Adverse Selection Plus Strategic Advantageous Selection Equals Trade

Hao Cen [1] and John R. O'Brien [2]

January 2009

Abstract

In this paper we apply a learning model from machine learning, to a human trading crowd to understand why the no trade theorem was rejected. Our results reveal that trading volume in a continuous double auction market is associated with inverse learning curves. Inverse learning results from adverse selection among market takers and strategic advantageous selection among market makers. In contrast to associating adverse selection with market failure in traditional competitive market theory, the rules of a double auction market efficiently exploit individual differences among the trading crowd to generate both large amounts of trading volume and relatively efficient prices.

---

[1] Hao Cen, Machine Learning Department, School of Computer Science, Carnegie Mellon University, Email: hcen@andrew.cmu.edu

[2] John R. O 'Brien Tepper School of Business, Carnegie Mellon University Email: jo0x@andrew.cmu.edu

1. Introduction

Trading volume in financial markets is large and difficult to explain within the context of a classical competitive market characterized by price taking behavior, rational expectations and no trade theorems, e.g. Tirole (1982). The financial economics literature responded to unintuitive no trade theorems by introducing some form of bounded rationality. For example, in finance "liquidity" (i.e., noise) traders (e.g., Black (1986); Hellwig (1980); Kyle (November 1985)) are introduced to resurrect the intuition that individual differences for processing information, leads to trading. Subsequently the behavioral finance literature has focused upon cognitive biases and information processing, motivated by the work of Tversky and Kahneman (1973) to provide plausible explanations of observed financial market phenomena[3]. However, as observed by Brav, Heaton and Rosenberg (2004).

"The most important philosophical concern facing financial economists today is the contest between rational and behavioral finance. Researchers in the rational paradigm assert that behavioral models employ ambiguous assumptions of irrationality undisciplined by rigorous mathematics ….. Researchers in the behavioral paradigm criticize the failure of rational finance to generate meaningful predictive successes and…. fail to identify measurable economic variables."

What makes this a fascinating philosophical concern is that both sides make positive contributions. The rational finance model has provided important and substantive insights into the equilibrium properties of financial markets. However, absence some type of bounded rationality it fails to provide a realistic description of the dynamic properties of price discovery

---

[3] Recent books provide a comprehensive introduction to this area e.g., Schleifer (April 2000); Shefrin (2007); Shiller (May 2006)

which in turn can impact upon equilibrium behavior. For example, Allan Greenspan, former Federal Reserve Bank Chairman, at a congressional hearing on Capitol Hill on Thursday October 23, 2008, conceded that he had put too much faith in the self correcting power of free markets (Andrews (October 23, 2008)).

Behavioral finance, on the other hand, provides sharper insights into the nature of the bounded rationality which drives trading and influences the success or otherwise of other market phenomena. In this paper we study the interaction between rational and behavioral finance in a market setting conducted under controlled conditions using a relative large trading crowd[4]. In the markets studied the no trade theorem was strongly rejected even though equilibrium price predictions were ultimately supported. As a result, this setting generated behavior that requires both the rational and behavioral finance approaches to describe how markets discover prices. In this paper we study how the market discovers prices by applying a learning model developed from the field of machine learning. This model provides a new method for studying behavioral finance and permits identifying more sharply the systematic behavior that led to the rejection of the no trade theorem and the spontaneous generation of noise. To our knowledge this is the first time the behavior of the trading crowd has been studied at this level of behavioral detail.

Our results reveal the unusual finding that efficient price discovery was associated with "negative learning." That is, contrary to a traditional single person learning environment where repeated practice improves performance we observe the opposite, and which resulted in efficient prices. We identify two fundamental drivers that underlie this result. First, the trading crowd

---

[4] The markets were conducted as a sponsored trading competition among the first and second year MSCF students at Carnegie Mellon University. Traders were competing for USD$2,500 prize money plus the sponsoring bank hired the top trader. The markets were conducted on the Financial Trading System (http://www.ftsweb.com) platform over the web between the New York and Pittsburgh campuses. There were 92 traders in the trading crowd and hundreds of thousands of transactions were executed.

exhibit individual differences in abilities to assess the intrinsic value of the securities traded which when combined with the price taking rules of the double auction market promoted a classical adverse selection problem among market takers. That is, Ackerloff's "market for lemons"(1970) was an important driver of trading volume generated from market takers as they learned to assess more accurately the intrinsic value of the securities[5]. Second, market makers who get to post the bid/ask spread, also exhibited negative learning but for different reasons. This subset of the trading crowd exhibits the opposite to the "lemons phenomena" which we refer to as "advantageous selection," because in this case the learning was associated with mastering the strategic price setting behavior. These results point to the importance of individual differences within the trading crowd constrained by the market's microstructure, for describing market dynamics.

The remainder of the paper is organized as follows. Section 2 describes the markets, section 3 the rules that govern trading in the markets studied. Section 4 then introduces the learning model we use to study the behavior of the trading crowd. Section 5 then develops the hypotheses and section 6 our methodology and results. Finally, section 7 provides some discussion and conclusions.

---

[5] The seminal paper on Adverse Selection is by Akerlof, George A. (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism". *Quarterly Journal of Economics* **84** (3): 488–500 for which he was later awarded the Nobel Prize.

2. The Trading Competition

2.1 The Market Data

The data under analysis is from the trading competition sponsored by a major bank conducted at Carnegie Mellon University. The trading crowd consists of both first and second year MSCF students who were competing for the $2,500 prize money[6]. The markets were conducted over the web using the Financial Trading System's platform[7] with price discovery using a trading crowd that consisted of 92 students enrolled in both the Pittsburgh and New York campuses of Carnegie Mellon[8]. The market comprised of four independent and identical trials. Each trial covers three years of calendar time and trading is broken up into three trading periods. At the beginning of each trial the trading crowd faces interest rate uncertainty and there are eighty one possible interest rate paths. Over the first two trading periods news is randomly and publicly released such that by period three there is a single implied interest rate path. Each trading period lasts for 300 seconds, and because there is no additional news in period three traders can form a unified expectation with respect to the intrinsic value of each security. As a result, in period 3 the strong no trade theorem is predicted to hold under the rational finance model. In period 3, two fixed-income security markets remain open referred to as security 1 and security 4. Security 1 is a three-year coupon bond with face value of $100 that pays a coupon payment at the end of each period and both a face value and the final coupon payment at the end of period 3. Security 4, is a 3-year zero-coupon bond that pays a $100 face value at the end of period 3.

---

[6] The sponsoring bank hired the top trader to their trading desk.

[7] These were conducted using the FTS interactive markets where the students were the trading crowd from which all price discovery originated from (http://www.ftsweb.com).

[8] The Pittsburgh students were full time and New York students were primarily part time usually from Wall Street.

Under the rational finance model the trading crowd is predicted to discover prices that equal the intrinsic value independently of individual preferences. This is because a rational trader will bid to buy at prices up to the point intrinsic value and ask to sell at prices down to the intrinsic value. If every trader is rational, no trades will happen because if every trader is quoting around true price, the bid ask spread will envelope the true price and no one is willing to buy or sell to avoid the opportunity cost associated with the trade[9]. However, we observed wide price fluctuations and huge trading volumes going on with the two securities in the last session as illustrated in Figures 1 and 2 respectively.

