Adaptive Market Design with Linear Charging and Adaptive k-Pricing Policy

Jaesuk Ahn and Chris Jones

Department of Electrical and Computer Engineering,
The University of Texas at Austin
{jsahn, coldjones}@lips.utexas.edu

Abstract. We describe a possible design strategy for a market specialist in the TAC Market Design Competition. Specifically, we describe both a linear charging policy that steadily increases the fees charged to trading agents interacting with the market, and an adaptive pricing policy that alters the trading price of goods based on the relative numbers of buyers and sellers in the market and the previous history of the market condition. We offer up two different sets of experimental data, the first to illustrate some basic properties about successful markets in the Market Design Competition, and the second to show the relative superiority of our chosen strategies over several possible alternatives.

Keywords: Auction Mechanism, Market Design.

1 Introduction

The Trading Agent Competition has, until now, primarily focused on developing strategies for agents proactively competing in various simulated economic activities. More particularly, the classic Trading Agent Competition features automated trading agents competing and trading with one another to maximize profit on a limited supply of goods on behalf of simulated clients [1], while the TAC Supply Chain Management competition features agents competing directly with each other for customers and supplier output in a multi-faceted supply chain simulation [2]. Recently, however, a group of researchers spearheaded by Gerding, Jennings, McBurney and Phelps proposed the development of a TAC Market Design Competition. Such a mechanism would allow participants to explore how different market design strategies promote or suppress certain global system properties, as well as which strategies might successfully attract various types of automated agents [3].

The TAC Market Design Competition represents a simulated double auction market, such as a stock or commodities exchange, where dedicated buyer and seller agents seek to find complementary partners to complete trades with. These trading agents are fully automated and are outside of the competition participants' control, but must trade their goods by participating in markets designed by the participants. More particularly, participants program specialist agents which control the fees a market charges, the pricing policy of the market, which controls the exact price a good will trade at, the clearing policy of a market, which controls when a market will execute trades between traders, and a quote accepting policy, which controls which buy or sell orders a market will accept [3].

Participants compete to create specialists which will extract the highest number of fees from a common pool of traders; however, since the automated traders are rational actors (ideally), markets must perform a careful balancing act between implementing policies which will charge high fees and drive traders towards other, cheaper markets, and policies which will charge modest fees but attract many traders.

Before the first official TAC Market Design Competition in Summer of 2007, students of Dr. Peter Stone's Agent-Based Electronic Commerce class at the University of Texas at Austin were offered the opportunity to work with an early version of the Market Design testbed software by developing their own market specialists and running a competition between themselves. The following sections describe one such strategy in detail. Specifically, Section 2 outlines a set of preliminary experiments which illustrate certain properties of the trading agents within the Market Design testbed. Section 3 describes in detail a strategy derived from these observations, and gives a brief overview of several alternative charging policies. Section 4 describes a set of experiments which compare the presented market specialists to the alternative strategies in various market conditions, while Section 5 provides analysis of the experimental results and presents some concluding remarks.

2 Market Design Competition Observation

2.1 Effect of Market Charging Policy

The two primary variables determining the income earned by any given market in the competition are the number of agents in the market, and the level of fees charged to each agent. Because agents have the ability to move between markets on different simulated "days" of the competition, relatively high fees might be expected to drive agents away from a given market, while lower fees might be expected to attract agents to a given market.

However, it is reasonable to suspect that fees lowered past a certain level cannot remain competitive, regardless of how many agents those fees attract. For example, suppose only two specialists existed in a competition – one controlling a first market that charged a flat fee of 1 credit per agent, and one controlling a second market that charged a flat fee of 500 per agent. In a trading competition with 100 agents, even a single agent selecting the second market would mean that the first market would lose the competition. Such a scenario could happen for any number of reasons – agents are not always perfectly rational, agents must operate without perfect information about fees, agents may find that a given good or service can only be obtained within a given market.

Accordingly, we performed a simple set of experiments with the basic "Fixed Charging Policy" specialists provided as default implementations by the testbed. Specifically, we took seven different specialists and placed them in competition with each other, with one specialist charging the maximum set of fees allowed by the student competition, and the six remaining specialist charging 90%, 75%, 50%, 25%, 10% and 0% of these maximum fees respectively. The experiments featured 100 trading agents evenly divided between both buyers and sellers, and divided between the various trading strategies provided by the testbed trader agent implementation. Figures 1 and 2 show the average results of 5 trials of these experiments.

As can be seen in Figure 1, although the 75% specialist charges 25% less in fees than the 100% specialist, it earns approximately as much in total. This is almost certainly due to the fact that the 75% specialist attracts a slightly higher number of traders, as seen in Figure 2. At the same time, however, 75% seems to represent some sort of "sweet spot" – specialists charging less than the 75% specialist earn significantly less than the top-earning specialist, even though they attract significantly more agents. These results would seem to confirm our initial hypothesis that the traders in the market design competition are not fully rational – obviously, if they were, all traders would congregate in the zero fee specialist. Instead, traders in the competition seem only partially bound to alter their behavior based on the fees being charged, and some traders can be expected to remain with any given specialist, no matter how high the associated fees might be.

