

Impact of Misalignment of Trading Agent Strategy across Multiple Markets

Jung-woo Sohn, Sooyeon Lee, and Tracy Mullen

College of Information Sciences and Technology,
The Pennsylvania State University,
University Park, PA 16802-6823, USA
jwsohn@ist.psu.edu, sul131@psu.edu, tmullen@ist.psu.edu

Abstract. We examine the effect of a market pricing policy designed to attract high-valued traders in a multiple market context using JCAT software. Our experiments show that a simple change to pricing policy can create market performance effects that traditional adaptive trading agents are unable to recognize or capitalize on, but that market-policy-aware trading agents can generally obtain. This suggests as parameterized and tunable markets become more common, trading strategies will increasingly need to be conditional on each individual market's policies.

Keywords: market design, trading strategy, market selection strategy, multiple markets, pricing policy

1 Introduction

The recent global financial crisis provides ample incentive for understanding market dynamics such as the spread of price volatility across multiple stock markets in different countries [2]. Yet, while real-world stocks can simultaneously be listed on the New York Stock Exchange (NYSE) and NASDAQ, classical economic theories and market microstructure typically assume a single market due to analytical complexity. The CAT market design competition [9] is aimed at encouraging more empirical investigation of such multiple market landscapes. Software market specialists compete against each other and are scored based on a combination of market share, profit from charging traders fees, and transaction rate. Each day software traders select a market to trade in, and then place shouts (i.e., bids or asks) in the market. Within this multiple market landscape, specialists must consider the intertwined effect of (1) market policies, (2) trader's market selection strategies and (3) trader's trading strategies.

In this paper, we show how a slight change in market microstructure aimed at attracting high-valued traders, whose bid is higher or ask is lower than the market equilibrium price, has unexpected implications both for trader's trading strategies and market selection strategies. In general, these high-valued traders are referred to as *intra-marginal traders*, while the low-value traders who bid lower or sell higher than the market equilibrium price are called *extra-marginal traders*. A market specialist who has numerous intra-marginal traders can more easily match bids and asks,

resulting in a higher transaction rate score. Since intra-marginal traders are more likely to be matched with another intra-marginal trader (depending on the matching algorithm), they can often achieve higher profit, and thus are more likely to return to the market. In addition, markets with more intra-marginal traders can charge higher fees without substantially reducing trader profit, thereby increasing the market profit per trader. However, some types of fees have the effect of driving away extra-marginal traders, thereby reducing the specialist's overall market share. For this initial work, we only consider free markets, and do not consider the impact of fees.

1.1 CAT Background

The CAT organizers provide a Java-based client-server platform, called JCAT [8] where software for specifying both trader and market specialist behavior is available. JCAT code can also be extended to incorporate new strategies and policies. A game parameter file determines numerous aspects of the game including how a trader's private value is set (randomly or fixed) or the number of trading days.

Traders are specified by selecting a trading strategy that governs traders' bidding behavior and a market-selection strategy that governs how traders choose between markets. Trading strategies in JCAT include truth-telling where each agent bids its private value, Zero-Intelligence-Plus (ZIP) [3], and Gjerstad-Dickhaut [5]. Examples of market selection strategies include a random strategy, where traders select a market to trade in randomly each day, and an adaptive learning strategy based on the N-armed bandit approach which selects the market with the best expected profit for $(1 - \epsilon)$ percent of the time, and randomly exploring other markets ϵ percent of the time.

Market specialists are specified by setting several market design parameters including the accepting, clearing, pricing, and charging policies. We focus here on the pricing policy, which determines how the transaction price is set when a bid and ask are matched.