Insert Figures 1 and 2 about here.

The research question is how to explain this trading volume and the dynamic price discovery process. The two security markets depicted represent two types of markets. Security 1 ultimately settles down to prices predicted from the rational competitive market theory but the trading volume is inconsistent with this theory. On the other hand the rational competitive market theory fails to describe both prices and trading volume for the security 4 market depicted in figures 1 and 2. Later we will see that this results from strategic as opposed to non strategic price taking market behavior. To understand this distinction in the next section we first describe the rules that govern trading in these markets. This is referred to as the market's microstructure.

---

[9] Technically, the money market rate (the realized interest rate path) is the opportunity cost facing the trading crowd. As a result, anyone buying at greater than the intrinsic value is giving up more by withdrawing cash from the money market than the yield to maturity they would receive from buying the fixed income security. Similarly, any trader selling below the intrinsic value is forgoing a greater yield to maturity from holding the fixed income security relative to what they receive by investing the proceeds from the sale in the money market.

3. Market Microstructure and Two Kinds of Selection

3.1 Continuous Double Auction Market

A continuous double auction market is composed of both market makers and market takers. Market makers post bids to buy and asks to sell up to some specified quantity. Market takers can sell to a bid or buy from an ask. A set of rules govern the allowable trading behavior for these two groups which is referred to as the market's microstructure. For market markers, two rules are most relevant here. The price improvement rule says that bids must be increasing and asks must be decreasing. The crossing rule says if the bid cross the ask or vice versa (i.e., Bid >= Ask or Ask <= Bid) then the order is processed at the minimum quantity (i.e., Min(Bid Quantity, Ask Quantity). For market takers, the "first in first served" rule is most relevant here. This rule says that the first market taker to submit a market order (i.e., sell to a bid or buy from the ask) gets the trade.

3.2 Adverse Selection versus Advantageous Selection

The learning problem facing the trading crowd is price discovery (i.e., learning to assess the intrinsic value). The market taking rule will promote adverse selection in the sense that the market takers who assess intrinsic value with the most error will accept a bid or ask first. On the other side of the market the market making rules (price improvement and crossing) will tend to promote advantageous selection because the price improvement rule combined with crossing rewards the most competitive market makers. Advantageous selection is the opposite to adverse selection, in the sense that the rules of market promote the survival of the most skilled market makers.

In a price taking competitive market, the market makers who are better able to assess the intrinsic value will remain as the active bids and asks in the market, under the price improvement and crossing rules. If the competitive price taking behavior assumption is relaxed in lieu of strategic market making behavior, then the learning problem facing this subset of the trading crowd changes. For this latter case, market makers are learning how to generate profits from market liquidity relative to the posted bid/ask spread. In this case either the bid or ask can systematically vary from the predicted intrinsic value if liquidity is available in the market. For example, see security 4 in figures 1 and 2 as providing an example of strategic market making behavior. Once again, the price improvement and crossing rules will promote advantageous selection among market makers for the case where the bid/ask spread is treated as a strategic variable.

4. The Learning Model

4.1 Motivation

The learning model (the Additive Factor Model) by Cen, Koedinger and Junker (2006) introduced in this section was used to model human performance in intelligent tutoring systems (a computer program with artificial intelligence to teach students). We applied this model to capture how the trading crowd learns the intrinsic value of the securities via trading in a continuous double auction market. Each trader's learning is inferred from whether she has correctly executed appropriate trading strategies, as either a market maker or a market taker, when repeated opportunities arise. Here we define "correct" relative to the intrinsic value of a security as defined by the rational finance model.

4.2 The Additive Factor Model (AFM)

This model also has a connection with Item Response Theory Linden and Hambleton (1997), the standard theory in intelligent tests. It starts with four assumptions that we adapted to the trading scenario as follows.

1. Different traders may initially know more or less. Thus, we use an *intercept* parameter (θ_i in equation 1.1) for each trader.

2. Traders learn at the same rate. Thus, *slope* parameters (γ_j in equation 1.1) do not depend on student. This is a simplifying assumption to reduce the number of parameters in equation 2. Importantly, it provides a description of the marginal trader learning behavior which is important to any concept of a competitive equilibrium in economics.

3. Some skills are more likely to be relatively easy or difficult to apply. That is, more skilled traders will find it relatively easier to apply than will less skilled traders. The parameters (β_j in

equation 1.1) for each production, provide a description of the marginal trader behavior for this dimension of the trading problem.

4. Some skills are easier to learn than others. Thus, we need a *slope* parameter for each skill. This is represented by the slope parameter being estimated relative to each of the skills in equation 1.1.

Based on the assumptions, AFM takes the form depicted by equation.

$$\log\left(\frac{P_{ijt}}{1-P_{ijt}}\right)=\sum\theta_i X_i+\sum\beta_j Y_j+\sum\gamma_j Y_jT_{jt} \tag{1.1}$$

where

P_{ijt} = the probability of the i^th^ trader applying the j^th^ skill right on its t^th^ opportunity

X_i = the i^{th} trader

Y_j = the j^{th} skill

T_{jt} = the t^{th} opportunities practiced on the j^{th} skill

Y_jT_{jt} = the interaction between a skill and its number of practice opportunities

θ_i = the i^{th} trader's prior knowledge

β_j = the easiness/difficulty of applying the j^{th} skill

γ_j = the learning rate of the j^{th} skill

Formally, this model captures that the probability for trader i to get trade j right relative to the rational finance model, is proportional to the trader's prior knowledge plus "easiness" of applying the skill and the learning of the intrinsic value through time.

4.3 Relating the Learning Model to Trading

For our settings, we defined the skills a trader may apply under the rules governing exchange in the market. For the two securities in the trading competition, we define two skills for market makers – bidding and asking, and two skills for market takers – buying and selling. A rational description of market making in the competition predicts that a market maker should not bid above an asset's intrinsic value or ask below the intrinsic value. Thus, we define a correct execution of a bid to be bidding at a price no higher than the asset's intrinsic value, and ask to be asking at price no lower than the intrinsic value. Accordingly, we have four market taking skills defined in this study – BidC3, AskC3, BidZ3, and AskZ3.

Similarly, we apply the same rule to buying and selling. According to the common notion of "Buy Low and Sell High", buying is defined to be correctly executed if a trader buys the security when the asset's intrinsic value is above the best ask. Selling is defined to be correctly executed if a trader sells when the asset's intrinsic value is below the best bid. Thus, we have four market taking skills defined in this study – BuyC3, SellC3, BuyZ3, and SellZ3.

5. Hypotheses

The problem facing the trading crowd is price discovery (i.e., learning to assess the intrinsic value). We start first with the traditional economic assumption that behavior in a competitive market is non strategic or "price taking." The market taking rule is predicted to promote *adverse selection* among market takers because the market takers who assess intrinsic value with the most error can accept a bid or ask first. The market making rules (price improvement and crossing) will tend to promote *advantageous selection* among non strategic market makers because those who are better able to assess the intrinsic value are predicted remain as the active bids and asks in the market under the price improvement and crossing rules. Hypotheses 1-3, stated in their null form, apply to a traditional competitive price taking market.

