MANAGEMENT SCIENCE
Vol. 50, No. 3, March 2004, pp. 336–351
ISSN 0025-1909 | EISSN 1526-5501 | 04 | 5003 | 0336

DOI 10.1287/mnsc.1040.0191

Suckers Are Born but Markets Are Made: Individual Rationality, Arbitrage, and Market Efficiency on an Electronic Futures Market

Kenneth Oliven, Thomas A. Rietz
Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242-1000
{kentoliven@beannet.com, thomas-rietz@uiowa.edu}

The Iowa Electronic Markets are specially designed futures markets that appear to aggregate information efficiently to predict events such as election outcomes. Yet, in theory, perfect information aggregation is impossible. Further, the markets are populated by a nonrepresentative sample of mistake-prone and biased traders. That is, traders are prone to the behavioral anomalies predicted by behavioral finance. How can this be reconciled with market efficiency? Here, we take a first step by analyzing the behavior of two self-selected types of traders. Dramatic differences in mistake rates across traders can help us answer the question. Market-making traders who set prices are less mistake prone and appear to be more rational than price-taking traders. This highlights an important feature of markets: marginal (in this case, market making), not average, traders set prices. This can drive the efficiency of market prices in spite of large numbers of traders who display patently suboptimal behavior.

Key words: prediction markets; market efficiency; arbitrage; individual rationality; Iowa electronic markets; experimental asset markets
History: Accepted by R. Jagannathan, former department editor; received July 22, 2002. This paper was with the authors 7 months for 1 revision.

1. Introduction

The recent wave of research in behavioral finance notes that human beings are prone to a host of biases and often make what appear to be outright mistakes (relative to "optimal" behavior). A reasonable inference is that, based on these biases, investors and traders in financial markets will behave in ways not predicted by traditional financial theory. Considerable evidence suggests that traders do buy, sell, and hold investments for "irrational" reasons predicted by behavioral biases and mistakes.[1] For convenience, we will call these "behavioral anomalies" as a group. Two tempting conjectures are that behavioral anomalies (1) affect market prices and (2) hurt traders (beyond transactions costs).[2] Do the traders' obvious biases and mistakes *actually* have (1) pricing implications for markets and (2) welfare implications for the traders themselves? Here, we argue that answers to these two questions are closely linked and behavioral anomalies do not necessarily imply answers for either. We provide evidence on these issues using the Iowa Electronic Markets (IEM), a unique set of large-scale, experimental futures markets.

We argue that whether behavioral anomalies have market pricing and welfare implications depends critically on whether markets are efficient or, more precisely, how efficiency arises in markets. To see why efficiency matters, suppose for a moment that markets are (strong-form) efficient. Consider a frequently cited implication of efficiency—that informed rational traders cannot "beat" the market by buying particular assets. Why? Efficient market prices should always reflect fundamental values. The flip side of this argument is that, so long as they remain well diversified, traders who buy, sell, or hold particular assets for "irrational" reasons, including documented behavioral anomalies, cannot underperform the market because market prices should always reflect fundamental values.[3] Of course, if markets are inefficient, then some traders may well gain, while others lose, when they buy, sell, and hold securities for either "rational" or "irrational" reasons.

Can we expect markets to be efficient? This is a matter of a long, well documented debate.[4] Frequently,

[1] See Nofsinger (2002) for a nice summary of much of the recent behavioral research in investments.

[2] For an example of predicted market effects, see Daniel et al. (2001).

[3] Anomalies that lead to poorly diversified portfolios of efficiently priced assets can lead to reductions in welfare because of increased risk.

[4] The idea that markets efficiently organize trade and production goes back at least to Smith's (1776) "invisible hand." Nearly any

the simple behavioral finance argument is that if large groups of traders are prone to pervasive biases, then these biases will affect prices.[5] There are numerous simple counterarguments. For example, a few arbitrageurs may drive out the effects of biased traders. Biased traders may be driven from the market by losses incurred as a result of their biased trading. Because even biased traders may learn from market prices, they may update their expectations in ways that defeat the bias.

