. Our agents trade very cautiously on the first round, with their decision-making based only on their private information and the bid-ask spread. After this, the RB agent adapts its risk factor much faster than the ZIP agent adapts its profit margin, which explains its faster convergence.

The performance of the different strategies, including the RB strategy for different values of θ, is shown in table 1. Using $\theta = -1$ in our RB strategy gives the best allocative efficiency with the fastest convergence towards the theoretical competitive equilibrium price. Generally, when θ is high, our RB agent exhibits a smooth and slow convergence. This contrasts with a rapid convergence and high-variance when θ is low. On average (over θ), our strategy has an allocative efficiency close to that of the ZIP strategy, but it exhibits a faster convergence.

3.3 Heterogeneous Populations

In these experiments, we compare RB with ZI. In this case, we measure performance by considering the strategy with the highest surplus

Strategy	Allocative Efficiency
ZI	97.47%
ZIP	99.69%
RB ($\theta = 2$)	99.29%
RB ($\theta = 1$)	99.29%
RB ($\theta = 0$)	99.09%
RB ($\theta = -1$)	99.63%

Table 1. Behaviour of a homogeneous population

in each period and the surplus difference is the margin by which one strategy outperforms the other.

θ	RB vs ZI wins	Surplus Difference	RB vs ZIP wins	Surplus Difference
2.0	100-0	12.20%	71-29	3.98%
1.0	100-0	14.08%	80-20	5.15%
0.5	100-0	9.81%	73-27	3.63%
0	100-0	9.38%	72-28	2.52%
-0.5	100-0	7.53%	69-31	3.68%
-1	100-0	6.60%	73-27	4.03%

Table 2. Behaviour of a heterogeneous population

As can be seen in table 2, the RB strategy dominates the ZI strategy by a very high margin and consistently outperforms the ZIP strategy (with a maximum of 5.15% surplus difference and winning in a maximum of 80 out of 100 periods, when $\theta = 1$). We observe a comparative drop in performance when $\theta = 2$, which suggests that when the gradient tends to 0 at risk factor 0, the RB agent is not adapting its risk behaviour fast enough to remain competitive in the market. We also observe that there is no correlation between θ and performance, and we can only state that the RB strategy is very effective in a heterogeneous environment when $\theta = 1$.

Moreover, it was observed RB strategy performs differently at different stages of the trading process. At the beginning of the period, when the trader needs to test the market, $\theta \leq 0$ allows the trader to adapt at a faster rate to market changes. However, when the transaction prices converge towards the theoretical CE price, $\theta \geq 1$ gives the best performance with less fluctuations in the target price. By constraining θ to a fixed value, we have an aggregate behaviour that is only effective at certain stages during the trading period, and is not necessarily the best behaviour.

3.4 Experimental Equilibrium

This set of experiments aim to determine whether there is any incentive for an agent to change from its current strategy to another one. When there is no such incentive for an agent to deviate from the strategy adopted by the population, as the chosen strategy provides the agent with the highest utility, then we describe this strategy as in experimental equilibrium[5] [1]. In these experiments, we put one agent of a particular strategy in an otherwise homogeneous population where agents use a different strategy. We then measured performance by how much more surplus the single agent obtains than its counterpart.

The results of table 3 show that in a ZIP or ZI population, a single RB agent performed considerably better than its counterpart which shows that there is an incentive for the ZIP or ZI agent to *defect* (i.e. use the RB strategy for higher profit). Conversely, the RB strategy

[5] An experimental equilibrium refers to an equilibrium which is measured experimentally with respect to a given set of strategies.

Strategy	Many ZIP	Many RB (θ=1)	Many RB (θ=0)	Many RB (θ=−1)	Many ZI
1 ZIP	-	-20.1%	-16.2%	-14.3%	-6.0%
1 RB$_{(\theta=1)}$	22.9%	-	-	-	9.5%
1 RB$_{(\theta=0)}$	16.5%	-	-	-	2.9%
1 RB$_{(\theta=-1)}$	13.3%	-	-	-	10.6%
1 ZI	7.8%	-21.3%	-24.6%	-21.8%	-

Table 3. Single agent in a homogeneous population of a different strategy

was less vulnerable to defection to other strategies as the single ZI or ZIP agent was outperformed by its RB counterpart. We conclude, therefore, that experimental equilibrium can only be achieved if all agents adopt the RB strategy.

4 Conclusions and Future Work

In this paper, we describe an adaptive bidding strategy, and where an agent can assess the risk associated with a bid or ask under current market conditions and bid accordingly. Specifically, the RB agent avoids making any initial *random* bid or ask (in contrast to the ZIP strategy) considering that there is no information other than its limit price. We demonstrated that our strategy outperforms two important benchmarks, namely the ZI and ZIP strategies, in a balanced heterogeneous population and that it performs well in a homogeneous environment. We also showed that there are incentives for an agent to defect from the ZI or ZIP strategy to the RB strategy, while it is not profitable to defect in a market populated by RB agents.

Further work includes the detailed study of the impact of the θ-parameter. In our current implementation, θ is fixed and we believe that by adapting it over the trading period the agent could better explore and exploit the market. In particular, we believe that a trader should have a low θ at the beginning of the trading period to attempt to increase its profit margin, and an increasing θ as the transaction prices converge to the competitive equilibrium.

ACKNOWLEDGEMENTS

We gratefully acknowledge funding from DIF-DTC 8.6 project (www.difdtc.com) and thank S. D. Ramchurn, A. Rogers and Minghua He for helpful comments.

REFERENCES

[1] Goldman C., Kraus S., and Shehory O., 'On experimental equilibria strategies for selecting sellers and satisfying buyers', *Decision Support Systems Journal*, (2003 (to appear)).

[2] Cliff D. and Bruten J. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91.

[3] Tesauro G. and Das R., 'High-performance bidding agents for the continuous double auction', *Proceedings of the Third ACM Conference on Electronic Commerce*, 206–209, (2001).

[4] Gode D. K. and 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).

[5] Smith V. L., 'An experimental study of competitive market behaviour', *Journal of Political Economy*, **70**, 111–137, (1962).

[6] He M., Leung H. F., and Jennings N. R., 'A fuzzy logic based bidding strategy for autonomous agents in continuous double auctions', *IEEE Trans on Knowledge and Data Engineering*, **15 (6)**, 1345–1363, (2003).

[7] Das R., Hanson J. E., Kephart J. O., and Tesauro G., 'Agent-human interactions in the continuous double auction', *Proceedings of the 17th Joint Conference on Artificial Intelligence*, 1169–1176, (2001).

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