Hypothesis 1 – market makers and market takers have the same prior knowledge, i.e. $\theta_{MM} = \theta_{MT}$

Hypothesis 2 – market makers and market takers have equal skills, i.e. $\beta_{MM} = \beta_{MT}$.

Hypothesis 3 – market makers and market takers have the same learning, i.e. $\gamma_{MM} = \gamma_{MT} = 0$.

6. Methodology and Results

6.1 Model Results

The large trading crowd generated a large number of trades (hundreds of thousands) when trading on the FTS electronic markets which allowed obtaining measurements of important variables using 1-second intervals of time. The learning model was fitted by trial (there were four trials) for each trader and the presence of systematic behavior was tested for. Formally, the dependent variable, the log odds of the probability of the i'th trader applying the j'th skill on it's t'th opportunity correctly is constructed using 1-second intervals of time. That is, time for each trading period is broken up into 1-second intervals and the learning model is estimated from the set of transactions within each 1-second interval.

6.1.1 Binomial Test on Trader Prior Knowledge Coefficients

We first compare the prior information of market makers and market takers across the four trials and the two securities.

Insert Table 1 about here

Table 1 summarizes the results from fitting the learning model[10] and provides the average knowledge parameters by MM/MT, security and trial. Table 1 reveals that in each cell the market making market taking pair was such that the average knowledge parameter was higher for the market maker. This was significant using a sign test at the p-value = 0.007812 level.[11] The results from table 1 provide a qualified rejection of hypothesis 1 and lend support to the intuitive

[10] We fit the model using a ridge logistic regression model fitting algorithm with the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method.

[11] Within each trial, the t-statistic for the mean difference of market maker and market taker prior knowledge did not attain significance given the size of the error among traders.

conjecture that informational differences provide one driver of trading volume. Further the sign test demonstrates that across trials and securities the market makers are more knowledgeable than market takers on average.

6.1.2 Binomial Test on Skill Coefficients

Insert Table 2 about here

Table 2 summarizes results from fitting the learning model and provides the learning and skill coefficients for the thirty two conditions assessed (4 trials times MM versus MT times 4 skills for each MM and MT). The bottom half of the table provides the coefficients for MT's and the top half the coefficients for MM's. First, a pronounced relationship is observed between skill differences for MM (Bid, Ask rows) versus MT (Buy, Sell rows). The skill parameters for MM's are positive and become increasingly positive across trials. The opposite is observed for MT's. The binomial test on the sign difference of the paired market making and market taking skills is highly significant with p-value = 3.052e-05. This rejects the second hypothesis. These results reveal that the market makers exhibit increasing mastery of the skills as indicated by the increasing positive coefficient trend. The market takers on the other hand revealed that the their skills became increasingly more difficult to apply. For MT's these trends are consistent with the presence of adverse selection, the "market for lemons" problem studied by Akerlof[12]. Just as the existence bad used cars drive out the good used cars in the Akerlof market, table 2 reflects that the existence of the less skilled MT's are driving out the more skilled MT's by accepting less

[12] The seminal paper on Adverse Selection is by Akerlof, George A. (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism". *Quarterly Journal of Economics* **84** (3): 488–500.

favorable bids/asks from the MM's. That is, under the price taking rules for a double auction market it is strictly "first in first served" and so when studying the set of active market takers it is the less skilled MT's who will move first. This gives rise to the equivalent of a "lemons" effect as revealed by Table 2 where the skill parameters for MT's are negative and become increasingly negative across the repeated independent trials.

On the MM side of the market the price improvement rule of the double auction institution provides a mechanism for trading that has induced the complete opposite effect. Under the price improvement rule bids have to increase and asks have to decrease in a continuous double auction institution. As a result, over time this rule ensures that the bid/ask spreads tend to decline. Table 2 reveals that under this rule it is the most competitive market makers who remain because MM's skills are positive and increasing across the trials. That is, in this trading mechanism it is the skilled MM's who are driving out the less skilled MM's.

6.1.3 Binomial Test on Learning Coefficients

In addition to the skill parameters, Table 2 also depicts the Learning coefficients by MM, MT, Trial and Security. The actual learning curves estimated from the learning model for both MM's and MT's by security and trial is provided in Figures 5-8, appendix 1. Inverted learning curves are immediately apparent from the results provided in Table 2. That is, there are a predominate number pf negative slope parameters observed associated with each cell of table 2.. Negative learning coefficients were observed for 13/16 of the MT cells and 10/16 of the MM cells. We test for the existence of negative learning using a statistical test of the null hypothesis that the probability "p" of a positive learning coefficient equals 0.5, against the alternative hypothesis that p is not equal to 0.5. An exact binomial test of the null hypothesis was conducted and which rejected the null hypothesis at a probability less than 0.0006.

Combined the results from Tables 1 and 2 are clear for market takers. In a price taking (i.e., non strategic) competitive market the adverse selection problem created by the market taking acceptance rule resulted in the less skilled MT's driving out the relative more skilled MT's over repeated exposures of trading opportunities. For the market making side, however, we observed the opposite to adverse selection, which we refer to as advantageous selection. Surprisingly this behavior also resulted in inverse learning being observed for MM's. In the next section we explore this behavior further by allowing for strategic or non price taking behavior among the MM's.

6.2 Inverted Learning and Strategic Market Making Behavior

In a single person learning environment (without competition), we expect to see an upward learning curve which represents increasing proficiency with increased exposure.

Insert Figure 3 about here

However, in the competitive markets studied in this paper as discussed in the previous section we usually observed an inverted learning curve as depicted in Figure 4, being generated from the price discovery dynamics. For MT's the driver of this was an adverse selection problem relative to ability for assessing the intrinsic value. For MM's however, we observe advantageous selection and in this section we explore advantageous selection in the presence of strategic or non price taking behavior.

Insert Figure 4 about here

By relaxing the price taking assumption we allow for the fact that the bid/ask spread is a strategic variable for MM's. This is plausible because MM's have two ways of making money in the market. First, they can attempt to buy or sell advantageously relatively to their assessed intrinsic value. Under non strategic advantageous selection we would expect to see more success with repeated exposures. Second, MM's can attempt to earn the spread by exploiting the available liquidity from MT's. In real world markets it is not uncommon for MM's to pay for liquidity and in the markets studied in this paper "payment" is endogenous to the strategic bid/ask variable. From this second approach the learning problem facing MM's is to assess the relationship among MT's, spreads and liquidity. We will explore strategic MM using a regression approach where the dependent variable is defined as:

$$BASpread = BestAsk - BestBid$$

We analyze the drivers of this spread formally using the following variables. Each variable is measured over successive 1-second intervals of time to capture the dynamic behavior of price discovery. The first set of drivers, are variables that are predicted to be important in a setting where prices arise from the strategic interactions between the different roles within the trading crowd themselves. These are:

MTPriorKnowledge – the average prior trading knowledge θ_i of active market takers in the market. The higher the statistic, the more competent the active market takers.

MTEasiness – the average easiness β_j of the skills used by the active market takers. The higher the statistic, the easier to do market taking.