We argue that outcomes in markets populated by biased and mistake-prone traders come from an interaction between traders and the market institution that, frequently, is not fully considered. A fundamental property of markets is that marginal, not average, traders determine prices. Who marginal traders are and how they set prices determine whether markets are efficient. Whether markets are efficient determines whether prices reflect fundamental values, and as a result whether buy, sell, or hold decisions based on behavioral anomalies hurt traders.

Using data from the IEM, we study how efficiency can arise in a market populated by biased, mistake-prone traders. We do this by studying the self-selected roles that participants choose for themselves, paying particular attention to the price-determining, marginal traders. Prior evidence (Forsythe et al. 1992, 1999) documents trader biases and that highly biased traders generally tend to buy and hold securities. Buy-and-hold strategies have, at most, transitory effects on prices. Here, we document evidence that many traders are prone to outright mistakes (buying and selling at prices that are not the best available or violating arbitrage restrictions). However, traders are most prone to mistakes when they place market orders and trade at prices set by other traders.

In contrast to traders who are biased and mistake prone, we document that, when setting at-market limit orders and setting prices, traders are much less mistake prone. That is, traders who "make" the market also appear more rational.[6] Other evidence suggests that active at-the-market traders are much less biased than buy-and-hold traders. As a result, the self-selected, marginal, price-setting, market-making traders appear to be a much more rational subset of traders than average. The prices they set also happen to be efficient. This should not be difficult to believe. After all, traditional financial theory argues that "rational" traders drive efficient markets. The recognition here is that not all trader actions have to be rational. At least in the IEM, rational, marginal trades can generate efficient prices in spite of biases, mistakes, and even theoretical arguments that suggest markets like the IEM cannot possibly be efficient.[7]

2. The Iowa Electronic Markets

The IEM are futures markets with contracts specifically designed so that prices can be used to make predictions. They are described in detail on the IEM website (www.biz.uiowa.edu/iem) and in Berg et al. (2004). Here, we summarize the major features of the IEM, focusing on the IEM political markets (designed to predict political outcomes) and the specific markets analyzed in this paper.

The IEM began as the Iowa Political Stock Market in 1988.[8] Here, we will focus on the markets run to predict vote shares in the 1992 presidential election. We use the 1992 vote share market as a prototypical example for several reasons. First, in the final analysis, the 1992 market was the most accurate to date in predicting the outcome. Second, it had a large number of active traders and the largest volume of any market run to date. Third, and most important, we have detailed survey information from an overwhelming majority of the traders, which we will use in our analysis.[9] While we will be brief here, the appendix, available on the *Management Science* website (mansci.pubs.informs.org/ecompanion.html), contains descriptions of these markets from the trader's manual.

2.1. Common Market Characteristics

Markets on the IEM share many common characteristics. They are real-money, real-time futures markets conducted through the Internet. Market traders invest their own cash, incurring the risks and returns of their trading behavior. There are no explicit transaction costs.

Each market contains a complete set, or "unit portfolio," of contracts. Each unit portfolio is a risk-free bundle of contracts that pays off exactly $1

general finance or investments textbook contains a discussion on the debate over market efficiency. Fama (1970, 1991) reviews the literature on efficiency.

[5] An entire vein of research cited in Daniel et al. (2002) argues that markets will be inefficient because of biases that few, if any, traders escape.

[6] Actually, as we shall show later, one can interpret the results in either of two ways. Our main direct evidence will show that traders appear more rational on an order-by-order basis when acting as market makers. We will also show evidence that traders who are market makers (that act only as market makers or more often as market makers) also appear more rational on a trader-by-trader basis.

[7] The IEM markets mirror closely those modeled by Grossman and Stiglitz (1980).

[8] Forsythe et al. (1992) describe the initial market.

[9] Practically speaking, surveys were mandatory at the time of the 1992 market and have since become optional, resulting in a much lower response rate. We focus on 1992 to avoid any sample biases caused by the lower response rates.