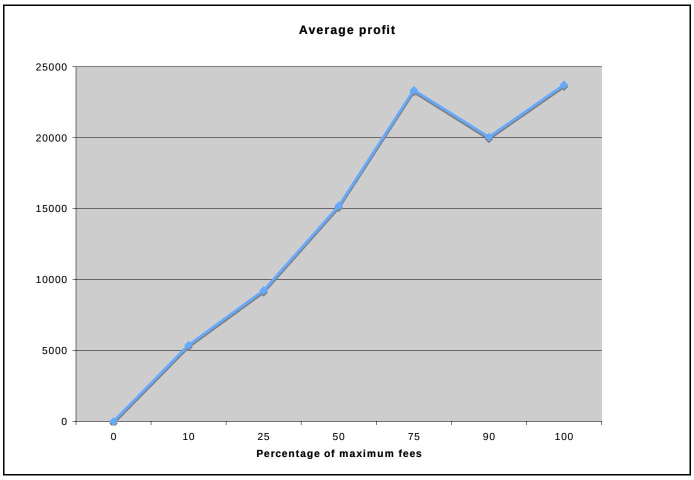

Figure 1. Average profit as a function of fees charged

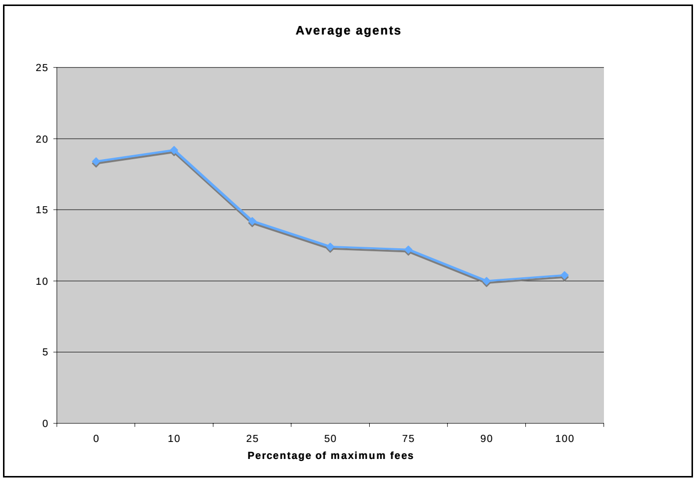

Figure 2. Average agents per market as a function of fees charged

Accordingly, we began to investigate how changing the fees charged over the course of the game effected trader behavior. More particularly, we wished to determine if a significant number of traders could be drawn in by initially low prices, and then stay even as prices were raised. In addition to the Fixed Price strategy mentioned above, where fees are held constant at the highest fees possible across the entire experiment, we also examined a Honeypot charging policy, where fees were set at zero until a given point in the game, after which they were raised to the maximum allowed, and a Linear Charging policy, where prices were raised from zero at a constant rate until they reached the maximum cap. Figure 3 shows the average profit per day for one such set of experiments, while Figure 4 shows the average number of agents registered with a given specialist per day for the same experiments.

As Figure 3 indicates, both the Honeypot and Linear policies outperform the Fixed Price policy. Furthermore, although the Honeypot policy earns higher returns than the Linear policy once it raises its prices (thanks to a higher average number of traders, as shown in Figure 4) its overall earnings are lower than Linear's since Linear earns a significant amount of credit during its "ramp up" period. Furthermore, as Figure 4 shows, the average number of traders per specialist does not change significantly, even as the specialists make major changes in their fee structures over time.

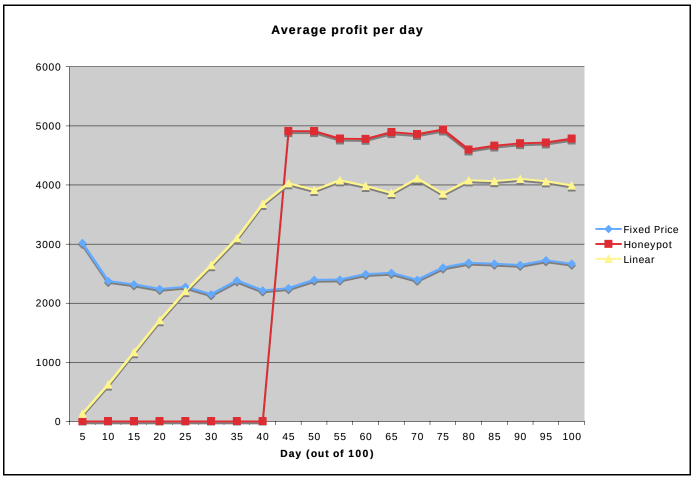

Figure 3. Average profit per day

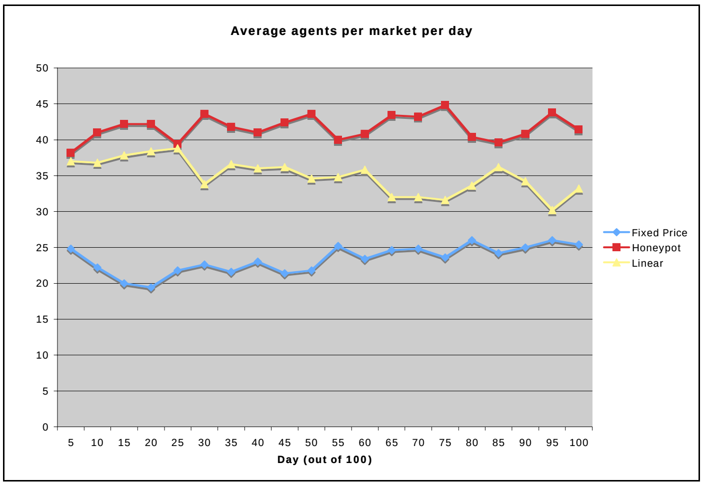

Figure 4. Agents per market per day

The Linear Charging policy was therefore a logical choice to investigate further. In addition to earning a high level of profits and attracting a significant number of agents in the preliminary testing described above, the structure of the competition itself also gave Linear certain advantages. Specifically, the number of days in each distinct game of the competition was both randomly selected and hidden from the competitors; in addition, a random number of days at the beginning and end of each game in the competition were simulated, but profits earned on those days were not counted towards the final score of each specialist.

Both the Linear and Honeypot policies benefit from a number of days at the beginning of each game being dropped from scoring consideration, since neither policy earns very much at the beginning of a game, focusing instead on attracting traders. However, while the "jump" in the Honeypot policy would probably have to be carefully set – setting it too early might drive away traders before they settled into the associated market, setting it too late might leave the specialist with no time to earn profits – the Linear policy constantly ramps up throughout its lifetime, and can therefore operate in a game of any length.

2.2 Effect of Market Pricing Policy

Market Pricing Policy in the testbed sets the clearing price of the goods in the market. The default pricing policy used by the default specialists in the testbed is the k-Pricing Policy, which uses a parameter *k* to set the transaction price between a buyer (bidder) and a seller (asker). More particularly:

$$Clearing\ Price = k \times bid + (1-k) \times ask \qquad \text{(Eq.1)}$$

To set the value of *k*, the default specialists use a Discriminatory Pricing Policy, where the value of *k* was set at 0.5. In practical terms, this means that the difference between what a bidder is willing to pay and what a seller is willing to take is split evenly.

To find out the effect of *k* value in the market, we set the simple preliminary experiments using "Fixed Charging Policy". Specifically, we took twelve different types of trading agents and two fixed charging specialists with the same amount of fees. Each specialist uses discriminatory charging policy with different value of *k*. The experiment featured total 156 trading agents with 47 sellers and 109 buyers. Our assumption is that a specialist can give more profit to the one type of trading agent by varying the value of *k*, thus changes clearing price favor of a specific type of trading agent.

The first specialist sets the value of *k* as 0.2 and the second specialist sets the value of *k* as 0.5. For example, if ask price is 10 and bid price is 20, clearing price of the first specialist would be 12 and clearing price of the second specialist would be 15. Since seller can earn more profit with clearing price of 15 than 12, we expected to see more sellers attracted to the second specialist.

Figure 5 shows the profit earned by each type of specialist per 5 iterations of experiment. From the simple preliminary experiment shown in Figure 5, it is clear that our assumption on the effect of the clearing price was proved to be correct. The second specialist with higher *k* value earned more than the first specialist by attracting more sellers thus increasing more trade in the market.