MTLearningRate -- the average learning rate γ_j of the skills used by the active market takers. The higher the statistic, the easier to do market taking over time

MMPriorKnowledge – the average prior trading knowledge θ_i of active market makers in the market. The higher the statistic, the more competent the active market makers.

MMEasiness – the average easiness β_j of the skills used by the active market makers. The higher the statistic, the easier to do market making.

MMLearningRate -- the average learning rate γ_j of the skills used by the active market making. The higher the statistic, the easier to do market making over time

MTSuccess – the average success rates of the active market takers computed across 1-second intervals of time.

MMSuccess – the average success rates of active market makers computed across 1-second intervals of time.

An additional set of potential drivers, resulting from the variables predicted under standard competitive pricing taking theory. These variables are listed next.

MTCount – the number of market taking trades

MMCount – the number of market making trades

MTVol – the average volume in thousands per market taking trade

MMVol – the average volume in thousands per market making trade

The strategy available to MM's is the ability to post the prices for which they are willing to buy and sell at subject to the price improvement rule. As a result, we conduct a multiple regression analysis using the posted prices (i.e., the bid/ask spread) as the dependent variable. This regression will allow the dynamics of bid-ask spreads and the strategic interactions between market takers and market makers to be more formally described. The results are provided in Table 3. This table displays the OLS parameter fits and their p-values. First, the MM success variable was significant and positive for three out of four trials. That is, market makers all other things being equal are more likely to apply skills correctly the wider they can maintain spreads.

Insert Table 3 about here

Prior knowledge of the active MM's was significant in three out of four trials and marginally significant in the other trial (trial number 2). The relationship is negative which implies that the more knowledgeable market makers are the more spreads tend to narrow. Thus across trials MM's skill, success and prior knowledge are influencing what bids and asks are being posted especially for trials 3 and 4. Note each of these variables reinforce the positive skill trends observed in Tables 1 & 2 for MM's driven by the advantageous selection among market makers.

Negative learning associated with spreads was significant for the MM in trials 3 and 4 and for MT's in trials 1-2 (trial 2 was at the 10% level of significance). This result combined with the trends for MM's in terms of skill, success and prior knowledge reveal the different nature of the learning problems facing the MT's and MM's. MM's are exploiting the adverse selection problem to earn a wider spread from MT's when they are learning to better assess the intrinsic value of the securities. In later trials competition among more informed/skilled MM's has increased and the nature of the MM's learning problem has changed to learning how to generate MM revenue from trading volume and spreads. . The negative learning is defined relative to the intrinsic value as spreads shrink, but observe that this coincides with trading volume (i.e., liquidity) becoming significant in trial 4. In other words, the second method of earning market making profits from liquidity and spreads is being detected in the regression analysis. The negative coefficient implies volume increase as spreads shrink. In the real world MM's pay for trading volume, and in the current markets negative learning coefficients from the skilled MM's provide the endogenous equivalent to paying for trading volume.

7. Conclusion

This paper studies the interaction between rational and behavioral finance in a market setting where the no trade theorem was strongly rejected even though equilibrium price predictions were ultimately supported. We describe how individual differences within the market's trading crowd combined with the rules of the double auction institution, resulted in relatively efficient price discovery. We adapt a learning model from machine learning, to study trading crowd behavior at both an individual and marginal trader level for market makers and takers. This model identified the importance of adverse selection for the market takers in the trading crowd and strategic advantageous selection for the market makers in the trading crowd.

For market takers, it was individual differences with respect to skills related to assessing the intrinsic value that created the adverse selection problem. Individual differences along this dimension resulted in a "market for lemons" effect whereby the skilled market takers were driven out by the relatively less skilled market takers under the price taking rules of the double auction market institution. On the MM side the price setting rule in the double auction is a price improvement rule (bids must increase, asks must decrease) combined with the use of spreads as a strategic variable that generated strategic advantageous selection on the market making side. That is, contrary to what was observed on the MT siding on the MM side it was the skilled MM's who drove out the less skilled MM's because the smaller the bid/ask spread the more competitive it is but it can imply violations of intrinsic value predictions to promote trading liquidity. Together, these dynamics led to the rejection of the strong no trade theorem and efficient price discovery.

Finally, in contrast to the usual implications that adverse selection can result in market failure, the rules of the double auction institution combined with a boundedly rational trading crowd, exploit these offsetting “selection problems” between the MT’s and MM’s to promote liquidity resulting in the rejection of the no trade theorem and relatively efficient price discovery behavior. In future work we will test for the presence of this phenomenon as a driver of price discovery in real world markets using the Nasdaq Level 2 quotes.

Trial	1		2		3		4	
Security	1	4	1	4	1	4	1	4
MM	1.31	0.27	0.40	0.48	0.37	0.36	0.55	1.49
MT	-1.48	0.03	0.33	0.01	-0.06	-0.10	-0.54	-0.14

Table 1 Compare the prior information of traders.The parameters in this table are the average of the individual trader prior knowledge parameters, θ_i resulting from fitting the learning model to the two active subsets of the trading crowd, market makers and market takers for each trial. This is the average knowledge parameter for the marginal market maker and taker.

	Trial							
	1		2		3		4	
SkillName	β	γ	β	γ	β	γ	β	γ
Bid1	2.4	0.1	2.65	-0.1	1.82	0	7.65	-0.01
Bid4	5	-0.06	3.46	-0.14	6.29	-0.1	7.81	-0.08
Ask1	2.59	0.03	3.08	-0.01	3.17	0.02	6.33	-0.03
Ask4	2.49	0	2.5	-0.02	3.04	0.01	5.27	-0.16
Buy1	-3.78	-0.01	-3.57	-0.03	-3.8	-0.01	-4.11	0.04
Buy4	-1.01	-0.02	-0.84	-0.13	-2.73	-0.04	-4.34	-0.02
Sell1	-2.97	-0.05	-1.88	0	-1.37	-0.05	-8.23	-0.02
Sell4	-2.72	-0.02	-2.33	-0.02	-3.83	-0.03	-7.81	0

Table 2 Skill easiness and learning rates over the four trials. β is the skill easiness and γ is the skill learning rate coefficients, resulting from fitting the learning model to trading crowd for each trial. The parameters in this table are the marginal trader parameters.