1.2 Biased k-pricing policy

The PSUCAT team in the 2008 CAT competition adopted a biased k-pricing policy related to k-double auctions [11]. K-pricing policy sets the transaction price by dividing the bid-ask spread profit by the parameter k with values between 0 and 1. For example when $k = 0.5$, the transaction price is set halfway between the bid and ask. Our biased k-pricing aims to attract intra-marginal traders by giving more of the bid-ask spread profit to intra-marginal shouts matched with extra-marginal shouts. In this case, if $k=0.9$, 90% of the bid-ask spread profit goes to intra-marginal trader and only 10% to the extra-marginal trader. For example, suppose the market equilibrium price is 100, and a buyer with a bid of 130 (i.e., intra-marginal bid) is matched with a seller who asks 120 (i.e., extra-marginal ask), then the intra-marginal buyer gets 90% of the bid-ask spread profit and the transaction price is set to 121. An unbiased k-pricing market policy with $k=0.5$ would have set the transaction price of 125, giving buyer and seller equal profit. However, we found that intra-marginal traders did not

effectively act on the biased k-pricing policy and thus did not favor our biased k-pricing market.

Further consideration showed that the biased k-pricing market introduces several inter-related issues. First, standard agent trading strategies are not attuned to market microstructure such as matching policy, pricing policy, and clearing policy. For example, the ZIP strategy assumes that the transaction price in the market is determined by the trading party who accepts the current market offer. When a buyer places a bid of \$100 and a seller accepts it, the transaction price becomes \$100. However, in the case of other market institutions such as sealed-bid auctions or when the market specialist pools and matches the shouts, the transaction price can differ. A seller who places an \$80 ask might be matched with a buyer who placed a \$100 bid, and the transaction price can be anywhere from \$80 to \$100.

Second, a related complication is that the adaptive trading strategies and the market selection strategies were not necessarily in alignment. For example, the ZIP trading strategy guides its adaptive behavior based on the profit margin, where profit is calculated as the difference between current shout price and private value. However, for market selection (other than random selection), traders calculate profit based on the difference between the transaction price and the trader's private value. So a buyer whose private value is 140, who bids 130, and who is matched with a seller for a transaction price of 121 will calculate its ZIP trading strategy profit as $140 - 130$ or 10, while its market selection strategy will calculate the profit as $140 - 121$ or 19. Thus adaptive traders cannot detect the extra profit achievable by the biased k-pricing policy and thus cannot use it to improve their trading results. On the other hand, intra-marginal truth-telling agents (who always bid their private value) are able to achieve high profits in the biased k-pricing market. Another inter-related problem is that the biased k-pricing market introduced a non-linear optimal bidding schedule for traders that simple adaptive behaviors did not handle well. We will discuss this aspect further in Section 3.

Lastly, even if the traders were able to discover the optimal bidding schedule, traditional trading agents do not distinguish their behavior based on what market they are in. Currently all trading strategies implicitly assume a single centralized market instead of conditioning their behavior on which market they are participating in.

In all of these cases, intelligent trading agents cannot fully exploit market microstructure to improve their profit. Since traders in the CAT tournament generally select markets based on the amount of the profit reaped so far, this can cause a trading agent to reach a suboptimal market selection decision. Thus our simple market policy change resulted in a more complex marketscape that simple adaptive behaviors were not able to optimize as well as a "stupid" truth telling trading strategy.

In this paper, we present a modeling approach that generalizes the inconsistency between the trading strategy and the market selection strategy. We also present a simulation result with a slight modification to the shout price variable in the original ZIP strategy, which we named ZIPK9Aware. The ZIPK9Aware agent was able to achieve higher overall profits by placing its shout price to take advantage of the market's biased k-pricing policy. We compare our ZIPK9Aware traders with standard ZIP and truth-telling traders.