(or, in some cases, another known fixed amount). For example, in the 1992 presidential race, one unit portfolio consisted of a contract that paid off $1 times the percentage of the two-party vote received by George Bush and one contract that paid off $1 times the percentage of the two-party vote received by Bill Clinton. As defined, the percentages must sum to 100%, so the payoff to the set of contracts must equal $1. Unit portfolios consisting of one of each contract could be purchased from, or sold to, the exchange for $1 (its certain value) at any time. This structure provides a means for traders to create and eliminate contracts. It also creates a risk-free portfolio from which all contracts can be priced according to arbitrage pricing theory. Finally, it creates complete contingent claims with fixed aggregate payoffs. This implies that prices should equal expected values.[10]

Besides buying and selling whole unit portfolios, traders could buy or sell individual contracts. The market for each individual contract ran as a continuous, electronic double auction with bid and ask queues in price increments of $0.001. Traders self-selected into (and out of) the markets and choose their own roles. Traders could act as market makers by submitting bids above the outstanding best bid or asks below the outstanding best ask. Traders could act as price takers by accepting an outstanding bid (making a sale) or accepting an outstanding ask (making a purchase). In addition, traders could place limit orders by submitting bids and asks outside the outstanding spread.

2.2. The 1992 Presidential Markets

For the 1992 presidential election, the creators of the market were allowed to operate the market internationally for up to 1,000 individuals with investments not exceeding $500 apiece.[11] They filled their allotment of traders, with traders investing more than $83,000 of their own funds. All bids, offers, trades, etc., were stored with time stamps (in one-second increments) for later analysis.

The specific market we choose to study is the 1992 Presidential Vote Share Market from July 17, 1992, through the election.[12] Contracts in this market traded in two submarkets each with one unit portfolio. The online appendix (mansci.pubs.informs.org/ecompanion.html) contains complete descriptions of the markets. In the two-party submarket, the contracts were:

Contract	Liquidation value
R.BU	$1 times % of two-party vote received by the Republican nominee
D.CL	$1 times % of two-party vote received by the Democratic nominee

Because prices should equal expected values in these markets (see Footnote 10), prices should be the market forecast of the expected two-party vote shares received by each candidate. In the Perot submarket, the contracts were:

Contract	Liquidation value
PERO	$1 times % of three-party vote received by Ross Perot
D&R	$1 times % of three-party vote received by the Democratic and Republican nominees together

Here, prices should be the market forecast of expected vote shares received by Perot versus the Democratic and Republican nominees combined.

2.3. The Case for Efficiency

The 1992 presidential market was the most active and liquid of the vote share markets conducted to date. It is also one of the most efficient ever conducted in terms of deviations of final (election eve) prices from true values and in terms of prices being better predictors of true values than other publicly available information such as polls. Final major poll predictions of the vote fractions received by the three candidates differed from the actual fractions received by average absolute errors of 1.2% (Harris) to 3.8% (Gallup). In contrast, the election-eve forecast generated by market prices differed from the actual fractions received by an average absolute error of 0.2%. Thus, in the final analysis, this market predicted the election outcome very efficiently.

The case for efficiency in these markets is not based solely on this single accurate point prediction. Berg

[10] See Berg et al. (2003) for a brief description on how prices in such markets should reflect expected values regardless of risk preferences. See Borch (1960), Caspi (1974), and Malinvaud (1974) for more detailed discussions. Intuitively, returns should reflect the risk-free rate, covariances with the market portfolio and the market risk premium (the difference between the market portfolio return and the risk-free rate). However, in these markets, the risk-free rate is zero and the market portfolio has a fixed payoff with a zero return, so all covariances and risk premiums are zero. Hence, all securities should have a zero expected return. As a result, all contracts should be priced at their expected future values.

[11] The market operated with a no-action letter from the staff of the Commodity Futures Trading Commission.