	1		2		3		4	
	Estimates	p	Estimates	p	Estimates	p	Estimates	p
(Intercept)	-3.50	0.02*	-0.86	0.54	-0.77	0.79	-2.94	0.12
MTPriorKnowledge	0.07	0.46	-0.01	0.47	0.03	0.64	-0.06	0.00**
MTEasiness	-0.84	0.00**	0.01	0.97	-0.88	0.07	0.08	0.45
MTLearningRate	-96.14	0.00**	-5.38	0.09	-23.39	0.45	4.66	0.50
MTSuccess	0.22	0.83	0.39	0.45	1.16	0.30	-0.11	0.88
MTCount	-0.05	0.27	-0.01	0.66	-0.02	0.67	0.03	0.19
MTVol	0.00	0.91	0.05	0.01**	0.05	0.14	0.02	0.29
MMPriorKnowledge	-0.12	0.01**	-0.03	0.09	-0.13	0.00**	-0.04	0.00**
MMEasiness	0.10	0.79	0.27	0.55	-1.09	0.02*	0.24	0.22
MMLearningRate	-6.27	0.42	5.75	0.04*	-38.74	0.01**	-15.87	0.00**
MMSuccess	4.39	0.00**	0.87	0.25	4.15	0.01**	4.78	0.00**
MMCount	-0.02	0.62	-0.01	0.71	0.00	1.00	-0.03	0.11
MMVol	-0.05	0.11	-0.01	0.64	-0.05	0.14	-0.05	0.01**

Table 3 The parameter estimates and their p-values with OLS from Model I across the four trials. The cells where the p-values are less than .05 are highlighted. The dependent variable in the regression results reported in Table 2 is the bid/ask spread for the market measured over successive 1-second intervals of time to capture the dynamic behavior of price discovery in the market. The predictor variables are the coefficients from the learning model (MT = Market Takers, MM = Market Makers) plus variables constructed from standard competitive market theory (Count = number of MM's or MT's, Vol = Traded volume for MM's and MT's, Success is the traded price defined relative to the predicted intrinsic value for MM's and MT's.

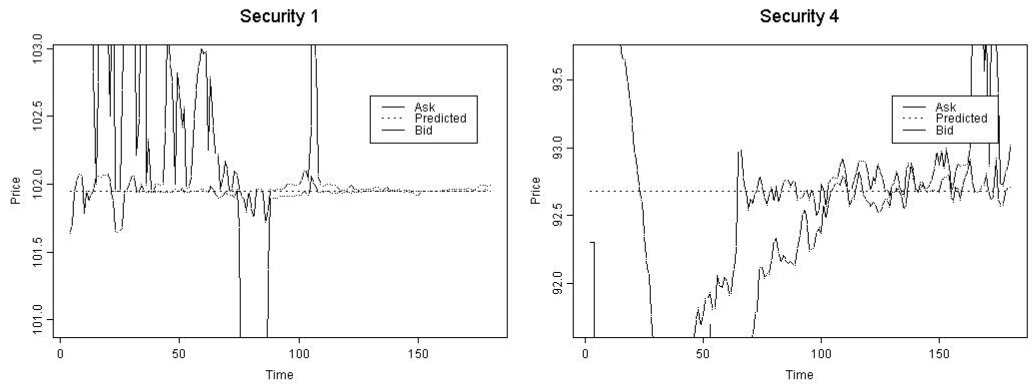

Figure 1 Prices of two securities. The y-axis in this graph represents the bid/ask deviations around the predicted intrinsic value of each security. The x-axis represents time and the path of bids and asks is plotted for period 3 of one trial. The Security 1 example, depicts a market where after initial volatility in the bids and asks the market settles down to the predicted intrinsic value. The Security 4 example depicts a market where bid/ask spread volatility sustained throughout the period and prices remained above the predicted intrinsic value.

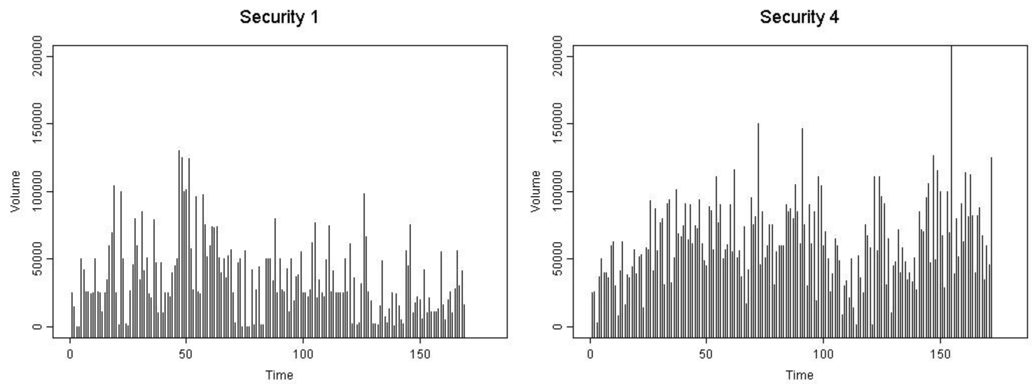

Figure 2 Trading volumes of the two securities. The y-axis represents trading volume in contracts exchanged for the security markets and trial depicted in Figure 1. The x-axis again represents time.

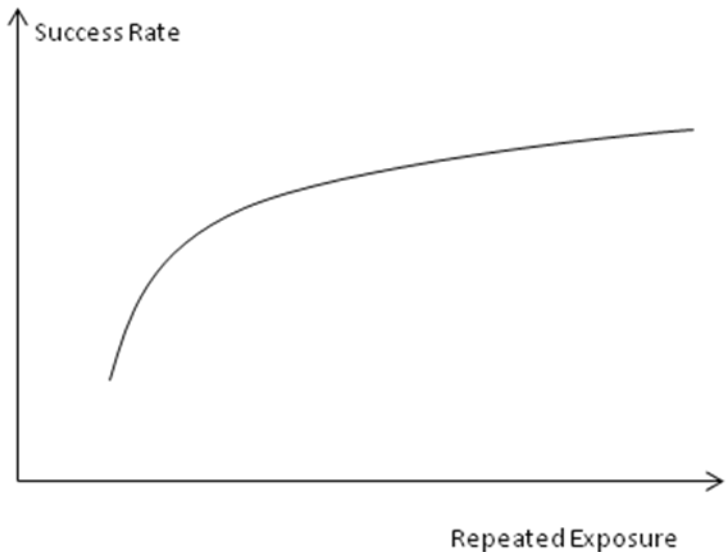

Figure 3 Classroom learning curve. The x-axis measures the number of times the user is exposed to the task and the y-axis measure the proficiency of performance.

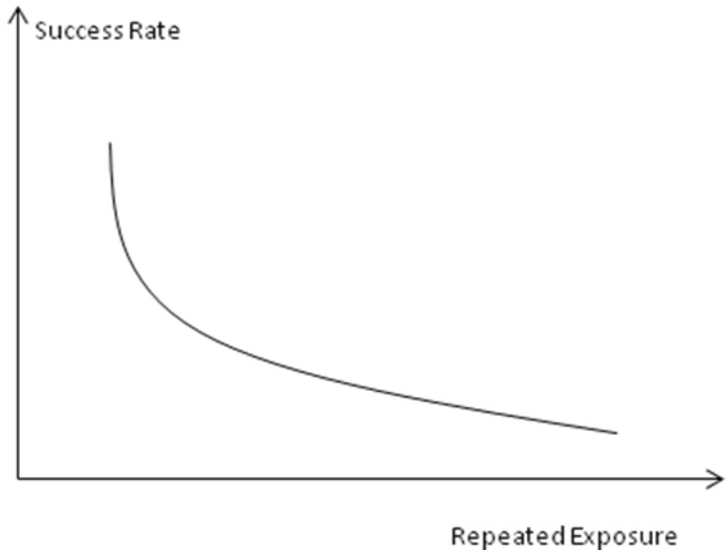

Figure 4 Inverted Learning in a Competitive Market. The x-axis measures the number of times the marginal trader is exposed to the bid/ask/buy/sell conditions and the y-axis measure the proficiency of performance.