2 Related Work

While continuous double auctions (CDA) are a well-established form of market institution, the complexity of even a single CDA means that human experiments and computer simulations are needed to more fully explore their properties. Smith [12] showed that even a few human traders in a continuous double auction would quickly converge to the equilibrium price. Gode and Sunder [6] introduced software zero-intelligence traders that randomly place shouts in a double auction market subject to a no-loss constraint. These zero-intelligence traders quickly converged to competitive equilibrium and suggested that equilibrium behavior can be achieved with extremely limited intelligence. However, Cliff [3] noticed that zero-intelligence traders owe their good trading behavior to regularities in the shape of the demand and supply curves, and thus did not generally perform well when the demand and the supply curve were asymmetric. Cliff designed his zero-intelligence plus (ZIP) traders on the assumption that traders should use at least some information about market conditions. ZIP traders collect information obtained from earlier shouts and trades and use it to adaptively set their target price margin for bidding. Similar to the ZIP strategy, the Roth-Erev (RE) strategy [10] is another adaptive trading strategy algorithm which adopts reinforcement learning. In the Gjestad-Dickhaut (GD) strategy [5], traders first collect the market history from an order book and use it to estimate the transaction success probability distribution, and then calculate the optimal shout price to maximize the expected profit.

All of the above work focused on traders trading a single good in a single market. If we extend this to trading a single good over multiple markets, some natural examples are financial markets and auction services. As Hasbrouck says [7] : "...Since a share of IBM is the same security whether purchased on the Midwest or Pacific Exchange, this is a particularly clear instance of multiple markets." Hasbrouck et al. focus on where price discovery occurs across multiple markets. In particular, they analyze NYSE and other regional exchanges seeking to establish "dominant" and "satellite" markets.

While the primary focus of this early research centered on price dynamics, Ellison et al. [4] view multiple market institutions as competing auctions. Their model analyzes what forces may cause markets to be concentrated. In particular, they consider why eBay, despite higher fees than its competitors, has achieved a dominant position in the online auction market while other auction services such as Yahoo have not acquired a large market share. Their preliminary results from two competing markets analysis imply that two markets will either co-exist or single market becomes dominant depending on various conditions.

Ellison et al. discuss the effects of attracting intra-marginal traders under various markets' charging policies. A typical strategy to acquire market share in auction services is to have zero listing fees. However, Ellison et al. found an unexpected effect of zero listing fees in Amazon or Yahoo was that they acquired numerous non-serious sellers with high reserve prices. In turn, serious buyers switched to other auction services. Ellison et al. also discuss the possibility that small markets can achieve high market efficiency by attracting high-value traders but do not proceed with further modeling analysis.

Westerhoff et al. [13] use an agent-based model to consider how small transaction taxes such as Keynes-Tobin tax can act to reduce market price volatility in a multiple market environment. They found that a transaction tax in one market reduces that market's price volatility but increases the volatility in the un-taxed market. However, when a transaction tax is imposed on both markets, then both markets show reduced price volatility. This suggests that market regulators in different markets could coordinate market transaction taxes to assist in controlling global market price volatility. Westerhoff's analysis also implies that once one market has imposed a transaction tax, other market regulators may wish to impose a transaction tax on their market so as not to face increased price volatility.

3 Problem Description

A trader in a multiple market scenario faces a two-stage process where the trader first selects the market in the first stage and places a shout in the market for the second stage. While the trading strategy and the market selection strategy are determined by the trading agent side, a market specialist selects one or more market policies such as pricing policy, in an effort to attract or keep certain types of traders based on their market selection strategy. However if a trader does not optimally trade in the given market, then a trader's market selection strategy will not necessarily operate optimally either. In this section, we characterize potential misalignments between agent trading strategies and market selection strategies due to the biased k-pricing policy.

3.1 Alignment of Trading Strategy and Market Selection Strategy

Trader profit is composed from two quantities. Let r be the transaction price determined by the market specialist, while s is the shout price and λ is the private value of the trading agent. The trading profit of $f(s)$ is composed from $|\lambda - s|$, the difference between the trader's private value and the shout price, and from $|s - r|$, the difference between the shout price and actual transaction price, so that $f(s) = |\lambda - s| + |s - r|$. Note that s is determined by the trader's bidding activity while r is determined by the market specialist.