In the next sections, the charging policy and pricing policy were proposed based on the observation made during the preliminary experiments.

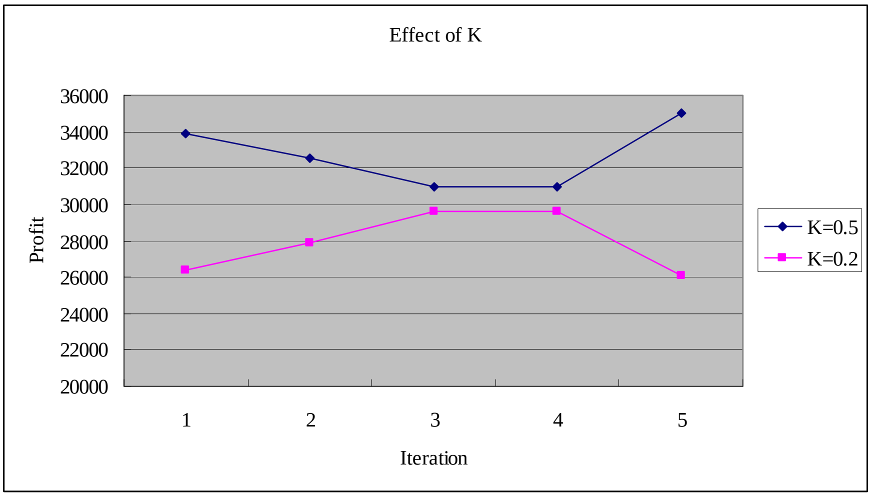

Figure 5. Effect of *k* in the market with more buyers

3 Market Design Competition Strategies

As stated in the previous section, our overall market design strategy focused primarily on two areas of control: the fees the market charged agents, and the price point at which trades between agents were executed. Although the market design competition allowed for specialists to control when trades occurred and what bids were accepted by the market, we could not immediately see any obvious improvements to make on the default strategies provides.

3.1 Linear Charging Policy

As described in the previous section, the Linear Charging policy simply raised fees at a constant rate as the game progressed. The Linear policy started at 0 for all fees, raised fees so that by the end of a set period the various fees would reach a given level (the "set point"), and continued to raise fees until a maximum cap was reached, at which point the fees would become constant. In the context of the student competition and in the experiments which follow, since the specialist is not able to know the length of the game, the Linear Charging strategy raised fees to the set point within 40 days[*].

In accordance with the results suggested by Figure 1 above, we found in preliminary experiments that Linear worked best when the set point was approximately 75% of the average of the fees being charged by all other markets in the competition. However, real fee values that this 75% level translated to was guessed at by the designer, because a fully accurate one-shot determination of average fees for the other specialists at the beginning of a game was extremely difficult, and because using a running average of the fees charged by other specialists would have warped the linear ramp-up of fees. For example, during a prototype competition between the various student-designed markets, the clear winner used a Fixed Price strategy with all fees set to the maximum level allowed; other markets set their fees much lower. Accordingly, we decided a set point of 50% of the maximum fees would work well during the competition.

3.2 Adaptive k-Pricing Policy

The types of traders in the market design competition are exclusively buyers or sellers; buyers cannot sell what they have purchased, vice versa. Therefore, trading agents have a greater chance to complete transactions when the ratio of buyers and sellers is approximately equal – when the opportunity for each agent to find a complementary partner is maximized. Since there is possibility to have more of one type over another type, it is important to attract the minority type of agent to the market to promote trading between seller and buyer. For example, if there is less number of sellers in the system, a specialist might want to have as much sellers as possible to increase the number of trade in the market. In this sense, one obvious improvement to the k-Pricing Policy is to set the value of k according to whether the market wishes to attract more buyers or more sellers.

For example, a scenario where buyers are more numerous than sellers may lead to several sellers being able to find a buyer for their goods or services even at high price. By setting the value of k to set the transaction price to favor sellers, we hope to create a market more attractive to sellers in the future, thereby achieving parity between buyers and sellers. In addition, setting the clearing price based on demand and supply model is one of the fundamental concepts of economics. In the Adaptive Pricing Policy (AP policy) used by our strategy, specifically, the value of k is set based on the ratio of buyers to sellers and proportion of "asks" (attempts to sell) to "bids" (attempts to buy) per day (Eq. 2). In Eq.2, when the average bid price per unit (Eq. 4) is equal to the average ask price per unit (Eq. 3) *offsetUnitprice* becomes 0.5 (Eq. 5).

[*] We assumed game length was not available for the specialist during the student competition, which was false assumption.

$$k = \begin{cases} (offsetShout + offsetTrader)/2 & when\ \ offsetUnitprice \neq 0.5 \\ 0.5 & when\ \ offsetUnitprice = 0.5 \end{cases} \quad \text{(Eq. 2)}$$

$$avgUnitprice_{ask} = \text{Price}_{ask} / Quantity_{ask} \quad \text{(Eq. 3)}$$

$$avgUnitprice_{bid} = \text{Price}_{bid} / Quantity_{bid} \quad \text{(Eq. 4)}$$

$$offsetUnitprice = avgUnitprice_{bid} \Big/ \left(avgUnitprice_{bid} + avgUnitprice_{ask} \right) \quad \text{(Eq. 5)}$$

$$offsetShout = Count_{bid} \Big/ \left(Count_{bid} + Count_{ask} \right) \quad \text{(Eq. 6)}$$

$$offsetTrader = Count_{buyer} \Big/ \left(Count_{buyer} + Count_{seller} \right) \quad \text{(Eq. 7)}$$

In this case, the AP policy considers that there is no difference between asking price and bidding price, thus set the value of k as 0.5. Otherwise, AP policy calculates the proportion of both asks to bids (Eq. 6) and registered sellers to registered buyers (Eq. 7) in the market per day.

As previously described, the AP policy changes the value of k based on the information available per each day using Eq.2. However, individual day's market condition might not reflect system-wise condition, meaning that there is possibility to have more buyers than sellers even though there are more sellers in the system. Since there is no way to figure out how many sellers or buyers exist in the system, the AP policy tracks the trends in its own market. AP policy tracks previous history on the value of *k* each day. If there is more than 10 days of same trends, for example, more buyers over 10 days, then the AP policy adjust the value of *k* by 0.1. In other word, the AP policy adds 0.1 to the calculated the value of *k* when there are more sellers over more than 10 days, and subtracts 0.1 from the calculated the value of *k* when there are more buyers over more than 10 days. The number of day used here (10 days) could be adjusted based on the length of the game if game length information were available for the specialist.

3.3 Alternative Charging Policies

Because other adaptations to the default pricing policy were not immediately apparent, other development work on alternative market specialists focused primarily on the charging policy. Specifically, we developed at least five other fee strategies, three of which we describe here and use in our experiments below.