[12] This was the first day in which two securities were traded in each submarket as discussed later. Before this date, securities that had values based on other potential Democratic nominees' vote shares were traded.

et al. (2004) show that election-eve prices are unbiased and accurate forecasts of election outcomes and that they are more accurate than election-eve polls across state, national, and international elections. Berg et al. (2003) show that prices predict well significantly in advance of elections and are significantly better than advance polls. Further, they show that prices appear to follow an efficient random walk across time as predicted by efficient markets theory.

2.4. The Puzzle of Efficiency

The apparent efficiency of election markets seems almost too good to be true. Grossman and Stiglitz (1980) argue that such informationally efficient markets are impossible. Without some profits available to motivate them, informed traders would have no incentive to trade. Further, evidence in Forsythe et al. (1992), Berg et al. (2003), and here shows that the traders in election markets are far from a random sample of the population and are unlikely to have opinions that are representative of the population. Forsythe et al. (1992, 1999) show that opinions of traders affect their actions biasing their beliefs about the election outcomes, their interpretations of news and their holdings in the markets. They document two behavioral anomalies (a false consensus effect and a cognitive dissonance effect) that lead to a clear bias (a wishful thinking bias) that colors traders' perceptions and actions in these markets. Thus, traders in these markets are prone to behavioral anomalies. In the next section, we document that traders in these markets make outright errors that appear inconsistent with that rational maximizing behavior driving models of efficient markets.

These observations result in the question underlying this paper: How can electronic markets be efficient in predicting election outcomes when theory predicts inefficiency and, in practice, the markets are populated by nonrepresentative traders who are prone to mistakes and behavioral anomalies?

3. Market Structure and Data

To understand the dynamics driving the apparent efficiency in these markets and reconcile it with the apparent biases and mistakes that characterize traders, we study the behavior of individuals who determine prices within the microstructure of these particular markets. In particular, we study trader behavior with respect to two assumptions that typically underlie models of rational traders and efficient markets: the law of one price and arbitrage-free pricing. We study how well traders conform to these assumptions and study how their interactions with the market create the efficiency we observe in the IEM.

3.1. Market Structure

IEM markets are ideally suited for studying individual rationality, trader dynamics, and efficiency issues for several reasons. First, the true values of the traded contracts are revealed at a single point in time (the election). Second, public information about contract values (e.g., polls) exists and is readily identified. Third, arbitrage constraints are easily identified and individual traders can exploit them with no explicit costs. Fourth, individual trades are stamped with trader identifications and, for the market we study, we have detailed demographic and survey response information from the traders. Finally, since the IEM runs continuously, there should be no irregularities associated with openings, closings, weekends, etc. These properties allow an in-depth analysis of the full range of issues surrounding trader rationality and market efficiency.

The way the market was structured, two separate acts determined each trade price. One trader set a limit order and another submitted a market order (accepting the limit order). The market order generated a trade at the price set in the limit order. We refer to traders who submitted at-market limit orders as *market makers*. We refer to traders who submitted market orders as *price takers*.

Because of the way unit portfolios were structured, there were always two different ways for price-taking traders to take desired positions. A price-taking trader could establish a position (1) directly by trading a single contract, or (2) indirectly by combining a unit portfolio transaction with a simple transaction in the complementary contract in the unit portfolio. Since there were no trading costs, these were equivalent. To see how this works, suppose that a price-taking trader wanted to take a position with one fewer of a particular contract. That trader could sell the contract immediately at the high outstanding bid. For example, a share of R.BU could be sold at the R.BU bid. Alternatively, that trader could purchase the complementary share(s) at the market ask(s) and sell the unit portfolio to the exchange for \$1. The alternative way to "sell" an R.BU was to purchase a D.CL at the market ask and then sell a unit portfolio to the exchange for \$1. The net trade is identical: one less R.BU. Similarly, a price-taking trader could acquire a contract by either purchasing it directly at the ask or by purchasing a unit portfolio and selling the complementary contract(s). The law of one price is based on the idea that rational traders will always seek out the lowest price when buying and the highest price when selling. We should not observe trades at anything other than the best available price. We call a violation of this premise a *price-taking violation* of individual rationality.