Appendix 1

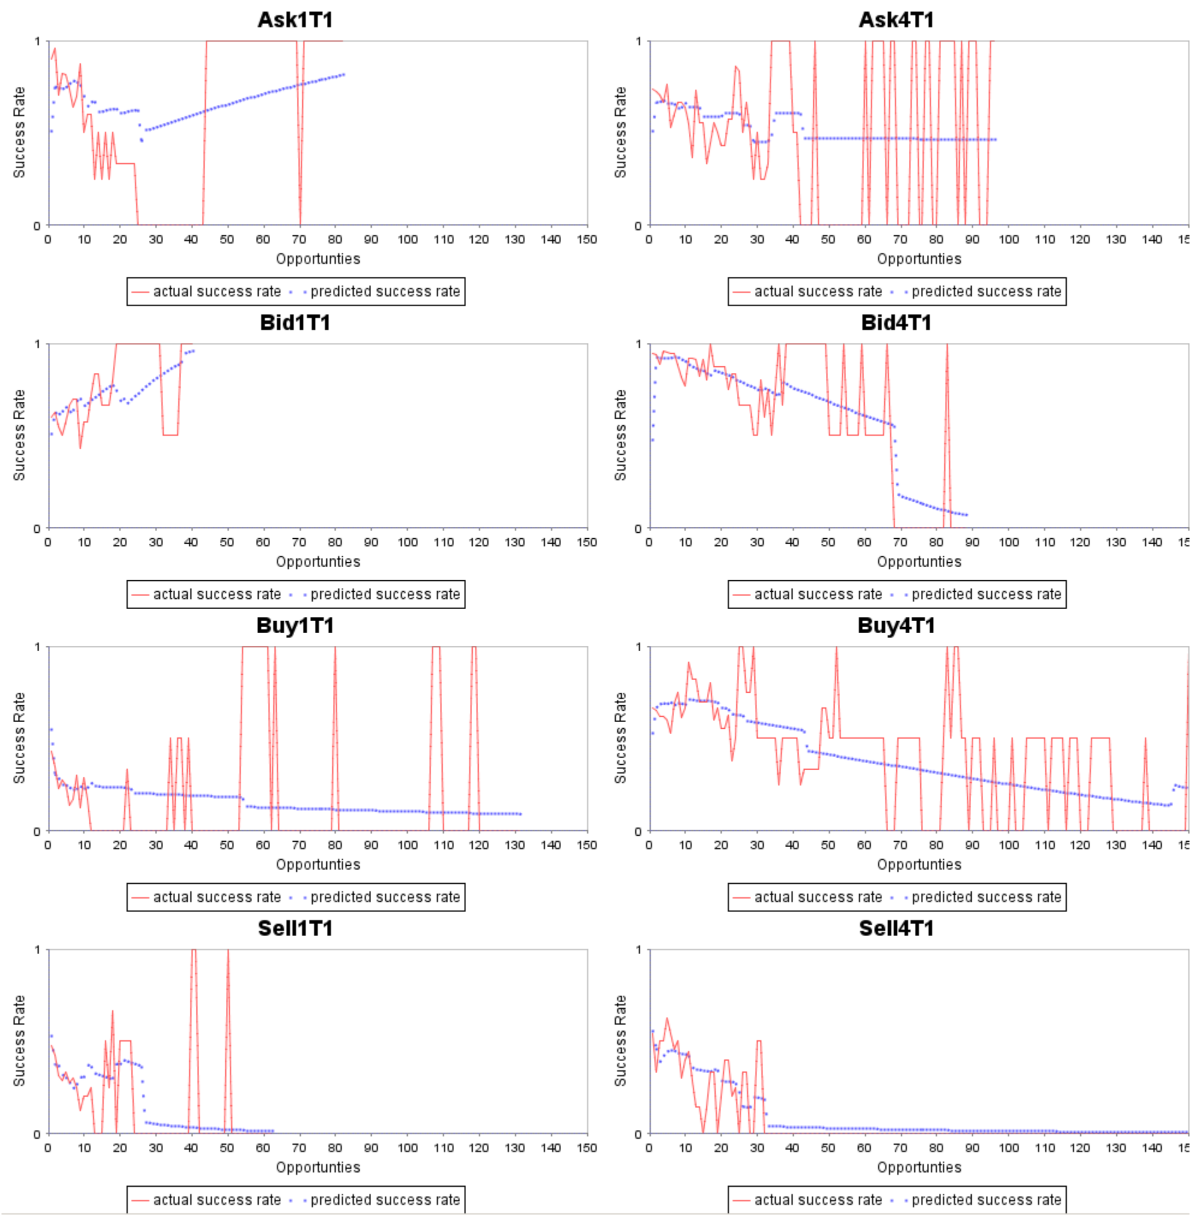

Figure 5 Learning curves in trial 1 period 3. The x-axis is the number of times the Bid/Ask/Buy/Sell opportunity was repeated during the period. The y-axis is success rate defined relative to the predicted intrinsic value of the security. The solid lines are the actual success rates of using each skill across the ordered number of practices. The dotted lines are the success rates predicted by the AFM model.