The Average charging policy follows directly from the results of Figure 1 above; the policy simply figures out the average fees being charged by the other markets in the competition on any given day, and sets its own fees to approximately 75% of that average. Note that because each market was required to post its fees for each day at the same time, the Average charging policy was actually averaging the previous day's fees for each competing market; however, it was felt in the aggregate that this would not significantly affect the strategy.

The Attractor policy was a slight variation on the Average Charging policy. Experiments suggested that the fees earned by a market as the number of agents increased did not increase in a linear fashion; instead, a high number of agents within a given market created a kind of critical mass wherein the increased opportunity for agents to find trading partners not only generated significantly more transactions (and therefore fees for the market) but created an attractive environment for agents to return to in the future. The Attractor policy attempted to take advantage of this phenomenon by dynamically changing its fees depending on both the average fees charged by other markets and the number of agents registered to the market at any given time. Like the Average charging policy, the Attractor market attempted to keep its fees at 75% of the average charged by other markets. However, when the number of agents in the Attractor market reached a certain set limit (roughly 2 * (total traders/total markets)) the profit fee began to scale rapidly as the number of agents increased, such that each additional agent past the set limit raised the profit fee by approximately 5 percent.

The Sawtooth Average charging policy was cross between the Average Charging policy and a modified version of the Linear Charging policy. Like the Average charging policy, the Sawtooth policy determined an average of the fees charged by competing markets each day. Furthermore, rather than raising fees indefinitely until the maximum cap was reached, the Sawtooth charging policy was set to raise fees until they reached 200% of the average, and then lower fees at a constant rate until they reached zero, at which point the cycle would begin again. Various periods for the Sawtooth policy were tried in preliminary experiments, and the end of which we remained with a period of 40 days for the complete zero fee-to-zero fee cycle.

In addition to the Linear, Average, Sawtooth Average and Attractor policies described above, the experiments below also feature two markets with Fixed Price fee policies, one where the fees were set permanently at the maximum level allowed, and one where no fees were charged.

4 Experiment Design and Analysis

Experiments were conducted to show how the proposed Linear charging and Adaptive pricing policy performs in different market settings. In each experiment, we examined the specific effect of each policy on the specialist's profits to find correlations between the policies and resulting profits given various market situations.

4.1 Experiment Design

We conducted four distinct sets of experiments to examine the relative utility of the Linear charging and Adaptive pricing policies in the market design competition. Each set of experiments further divided into two subsets: one where we deployed a combination of the Adaptive Pricing policy and Linear Charging policy (adaptive-linear) and another where we deployed the Linear Charging policy with the Discriminatory pricing policy (discriminatory-linear).

As shown in Table 1 below, the competing specialists exclusively used the Discriminatory pricing policy in combination with the Attractor, Average, Sawtooth average, Fixed with maximum fess, and Fixed with zero Fees. Traders were a mixture of the GD, ZIP, and Random Constrained market selection strategies with both Softmax and Epsilon-Greedy Learner. Each experiment was executed for 10 trials, and a significance test with level of 0.05 was conducted to the results to see if two set of results are statistically significant.

The first experiment (experiment 1) was conducted with randomly assigned parameters as shown in the table 1 to reflect the real competition environment, however, the parameters of other experiments (experiment 2, 3, and 4) was intentionally set to see how the proportion of sellers to buyers effect the outcomes of each specialist. Experiment 2 was specified to see the case of equal number of sellers and buyers in the system and Experiment 3 and 4 were specified to investigate the effect of inequality in the proportion of sellers and buyers.

Additional experiments controlling the types of trader were conducted to see the effect of each types of trader on the proposed policies. GD-Epsilon, ZIP-Epsilon, and Random-constrained-Softmax traders were used in the additional experiments. The result from the experiments can be found in the Appendix section.

	Experiment 1 (random)	Experiment 2 (equal)	Experiment 3 (more sellers)	Experiment 4 (more buyers)
Specialist	Discriminatory Pricing Policy / Adaptive Pricing Policy Linear Charging Policy			
Competing Specialists	Charging Policy: Discriminatory Pricing policy: Attractor, Average, Sawtooth, Fixed with Maximum Fee, Fixed with Zero Fee			
Sellers	ZIP[†]-EG[‡] (.37)=14[§] ZIP-SM (.14)=40 GD-EG (.21)=3 GD-SM (.12)=1 RC-EG (.28)=3 RC-SM (.27)=9 Total: 70	ZIP-EG (.37)=8 ZIP-SM (.14)=8 GD-EG (.21)=8 GD-SM (.12)=8 RC-EG (.28)=8 RC-SM (.27)=8 Total: 48	ZIP-EG (.23)=18 ZIP-SM (.22)=23 GD-EG (.18)=7 GD-SM (.20)=7 RC-EG (.07)=36 RC-SM (.27)=41 Total: 132	ZIP-EG (.26)=7 ZIP-SM (.25)=5 GD-EG (.08)=12 GD-SM (.12)=7 RC-EG (.06)=8 RC-SM (.15)=8 Total: 47
Buyers	ZIP-EG (.33)=7 ZIP-SM (.24)=7 GD-EG (.29)=1 GD-SM(.23)=38 RC-EG (.30)=2 RC-SM (.14)=3 Total: 58	ZIP-EG (.37)=8 ZIP-SM (.14)=8 GD-EG (.21)=8 GD-SM (.12)=8 RC-EG (.28)=8 RC-SM (.27)=8 Total: 48	ZIP-EG (.29)=2 ZIP-SM (.10)=2 GD-EG (.26)=8 GD-SM (.22)=4 RC-EG (.26)=35 RC-SM(.16)=2 Total: 52	ZIP-EG (.21)=8 ZIP-SM (.16)=14 GD-EG (.24)=39 GD-SM (.21)=1 8 RC-EG (.16)=14 RC-SM (.23)=16 Total: 109
Game Setting	Game Length: 107 Days / Day Length: 5 rounds / Round Length: 1000			
Critical period	Profit recorded between 16th day and 91st day			

Table 1. Parameters for the Experiment

4.1 Experiment Analysis

Experiment 1 (Random). The objective of Experiment 1 is to show how the proposed adaptive-linear specialist performs in the various environments involving many different types of trader and specialist.

More particularly, Figure 7 shows the profit earned by each specialist during the Experiment 1, where the number of buyers and sellers was randomly selected. Unsurprisingly, the Linear pricing policy out-competes all other strategies (statistically significant) during the critical period during the middle of the game, regardless of whether the Adaptive pricing policy or not. However, the Linear pricing policy was, from a statistical standpoint, in a tie with the Sawtooth and Average pricing policies. This can almost certainly be explained by the fact that Sawtooth and Average earn significant amounts of profit in the beginning of the game, when Linear is slowly ramping up fees. The right side of Figure 6 shows the average profit earned by each specialist during the entire game. The average profit of discriminatory-linear specialist was a bit less than average policy specialist; whereas, the average profit of the adaptive-linear specialist was higher than average policy specialist. The result was not statistically significant; however, it shows the possibility of adaptive pricing policy was playing a role in the performance of the specialist.