Introducing the biased k-pricing policy to this framework, r becomes a function instead of a constant. Let s' be the price of the matched shout, then the profit now becomes $f(s) = |\lambda - s| + |s - r(s, s')|$. For the biased-k pricing, $r(s, s') = \{1 - k(s, s')\} s + k(s, s') s'$ where $k(s, s')$ becomes 0.9, 0.5, or 0.1 depending on whether s or s' are intra-marginal or extra-marginal shouts with respect to the current estimated equilibrium price p^* . Table 1 summarizes the k value assigned to the traders depending on their shout price s and s' . Note that traders can have some control over k by manipulating s in this framework.

However, the concept of trader's shout prices having an effect on market transaction price is not in alignment with most adaptive trading strategies. Traders consider k as exogenously determined by the market specialist. Thus, a trader's profit from an adaptive trading strategy actually stays as $f(s) = |\lambda - s| + |s - r|$, where r is

a constant and thus $|s - r|$ cannot be easily maximized by varying s . ZIP traders even take this simplification one step further by ignoring any transaction-related profit so that their profit function becomes simply $f_{zip}(s) = |\lambda - s|$.

Table 1. k value assigned to each traders depending on the matched bid s and the ask s'

$(k(s,s')$ for bid, $1-k(s, s')$ for ask)	Intra-marginal ask ($s' < p^*$)	Extra-marginal ask ($s' > p^*$)
Intra-marginal bid ($s > p^*$)	$(k(s, s') = 0.5, 1-k(s, s') = 0.5)$	$(k(s, s') = 0.9, 1-k(s, s') = 0.1)$
Extra-marginal bid ($s < p^*$)	$(k(s, s') = 0.1, 1-k(s, s') = 0.9)$	trade not possible

Non-intelligent traders such as truth-tellers, who do not control their shout price s , implicitly get the best k . Since $s = \lambda$ for truth-tellers, their profit function $f(s)$ is reduced to $f_{TT}(\lambda) = |\lambda - r(\lambda, s')|$. For intra-marginal truth-tellers, k will be either 0.9 or 0.5 but never $k=0.1$. Now consider an imaginary price-taking trader for another example. Suppose the market provides a buyer with a quote of $p^* + \varepsilon$ and provides a seller with the quote of $p^* - \varepsilon$ in the hope that traders will place shouts honoring the biased k -pricing policy. The profit now becomes $f_{PT}(p^*) = |\lambda - p^*| + |s - r(p^*, s')|$ where ε is infinitesimally small and can thus be ignored with k either 0.9 or 0.5 for intra-marginal traders. This exhibits an interesting implication that less intelligent trading strategies might actually optimize the profit coming from the transaction price set by the market in some cases.

3.1 Modifications in Trading Strategy under Biased K-Pricing Policy

Now we consider how an adaptive trading strategy could successfully act on the biased k -pricing policy. Figure 1 shows the profit schedule of an intra-marginal buyer when faced with the biased k -pricing policy. For simplicity, the seller is assumed to be an intra-marginal truth-teller. Note that the profit drop occurs around the market equilibrium price p^* when the intra-marginal buyer is matched with another intra-marginal seller.




Fig. 1. Profit drop caused by biased K9 pricing policy for a buyer bidding at price (x-axis) with private value 150. (Seller is an intra-marginal truth-teller with the private value 50)

The buyer's profit drop from placing an intra-marginal bid at $p^* + \varepsilon$ to placing an extra-marginal bid at $p^* - \varepsilon$ is $f(p^* + \varepsilon) - f(p^* - \varepsilon) = (0.9 - 0.5) (p^* - s') = 0.4 (p^* - s')$

where ε is infinitesimally small. (Since p^* itself is on the line between the intra-marginal and extra-marginals, we use ε to force the shout to be either intra-marginal or extra-marginal.) Now the profit drop depends on the matching shout price of s' since p^* can be assumed to be constant. Next consider the profit drop, δ in Figure 1, where the bidder is penalized for placing a bid in the range $[p^* - \delta, p^*]$. To avoid being penalized, the trader must estimate δ . By equating $f(p^*) = \lambda - (0.5 p^* + 0.5 s')$ with $f(p^* - \delta) = \lambda - (0.9 (p^* - \delta) + 0.1 s')$, we have $\delta = (4/9)(p^* - s')$. Since s' cannot be directly calculated by the trader, its trading strategy has to estimate s' . Figure 2 shows a 3-d plot of an intra-marginal buyer's profit across various possible combinations of (s, s') . As can be seen from the graph, profit estimation can become complex, which in turn results in the difficulty of estimation of δ .