It was possible for market makers to violate individual rationality as well. To see how, consider how the unit portfolio/market structure resulted in two different means for market-making traders to set prices for any given position. The trader could have committed to a price directly by submitting a limit order for a single contract. Alternatively, the trader could have committed to it indirectly by submitting an order for the complementary contract and making a unit portfolio transaction. To see how this works, suppose that a market maker wanted to quote a price for a position with one fewer of a particular contract. That trader could submit an ask for that contract. For example, a trader could offer to sell a share of R.BU by submitting an ask for R.BU. Alternatively, that trader could bid to purchase the complementary share(s) and, if there was a purchase, sell a unit portfolio to the exchange for $1. Thus, the alternative way to offer to "sell" an R.BU was to bid to purchase a D.CL. If that bid were accepted, the trader would sell a unit portfolio to the exchange for $1. Conditional on the order being accepted, the net trade was identical: one less R.BU. Thus, both were means of offering the same position. Similarly, a market-making trader could offer to buy a contract either by submitting a bid for it or by purchasing a unit portfolio and submitting an ask (or asks) for the complementary contract(s). Finally, instead of committing to a purchase or sale by submitting a limit order, market-making traders could have simply traded to their desired position by accepting available bids or asks. We should not observe market makers who commit to buy or sell at prices worse than those immediately available at current outstanding bids and asks. We call a violation of this premise a *market-making violation* of individual rationality.

The unit portfolio/market structure also creates arbitrage free pricing relationships. There are two different ways to purchase a portfolio: (1) directly from the exchange for $1, or (2) at each of the outstanding asks at the sum of the ask prices. There are also two different ways to sell a unit portfolio: (1) directly to the exchange for $1, or (2) at each of the bids at the sum of the bid prices. Arbitrage-free pricing restricts the bids and asks. The sum of the bids must be less than $1 or else a trader could profit by buying the unit portfolio from the exchange for $1 and simultaneously selling it for the sum of the bids. The sum of the asks must be greater than $1 or else a trader could profit by buying the unit portfolio at the sum of the asks and simultaneously selling it to the exchange for $1. As we will see below, violations of arbitrage-free pricing result from the market-making violations of individual behavior described in the last paragraph.

Thus, overall, the unit portfolio/market structure allows us to:

(1) identify traders according to role (market making versus price taking);

(2) determine when traders are making mistakes by not trading at the best possible price (because there are always two ways of trading to any position); and

(3) determine when available prices violate a basic underlying assumption of efficient market models (arbitrage-free pricing).

The survey information allows us to correlate characteristics of traders with behavior.

3.2. Data

Our data set consists of actions we can unambiguously classify as price-taking or market-making behaviors in the presidential market from July 17, 1992, until the election. Price-taking behaviors are immediate purchases or sales. Thus, we classify accepting an outstanding best bid or ask as price taking. We also classify making a bid that exceeds the outstanding ask as price-taking behavior because it results in an immediate purchase at the ask. Similarly, we classify making an ask that falls below the current outstanding bid as price-taking behavior because it results in an immediate sale at the bid.[13] Market-making behaviors are bids to buy or asks to sell that set new best limit prices. Thus, we classify a newly submitted bid that strictly exceeds the outstanding best bid as market-making behavior. Similarly, we classify a new ask that is strictly less than the outstanding best ask as market-making behavior. We do not evaluate off-market limit orders for two reasons. First, having this order rise to the top of the queue is not necessarily an active choice of the trader. Second, the only way such an order could become a violation of arbitrage conditions is if another trader submits a new best bid or ask order in the complementary contract that also violates the restriction. Thus, we capture such violations as market-making violations when this other trader submits his or her order. To do otherwise would double count violations. Selecting data using these criteria leaves us with a data set of 5,713 observations.

For each observation in the data set (from now on referred to as each "order"), we have the type of order (bid, ask, purchase, or sale[14] which allows us to determine whether it was a market-making or price-taking

[13] If the quantity demanded in such orders exceeds the quantity available at the outstanding bid or ask, the transaction runs through the opposing queue and we account for each transaction at its actual price. This minimizes any violations that occur. Thus, we classify such events in as conservative a manner as possible.