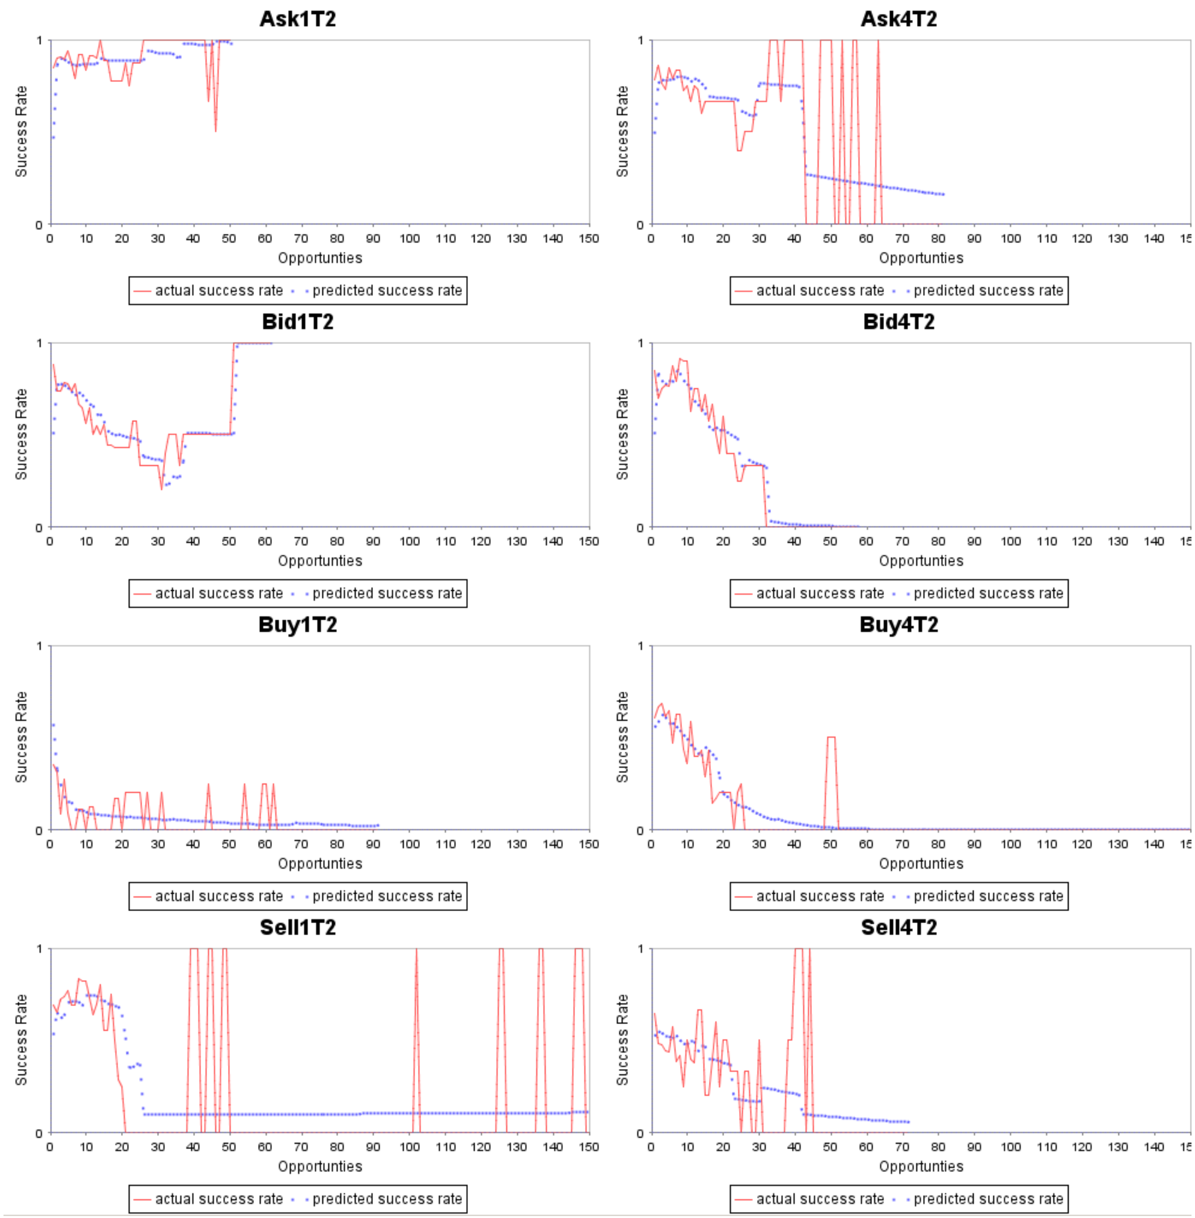

Figure 6 Learning curves in trial 2 period 3. The x-axis is the number of times the Bid/Ask/Buy/Sell opportunity was repeated during the period. The y-axis is success rate defined relative to the predicted intrinsic value of the security. The solid lines are the actual success rates of using each skill across the ordered number of practices. The dotted lines are the success rates predicted by the AFM model.

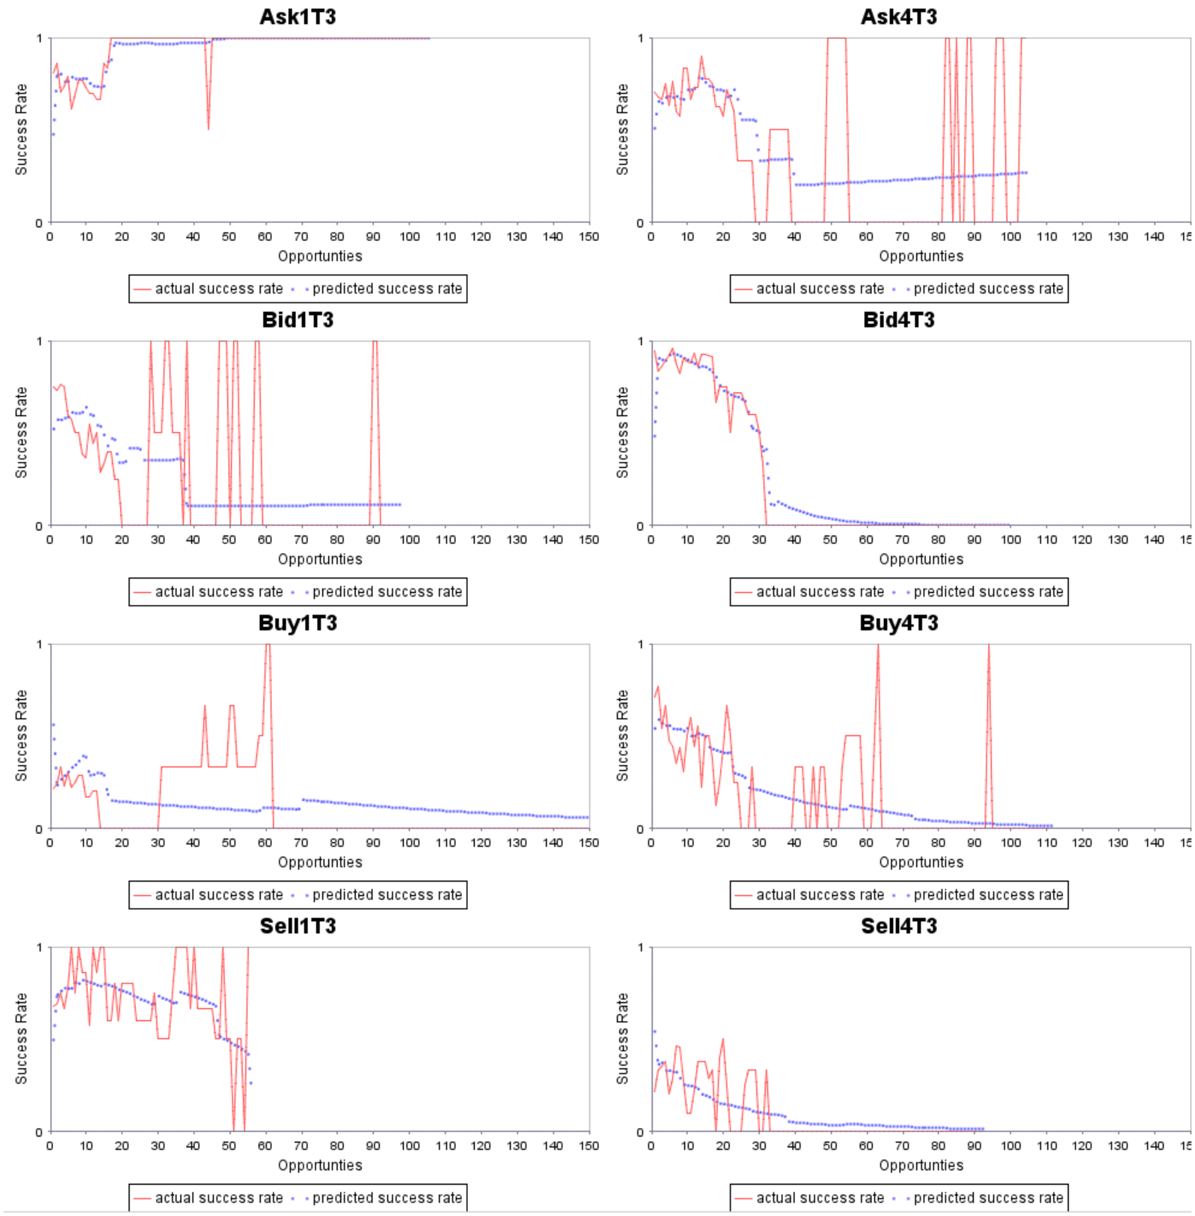

Figure 7 Learning curves in trial 3 period 3. The x-axis is the number of times the Bid/Ask/Buy/Sell opportunity was repeated during the period. The y-axis is success rate defined relative to the predicted intrinsic value of the security. The solid lines are the actual success rates of using each skill across the ordered number of practices. The dotted lines are the success rates predicted by the AFM model.