Fig. 2. 3-dimensional plot of the trade profit for intra-marginal buyer with the private value of 150 under K9 pricing policy. Market equilibrium price $p^* = 100$.

One simple way to tackle this situation is to assume a constant average δ for trading agents regardless of their private value and to check its effect with simulation experiments. In the next section, we show that adjusting the trading strategy so that traders do not place shouts in the interval of $[p^* - \delta, p^*]$ for buyers and $[p^*, p^* + \delta]$ for sellers can lead to a significant increase in trader profit and improve the decisions of the trader's market selection strategy.

4 Experimental Setup

For the experimental setup we used the Java-based JCAT market simulation platform. The trading population consisted of 50 buyers and 50 sellers. We compare the trading strategies truth-telling, ZIP, and our ZIPK9Aware strategy. The ZIPK9Aware strategy is a simple modification to the original ZIP strategy that avoids placing a shout in the range of $[p^* - \delta, p^*]$ for buyers and $[p^*, p^* + \delta]$ for sellers. For each shout generated in that range, the ZIPK9Aware strategy recasts it as a bid for p^* to avoid facing a profit drop.

Private values for buyers and sellers are randomly drawn from a uniform distribution of $[50, 150]$, giving a theoretical market equilibrium price of 100. Each trader is endowed with 5 goods to trade per game day. A single game lasts for 100 days and each game day has 10 rounds. To ensure our results are consistent, each game result is averaged over 10 trials.

We used a market selection strategy based on N-armed bandit problem in which the trader selects the market with the highest expected profit for 90% of the time and randomly selects markets 10% of the time. The market accepting policy is the same as the New York Stock Exchange (NYSE) spread-improvement rule, which is typically found in CDA experiments, and requires that new bids (or asks) must be higher (or lower) than the current best bid (or ask) in market. The market clears matching bid and ask whenever it finds the best matching pair.

4.1 ZIP strategy and ZIPK9Aware strategy

To make trading strategies detect the profit drop caused by our biased k-pricing policy, we modified the original ZIP strategy into our ZIPK9Aware strategy. See Bagnal and Toft [1] for further algorithmic details on ZIP strategy.

One issue related to the experimental design with ZIPK9Aware strategy is that a ZIPK9Aware trader needs to have an appropriate estimate of the profit drop δ in advance to determine the threshold where it changes the bidding strategy from ZIP to ZIPK9Aware and vice versa. Under the assumption that $p^*=100$ and the matching shout price of s' is drawn from a uniform distribution of $[50, 100]$, the expected value of δ becomes $100/9$ since $E[\delta] = (4/9)(p^* - E[s']) = 4/9 (100 - 75)$. However, we chose 15 which is slightly greater than $100/9$ so that ZIPK9Aware traders are more likely to behave as ZIPK9Aware traders in uncertain situations, not as the original ZIP traders. Thus a ZIPK9Aware buyer will update its shout to $p^* + \varepsilon$ when its shout price is within $[p^* - \delta, p^*]$ interval to avoid being penalized by $k = 0.1$ pricing when it crosses the K5/K9 threshold. Similarly, a ZIPK9Aware seller will lower its shout price to $p^* - \varepsilon$ for the shout price interval of $[p^*, p^* + \delta]$. As discussed in the previous section, finding the optimal δ can be complicated so we initially took a simplification approach by using constant δ to check if our general idea is feasible.