[14] As mentioned above, bids that cross the ask queue are classified as purchases at the ask and actual traded quantity. Similarly, asks that cross the bid queue are classified as sales at the bid and actual traded quantity.

order and whether the order was on the buy or sell side of the market), the date and time of the order, the trader's identification number, the trader specific order number (1 for the trader's first order submitted to the market, 2 for the second, etc.), the quantity ordered or actually traded, the price of the order, the total dollar size of the order (price × quantity), the total dollar volume on the day of the order, the relative spread in the ordered contract to the spread in the complementary contract and the outstanding best bids and asks for all securities in the market at the time of the order. Using the trader identification number, we link each order to available demographic information about the trader provided by the trader in response to online polls conducted by the IEM. Most traders provided information about their sex, race, religious affiliation, age, income class, education level, academic status, college major if appropriate, and levels of other financial market experience, participation, and knowledge. Table 1 gives brief descriptions of each of these variables and all transformed variables that we use in the later data analysis.

4. Results

The behavioral anomalies we document are violations of individual rationality embodied in the law of one price and arbitrage-free pricing. We categorize results into five areas. First, we document the existence of violations. Second, we show how violation

Table 1 Variables in the Data Set

Variable name	Description	Obs.	Mean	Std. dev.
Order specific variables				
ID	Trader's ID #	5,713	NA	NA
Order number	Trader's order #	5,713	950.07	1,839.18
ln(Order number)	Natural log of trader's order #	5,713	4.86	2.35
Role	1 if order is new best bid or ask	5,713	0.64	0.48
Side	1 if order is to buy (bid or purchase)	5,713	0.63	0.48
Quantity	Ordered or traded quantity	5,713	19.53	53.14
Price	Order bid, ask, or trade price	5,713	0.47	0.25
Order size	Price × quantity	5,713	8.54	22.96
ln(Order size)	Natural log of order size	5,713	0.93	1.54
Daily dollar volume	Daily total dollar volume	5,713	142.08	115.02
Bid	Outstanding best bid at order	5,713	0.46	0.25
Ask	Outstanding best ask at order	5,713	0.49	0.25
Cross bid	Bid in complementary contract	5,713	0.51	0.25
Cross ask	Ask in complementary contract	5,713	0.53	0.26
Relative spread[I]	Spread in ordered market over spread in complementary market	5,713	1.57	2.84
ln(Relative spread)[I]	Natural log of relative spread	5,713	−0.22	1.20
Trader demographic variables (self-reported)				
Sex	Sex 1 = F, 0 = M	5,711	0.10	0.30
Race	Race[II]	5,439	NA	NA
Religion	Religion[III]	5,353	NA	NA
Religious affiliation dummy	1 if any religious affiliation reported	5,713	0.68	0.47
Age	Age	5,061	30.09	9.54
Family income class	Family income class[IV]	5,139	1.54	0.70
Education level	Highest degree attained[V]	4,817	2.50	1.02
Academic status	Academic status (college)[VI]	5,264	NA	NA
College major	Major (if student)[VII]	4,680	NA	NA
Financial market knowledge	Knowledge of financial markets[VIII]	5,021	2.13	0.79
Financial market experience	Financial market experience[IX]	5,039	2.00	1.02
Financial market participation	Other financial market participation level (hours per week)	5,048	6.98	14.86

[I] To avoid what would appear to be a locked market when bid and ask queues cross, the spread is defined as the minimum of the ask minus the bid and 0.001 (the smallest pricing increment in the market). The relative spread is the spread in the contract divided by the spread in the complementary contract in the unit portfolio.

[II] 1 = White, non-Hispanic origin; 2 = Black, non-Hispanic origin; 3 = Hispanic; 4 = Asian or Pacific Islander; 5 = American Indian or Alaskan native; 6 = Other.