[†] ZIP (Zero Intelligent Plus), GD (Gjerstad Dickhaut), RC (Random Constrained)

[‡] EG (Epsilon Greedy), SM (Softmax)

[§] ZIP-EG (.37) = 14 → Strategy-Learner (epsilon value or temperature value) = number of trader

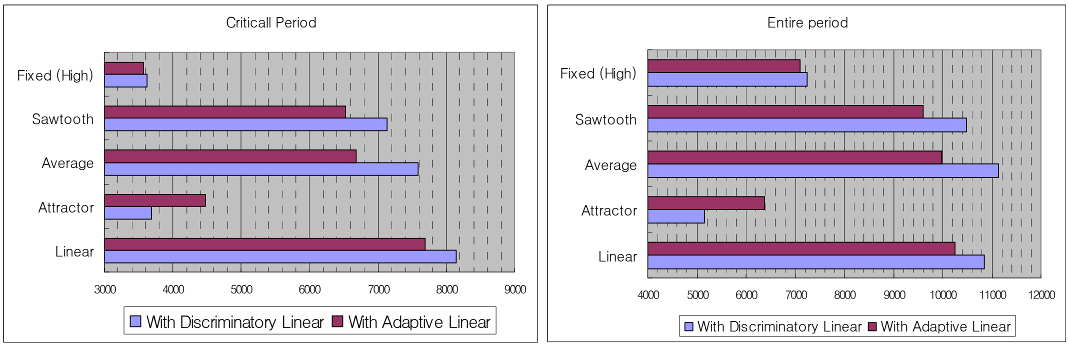

Figure 6. Average Profit of each specialist (Experiment 1)

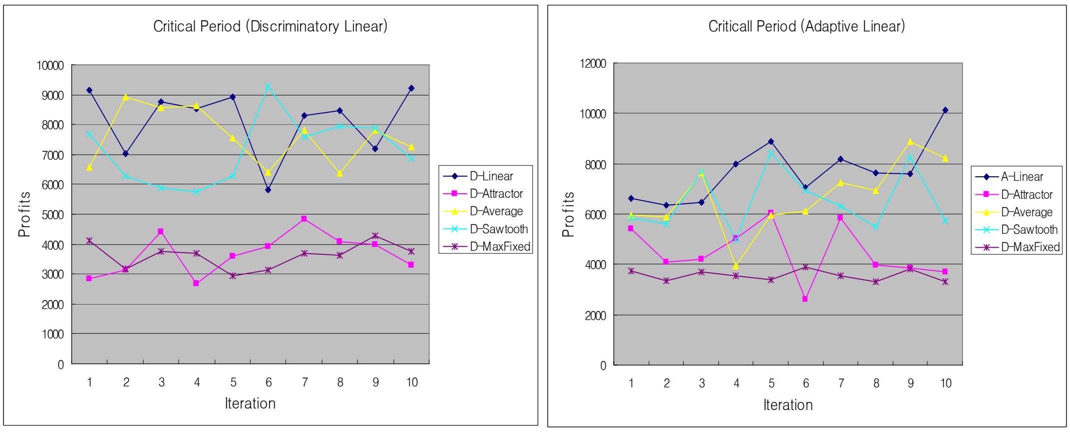

Figure 7. Specialist Profits in Experiment 1

Experiment 2 (same number). The objective of Experiment 2 is to show how the proposed adaptive-linear specialist performs in the market where there is the same number of buyers and sellers. Figure 8 and Figure 9 shows the results of Experiment 2, where there is equal number of seller and buyer. When paired with the Adaptive pricing policy, Linear performs statistically significantly better than other strategies during the critical period at the middle of the game (Right chart at Figure 9). However, Linear without the Adaptive pricing policy does not perform statistically significantly better than the Average pricing policy (Left chart at Figure 9). However, it is possible that a higher number of trials would have revealed a significant difference between the two.

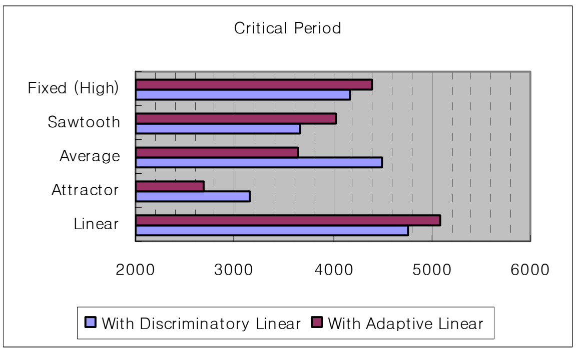

Figure 8. Average Profit of each specialist (Experiment 2)

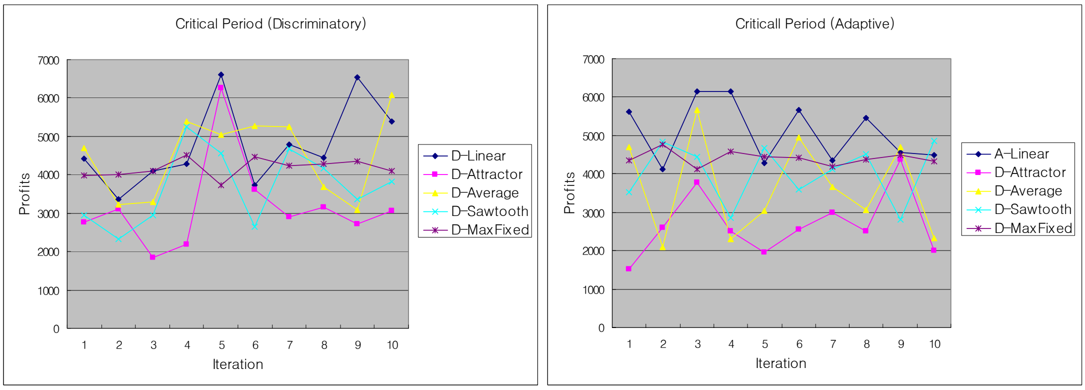

Figure 9. Specialist Profits in Experiment 2 (Same Number)