5 Experimental Results

5.1 Truth-teller Case

We start our experiments using truth-telling traders as a baseline case. Figure 4 shows the CAT tournament score for the averaged result of 10 experimental runs where M1 is a free market with K5 pricing policy and PSUCAT is a free market with the biased K9 pricing policy.¹ In CAT tournaments, each market's scores are composed of market share, profit ratio earned by the market, and the transaction success rate for the shouts placed in the market. These three factors are weighted equally and added together to evaluate the market's performance for a game day. Since neither market charges fees, their profit score is zero. The total score becomes

¹ For convenience, we introduce the term K5 market for the market with the biased k-pricing policy of $k=0.5$ and K9 market for the market with the biased k-pricing policy set to $k=0.9$. The data shown in the Figures are all averaged on 10 experimental runs.

the sum of each market's market share and transaction success rate scores. Figure 4 shows that the K9 market acquires both a larger market share and a higher transaction ratio than the K5 market. This implies that more truth-tellers prefer the K9 market, which in turn implies that traders earn more profit in the K9 market. However, large number of traders do not necessarily result in better market performance because markets cannot match trades when there are a high percentage of extra-marginal traders as pointed out by Niu et al. [9]. Figure 5 shows buyer's average private values for each market, while Figure 6 shows seller's average private values.




Fig. 4. JCAT scores for K5 (M1) and K9 (PSUCAT) markets with truth-telling traders




Fig. 5. Average private values for K5 (M1) and K9 (PSUCAT) markets with truth-telling buyers

As shown above in Figure 5 and 6, the K9 market attracts intra-marginal traders who have higher private values for buyers and lower private values for sellers than the theoretical private value average of 100. The separation effect looks more evident from the seller side in figure 6.

4.2 ZIP trader case

Figure 7 shows the game scores for standard ZIP traders, who do not recognize the biased k-pricing policy. Unlike the previous case, the market share between K5

market and K9 market does not show any significant difference nor does the transaction rate score. Clearly the K9 pricing policy did not make a significant difference in market selection behavior for ZIP traders.

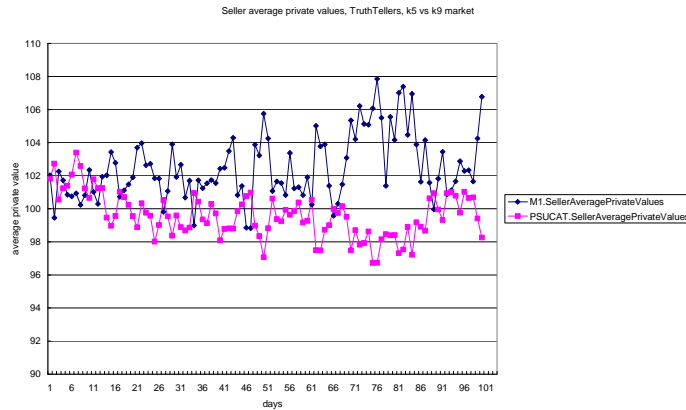


Fig. 6. Average private values for K5 (M1) and K9 (PSUCAT) markets with truth-telling sellers




Fig. 7. JCAT scores for K5 (M1) and K9 (PSUCAT) markets with ZIP traders

Average trader private value plots shown in Figures 8 and 9 do not demonstrate a significant difference either. While the K9 market seems to attract intra-marginal traders until day 20, the separation effect is weak after that. This same pattern can be observed for both buyers and sellers.

4.3 ZIPK9Aware Case

Figure 10 shows the CAT game scores using ZIPK9Aware traders. The K9 market acquires traders from the K5 market, which can be seen from increasing market share and the transaction success rate drop for K5 market. Investigation of the log file for individual trial runs revealed that the K5 market temporarily experienced zero transactions in several runs which drove down the averaged value.




Fig. 8. Average private values for K5 (M1) and K9 (PSUCAT) markets with ZIP buyers




Fig. 9. Average private values for K5 (M1) and K9 (PSUCAT) markets with ZIP sellers

Figure 11 shows a different pattern than the previous cases where average private values for K5 market buyers show a steep drop compared to the increase in average for K9 market buyers. This implies that the K9 market not only attracts intra-marginal traders but also extra-marginal traders as well. Given the sudden transaction success rate drop in K5 market, apparently the K9 market successfully attracted both intra-marginal and extra-marginal traders by providing minimum trading profit for extra-marginal traders and not harming trading profits for intra-marginal traders. The result is that both intra-marginal and extra-marginal traders are attracted to the K9 market since extra-marginal traders can sometimes make trades with intra-marginals when attracted to the K9 market.