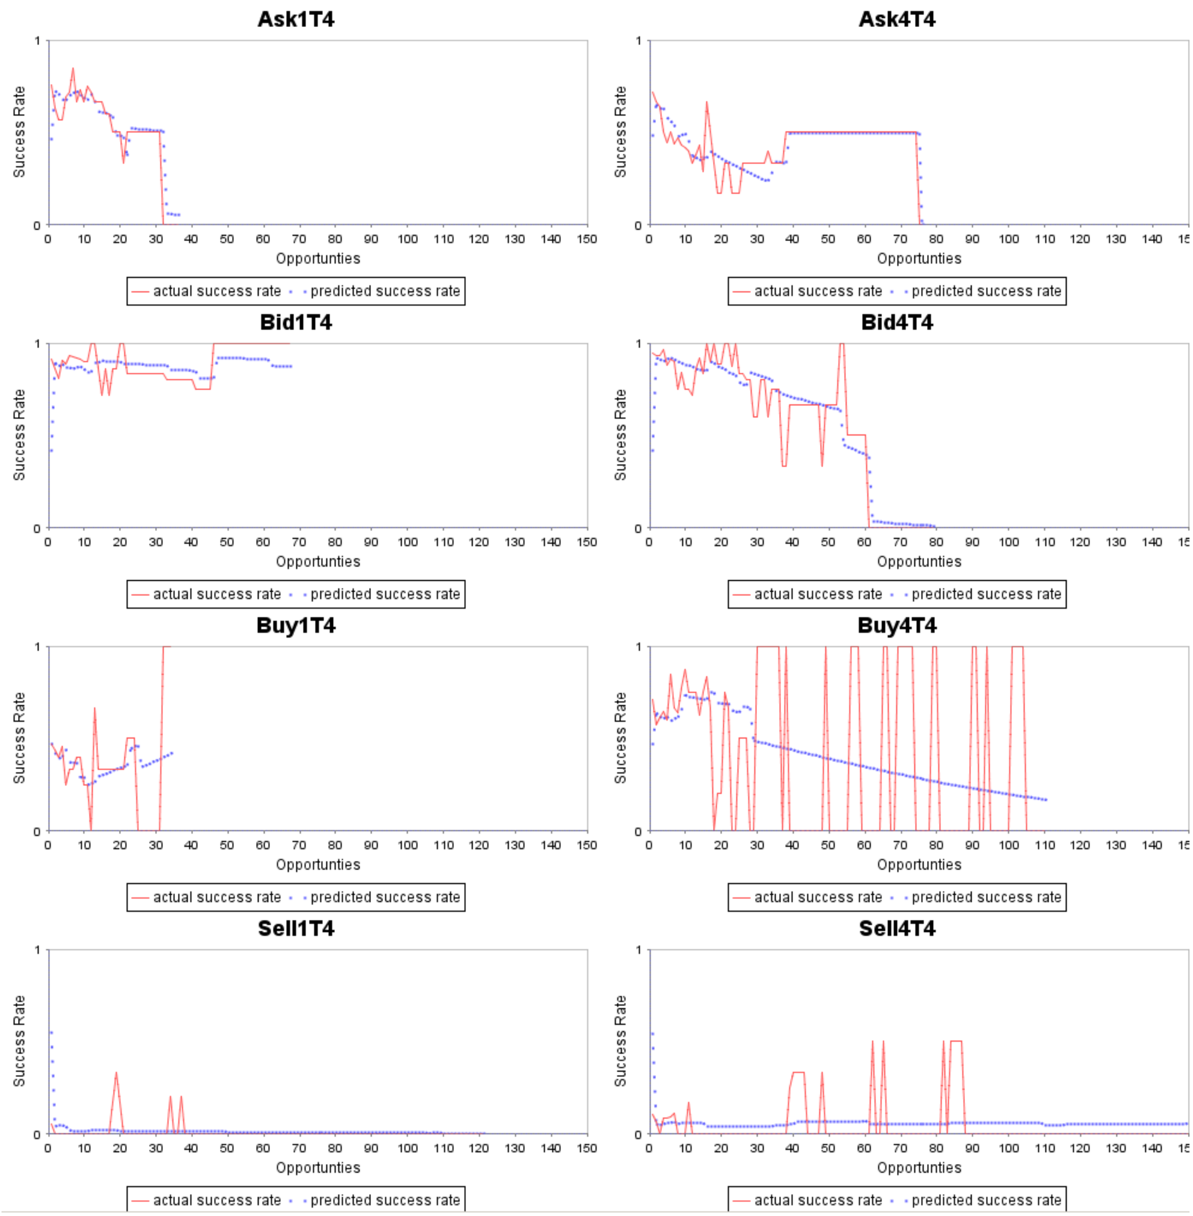

Figure 8 Learning curves in trial 4 period 3. The x-axis is the number of times the Bid/Ask/Buy/Sell opportunity was repeated during the period. The y-axis is success rate defined relative to the predicted intrinsic value of the security. The solid lines are the actual success rates of using each skill across the ordered number of practices. The dotted lines are the success rates predicted by the AFM model.

Bibliography

Akerlof, G.A., 1970. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* **84**, 488-500

Andrews, E.L., October 23, 2008. In: New York Time

Black, F., 1986. Noise. Journal of Finance 41, 529-543

Brav, A., Heaton, J.B., Rosenberg, A., 2004. The rational-behavioral debate in financial economics. Journal of Economic Methodology 11

Cen, H., Koedinger, K., Junker, B., 2006. Learning Factors Analysis - A General Method for Cognitive Model Evaluation and Improvement. In: 8th International Conference on Intelligent Tutoring Systems

Hellwig, M., 1980. On the Aggregation of Information in Competitive Markets. Journal of. Economic Theory 22, 477-498

Kyle, A.S., November 1985. Continuous Auctions and Insider Trading Econometrica 53, 1315-1335

Linden, W.V.D., Hambleton, R.K., 1997. *Handbook of Modern Item Response Theory* Springer.

Schleifer, A., April 2000. Inefficient Markets : An Introduction to Behavioral Finance. Oxford University Press.

Shefrin, H., 2007. Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing. Oxford University Press.

Shiller, R.J., May 2006. Irrational Exuberance. Bantam Books.

Tirole, J., 1982. On the Possibility of Speculation under Rational Expectations. Econometrica 50, 1163-82

Tversky, A., Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. Cognitive Psychology 5