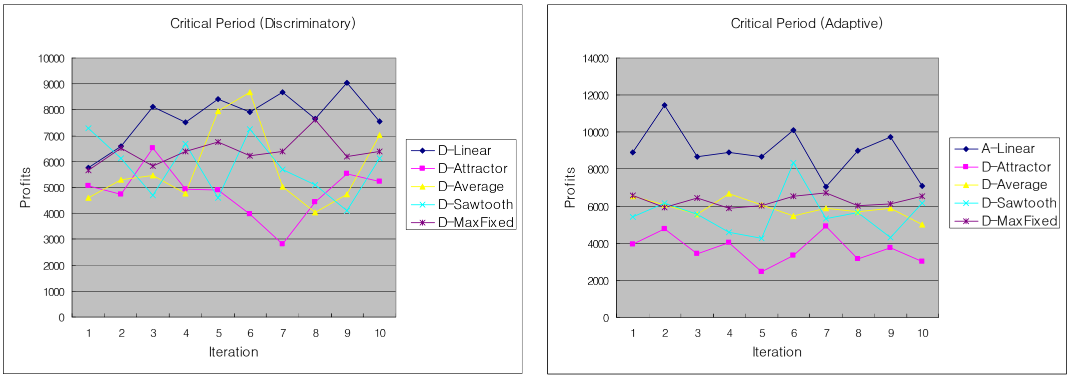

Figure 10. Specialist Profits in Experiment 3 (More Seller)

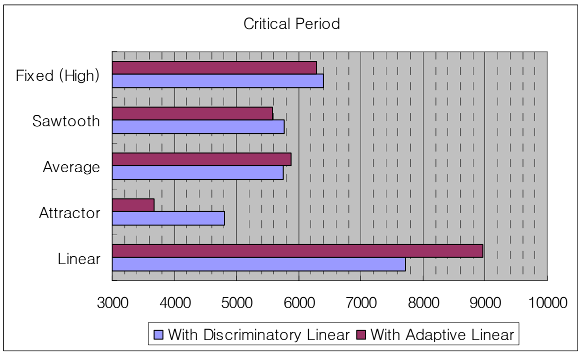

Figure 11. Average Profit of each specialist (Experiment 3: more sellers)

Experiment 3 (more seller) and Experiment 4(more buyer). Similar results can be seen in Experiment 3 and 4, where, when paired with the Adaptive policy, the Linear policy is clearly (and statistically) superior to the other strategies during the critical period during the middle of the game (top-right chart in Figure 10, and Figure 12), and less so (top-left chart in Figure 10 and Figure 11) or not at all (Figure 12) without the Adaptive policy. Interestingly, Linear without Adaptive fails to clearly distinguish itself in both Experiment 2, where the number of buyers and sellers was equal, and in Experiment 4, where the number of sellers outstripped the number of buyers. This seems to imply that the value of the Adaptive pricing policy is not that it rectifies system-wide imbalances in the number of buyers and sellers, but that it helps to prevent such imbalances occurring within a given market, regardless of the greater balance between buyers and sellers. Future research in strategies for the TAC Market Design competition might benefit from examining how often such imbalances between buyers and sellers within a market occur, and what further policies can be introduced to better prevent such imbalances.

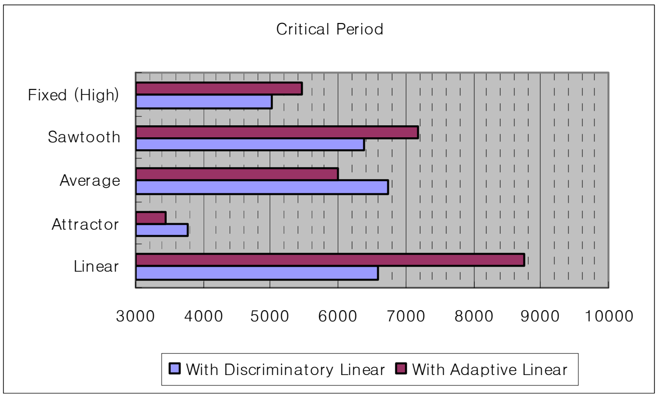

Figure 12. Average Profit of each specialist (Experiment 4: more buyers)

5 Conclusion

The combination of Linear Charging and Adaptive pricing seems to provide a strategy which is clearly superior to the others observed, at least within the context of the critical period used for the student competition. Given that the Adaptive policy, when used, gave a statistically significant victory in every experiment, but that Linear by itself only achieved a statistically significant victory in two out of the four experiments, it seems reasonable to conclude that the Adaptive pricing policy plays a more important role in making our strategy a success. Especially given that most of our efforts thus far have been dedicated to exploring the fee policy space, further research into other policy areas, especially pricing policy, would seem warranted.

It should also be noted that the agents used in the current (very early) version of the Market Design testbed are far from fully rational. This is indicated by both the results of Figure 2, where the zero-fee market did not show a decisive advantage in traders attracted over other, far more expensive markets, and by anecdotal evidence which suggests that traders in the testbed constantly lose significant amounts of money, even if a “no play” option is available. Although the strategy outlined above works well with the current testbed, it seems probable that it would not necessarily work so well with the increasingly rational agents that are sure to be developed as the testbed software matures.

References

1. M. P. Wellman, A. R. Greenwald, P. Stone, and P. R. Wurman, "The 2001 Trading Agent Competition," *Electronic Markets*, vol. 13, no. 1, 2002.
2. Raghu Arunachalam and Norman M. Sadeh. The supply chain trading agent competition. Electronic Commerce Research and Applications, 4:63–81, 2005.
3. Enrico Gerding, Nicholas R. Jennings, Peter McBurney, and Steve Phelps. Outline Proposal for a New Trading Agent Competition: TAC Market Design.

Appendix: Additional Experiment Results

A.1 ZIP with Epsilon-Greedy

Experiment	Experiment A1-1 (Same Number)	Experiment A1-2 (More Seller)
Specialist	Discriminatory / Adaptive Pricing Policy	Discriminatory / Adaptive Pricing Policy
	Linear Charging Policy	Linear Charging Policy
Seller	ZIP / Epsilon (0.1) = 50 Total Number: 50	ZIP / Epsilon (0.1) = 75 Total Number: 75
Buyer	ZIP / Epsilon (0.1) = 50 Total Number: 50	ZIP / Epsilon (0.1) = 25 Total Number: 25