However, in examining the 10 trials individually, 2 trials produced the opposite result where the main market share went to the K5 market and intra-marginal traders were attracted to the K5 market. While we do not understand exactly why these anomalous cases occur, one possibility is that we applied the same δ of 15 for all traders. The optimal size of δ depends on the matching shout, but we simplified the estimation process for the purpose of our experimental trials. In future work, we plan to increase the sophistication of the ZIPK9Aware trading agents to see if this produces the desired effect.

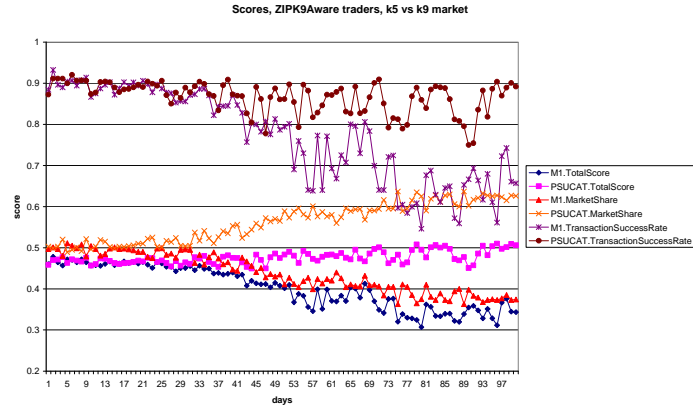


Fig. 10. JCAT scores for K5 (M1) and K9 (PSUCAT) markets with ZIPK9Aware traders




Fig. 11. Average private values for K5 (M1) and K9 (PSUCAT) markets with ZIPK9Aware buyers




Fig. 12. Average private values for K5 (M1) and K9 (PSUCAT) markets with ZIPK9Aware sellers

4.4 Comparison of Total Profit

Table 2 shows the comparison of average total profit reaped in K5 and K9 markets. ZIPK9Aware traders are making the highest total profits in K9 market. While ZIP traders still seem to make more profits on average in K9 market, the profit increase about 38,513 is relatively small compared to truth-teller and ZIPK9Aware trader cases of 172,193 and 132,993 respectively.

Table 2. Average total profit (and the stdev) earned in K5 market vs. K9 market over 10 runs. Numbers preceded by * denote that comparison is significant at 95% confidence interval under the null hypothesis that average total profit in K9 market is the same to that of K5 market. # denote significance at 90% confidence interval.

	K5 market			K9 market		
	Buyer	Seller	Total	Buyer	Seller	Total
Truth-teller	72105 (11157)	72105 (11157)	144207	*162715 (22014)	*153684 (24384)	316400
ZIP	108736 (35416)	110771 (22531)	219507	137581 (25907)	120439 (20127)	258020
ZIPK9Aware	108108 (80470)	99682 (66016)	207790	#159721 (83000)	*181061 (76755)	340783

In the K9 market, more profit is made by the agent trading strategies of truth-telling and ZIPK9Aware than the original ZIP traders. This implies that trading policies in alignment with the K9 pricing policy actually allows traders to acquire larger total profit.

5 Conclusion and Future Work

Our preliminary research demonstrates that market policy and agent trading behavior need to be aligned to perform effectively. We explore the implications of a biased k -pricing policy with $k=0.9$, called K9 pricing policy. This policy was aimed at incentivizing intra-marginal traders to favor the K9 market. We showed that the ZIP trading strategy is not in alignment with the K9 pricing policy since it does not consider the actual transaction price, which in turn leads to sub-optimal bidding decisions and market selection decisions. We developed a ZIPK9Aware trading strategy, a simple modification on ZIP strategy to verify our argument that trading strategies should be in alignment with market policies. With ZIPK9Aware traders, the K9 market was able to attract more intra-marginal traders than the K5 market. The K9 market was also able to attract more market share and total profit. Thus our experiments show that a simple change to market pricing policy can create market performance effects that traditional adaptive trading agents are unable to recognize or capitalize on, but that market-policy-aware trading agents can obtain (most of the time). This suggests as more parameterized and tunable markets become more common, trading strategies will increasingly need to be conditional on a specific market's policies.

Alternatively from the market design point of view, our experimental results also suggest the not-surprising idea that market specialists should consider participant's trading strategies when selecting a market policy. In our experiment, truth-tellers and ZIPK9Aware traders were able to take advantage of the biased k-pricing policy, while ZIP traders were not. An interesting question is whether human traders would recognize the impact of a $k=0.9$ pricing policy and act more like ZIPK9Aware traders or if they would act more like ZIP traders. A more general question is how well humans (or trading agents) can recognize and act on various market policies to maximize profits.

Finally, our results suggest that traditional adaptive trading strategies should be extended to keep a separate trading history for each market. This may be especially important if markets have substantially different market policies.

For future work, we plan to test our result with other intelligent trading strategies such as GD and RE. In addition, we will investigate further anomalous cases where the ZIPK9Aware trading agents prefer the K5 market over the K9 market. While our preliminary work used the same δ values for all ZIPK9Aware traders, we will optimize δ for each individual trader's private value. We also plan to investigate our results using human subjects to see if they can recognize the K9 profit drop around the market equilibrium price.

Acknowledgements

We would like to thank the CAT tournament organizers and developers for stimulating our research work as well as our reviewers for their helpful comments.

References

1. Bagnall, A. J., and Toft, I. E., Zero Intelligence Plus and Gjestad-Dickhaut Agents for Sealed Bid Auctions, Workshop on Trading Agent Design and Analysis, part of International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), pp. 59--64, (2004)
2. Buchanan, M., This Economy Does Not Compute, The New York Times, October 1 (2008)
3. Cliff, D., Minimal-intelligence Agents for Bargaining Behaviors in Market-based Environments, Technical Report, HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England (1997)
4. Ellison, G., Fudenberg, D., and Mobius, M., Competing Auctions, Journal of the European Economic Association, Vol. 2 (1), pp. 30--66, MIT Press (2004)
5. Gjestad, S., and Dickhaut, J., Price formation in Double Auctions. Games and Economic Behavior, Vol. 22, pp. 1-29 (1998)
6. Gode, D. K., and Sunder, S., Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as partial substitute for Individual Rationality. The Journal of Political Economy, Vol. 101 (1), pp. 119--137, February (1993)
7. Hasbrouck, J., One Security, many Markets: Determining the Contributions to Price Discovery, Journal of Finance, Vol. 50 (4), pp. 1175, Waverly Press Inc., New York (1995)

8. <http://jcat.sourceforge.net>
9. Niu, J., McBurney, P., Gerding, and E., Parsons, S., Characterizing Effective Auction Mechanisms: Insights from the 2007 TAC Market Design Competition. In Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems, Estoril, Portugal (2008)
10. Roth, A., and Erev, I., Learning in Extensive-form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term, *Games and Economic Behavior*, Vol. 8 (1), pp. 164--212, (1995)
11. Satterthwaite, M. A., and Williams, S. R., Bilateral Trade with the Sealed Bid K-Double Auction: Existence and Efficiency, *Journal of Economic Theory*, Vol. 48, pp. 107--133 (1989)
12. Smith, V. L., An Experimental Study of Competitive Market Behavior, *Journal of Political Economy*, vol. 70 (2), pp. 111-137, April (1962)
13. Westerhoff, F. H., and Dieci, R., The Effectiveness of Keynes-Tobin Transaction Taxes when Heterogeneous Agents can Trade in Different Markets: A Behavioral Finance Approach, *Journal of Economic Dynamics and Control*, Vol. 30 (2), pp. 293--322, February (2006)