Table A1. ZIP with Epsilon-Greedy Learner

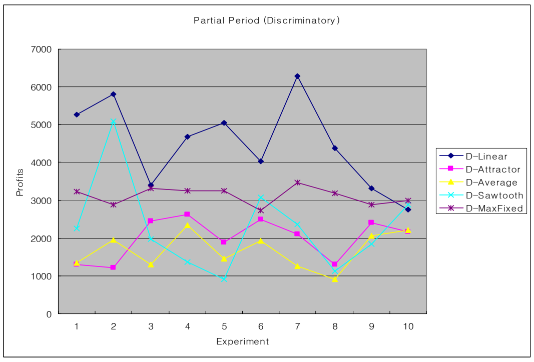

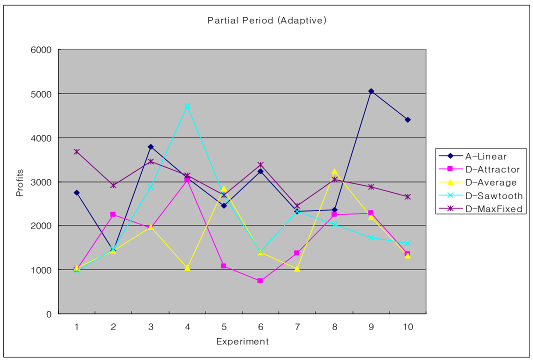

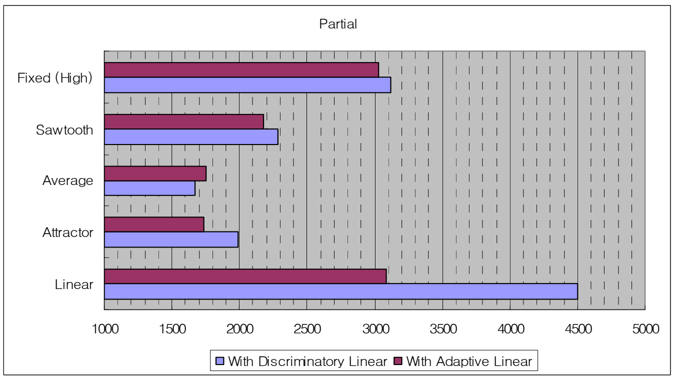

Figure A1. Experiment A1-1: ZIP-Same Number

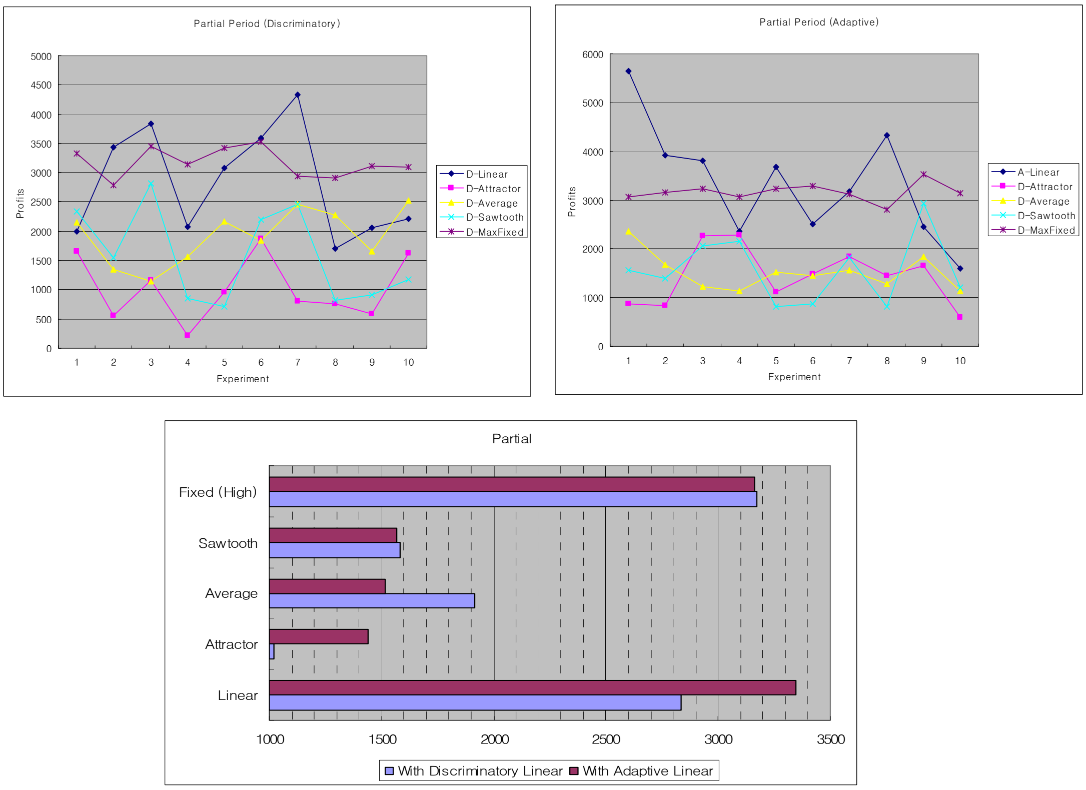

Figure A2. Experiment A1-2: ZIP-More Sellers

A.2 GD with Epsilon-Greedy

Experiment	Experiment A2 (Same Number)
Specialist	Discriminatory / Adaptive Pricing Policy
	Linear Charging Policy
Seller	GD / Epsilon (0.1) = 50 Total Number: 50
Buyer	GD / Epsilon (0.1) = 50 Total Number: 50

Table A2. GD with Epsilon-Greedy Learner

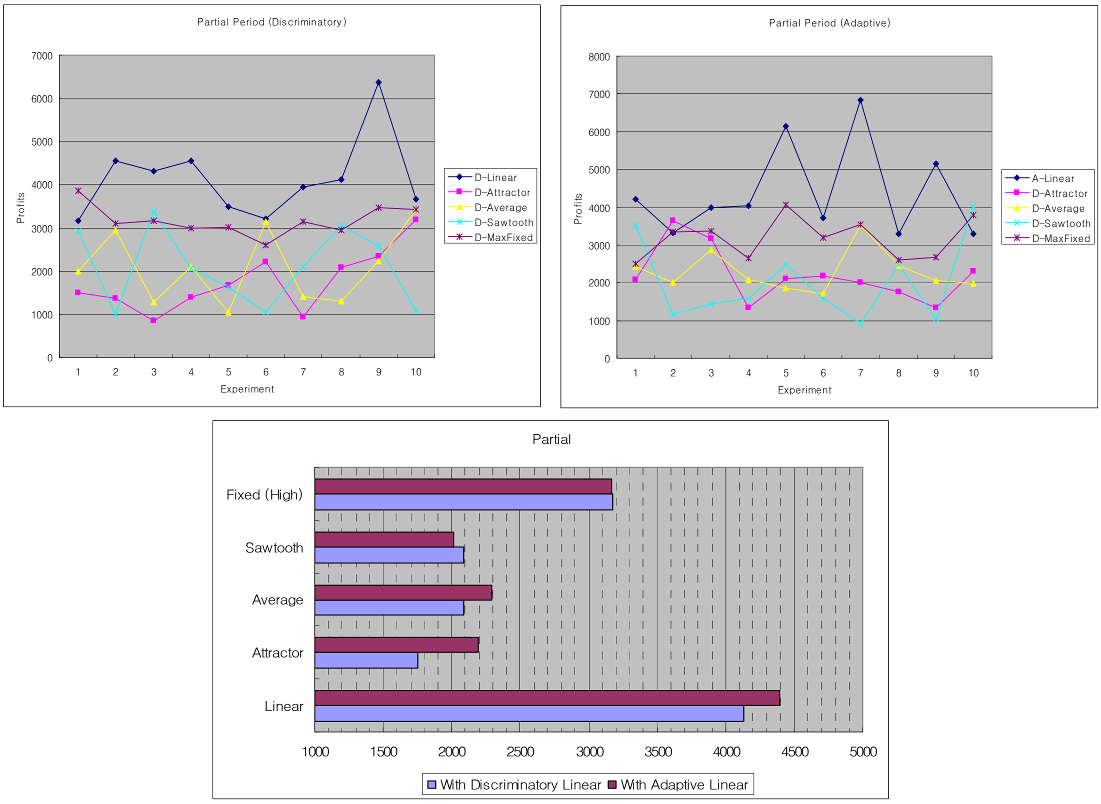

Figure A3. Experiment A2: GD-Same Number

A.3 Random Constrained with Softmax

Experiment	Experiment A3-1 (Same Number)	Experiment A3-2 (More Seller)
Specialist	Discriminatory / Adaptive Pricing Policy	Discriminatory / Adaptive Pricing Policy
	Linear Charging Policy	Linear Charging Policy
Seller	Random Constrained / Softmax (0.1) = 50 Total Number: 50	Random Constrained / Softmax (0.1) = 75 Total Number: 75
Buyer	Random Constrained / Softmax (0.1) = 50 Total Number: 50	Random Constrained / Softmax (0.1) = 25 Total Number: 25

Table A3. Random Constrained with Softmax Learner

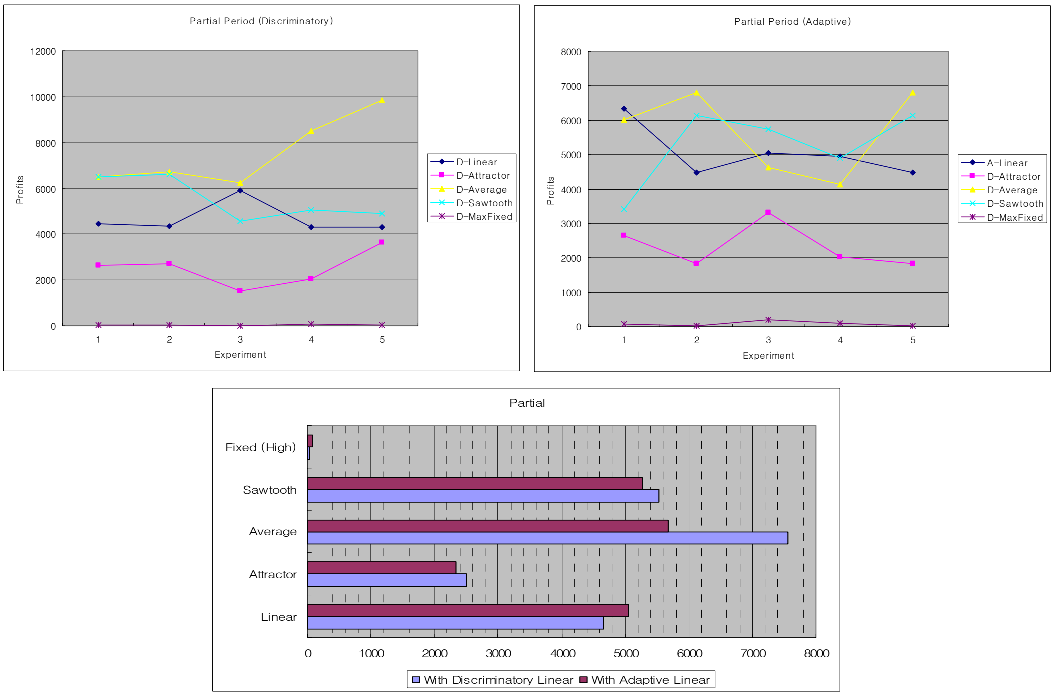

Figure A4. Random Constrained with Softmax (Experiment A3-1: Same Number)

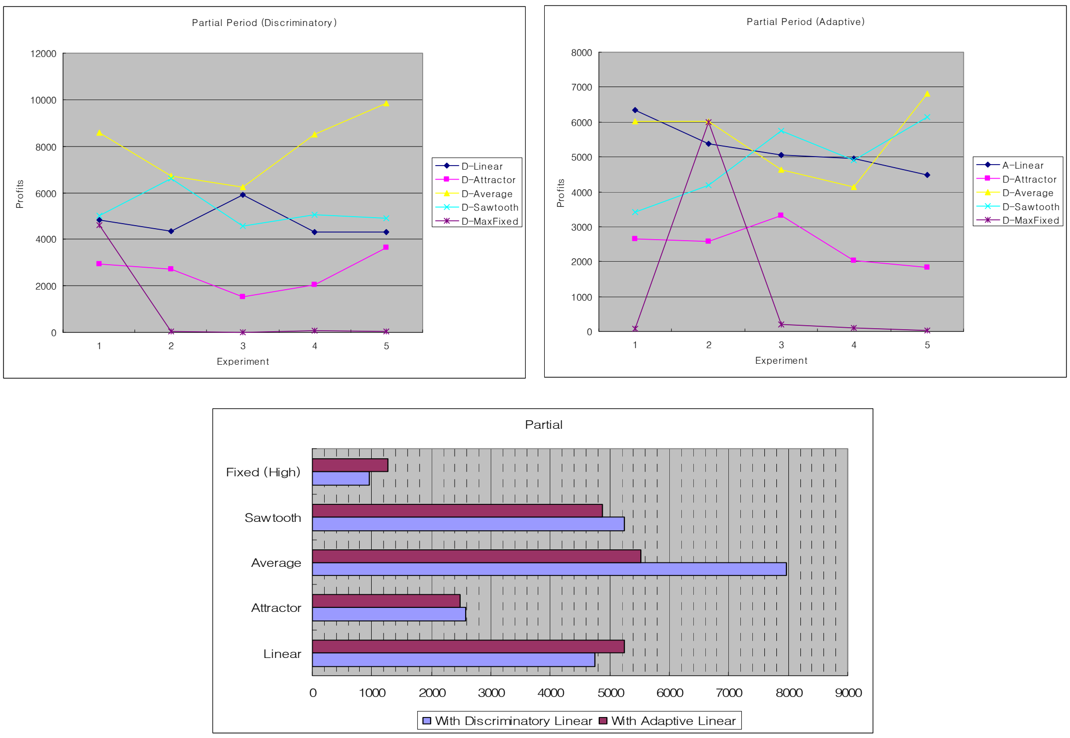

Figure A5. Random Constrained with Softmax (Experiment A3-2: More Seller)