[III] 1 = Protestant; 2 = Catholic; 3 = Jewish; 4 = Other; 5 = No religious affiliation.

[IV] Best estimate of family's income class: 1 = Lowest 1/3 of US families; 2 = Middle 1/3; 3 = Top 1/3.

[V] 1 = High school; 2 = Bachelor's; 3 = Master's; 4 = Doctorate.

[VI] 1 = Freshman; 2 = Sophomore; 3 = Junior; 4 = Senior; 5 = MA/MBA candidate; 6 = Law or Medical student; 7 = Ph.D. candidate; 8 = Faculty; 9 = Other.

[VII] 1 = Business; 2 = Social Science; 3 = Humanities; 4 = Natural Science; 5 = Mathematics or Engineering; 6 = Other; 7 = Not a student.

[VIII] 1 = Beginner level; 2 = Intermediate level; 3 = Advanced level.

[IX] 1 = Novice; 2 = Limited; 3 = Experienced Amateur; 4 = Professional.

frequencies vary across types of activities. We show that market makers are much less likely to produce violations than price takers. We also show traders on the buy side of the market are less likely to produce violations than sell-side traders. Third, we investigate the effects of some self-reported trader demographics and market-specific experience on violation frequencies. We show that decreases in violation frequencies are associated with increases in measured market-specific experience levels and self-reported financial market knowledge levels. In contrast, increases in self-reported family income class are associated with increased violation frequencies. We find that no other collected demographic information helps significantly in explaining violations in the multivariate analysis. Fourth, we discuss the impact of market liquidity on violation frequencies. We find that more competitive markets (those with lower relative bid-ask spreads) produce fewer violations. We also find that more hectic markets (those with higher total daily dollar volumes) produce more violations. Fifth, we look to the collected demographic information to help identify who chooses to make markets versus those who choose to take prices. While it remains largely unexplained, we find this choice is significantly affected by market-specific experience and general financial knowledge, education, sex, and religious affiliation. Finally, we show that, while our main analysis is on an order-by-order basis, the results hold up on a trader-by-trader basis. This allows a broader interpretation of the results. These results drive our observations on market dynamics and efficiency in the conclusion.

4.1. Violations of Individual Rationality and No-Arbitrage Restrictions

As discussed above, we classify actions as violations of individual rationality or the no-arbitrage restrictions according to two very conservative definitions that we call market-making violations and price-taking violations. The following notation will help frame the discussion. Let the contracts constituting a unit portfolio be indexed by i, for $i = 1$ to n.[15] The highest (or "best") bid price, at time t, associated with contract i is denoted B_t^i. The lowest (or "best") ask price, at time t, of that contract is denoted as A_t^i. If contract i trades at time t, the price is denoted by P_t^i. Given this notation, it is simple to describe violations of no-arbitrage: $\sum_{i=1}^n B_t^i > \$1$ or $\sum_{i=1}^n A_t^i < \$1$. Similarly, it is simple to identify a trade that violates individual rationality as embodied in the law of one price: $P_t^i > \min[A_t^i, 1 - \sum_{j\neq i} B_t^j]$ for a purchase violation, and $P_t^i < \max[B_t^i, 1 - \sum_{j\neq i} A_t^j]$ for a sale violation. Both of these are violations of assumptions underlying the efficient market hypothesis. Since the current bids and asks are in each trader's information set, they can profit immediately from violations of the no-arbitrage condition. Similarly, traders could strategically place bids to profit on average from patterns of violations of individual rationality where traders systematically do not trade at the best available prices.

4.1.1. Market-Making Violations. A market-making violation occurs when a trader acts as a market maker (posting the best bid or ask) and this action forces a violation of the no-arbitrage restriction. If a market maker submits a new best bid, we classify it as a violation if the bid submitted forces $\sum_{i=1}^n B_t^i >$ \$1. This bid creates a violation at the market level because it creates a no-arbitrage violation.[16] This bid is also a violation at the individual level. The trader could have effectively purchased the contract *immediately* at a *lower* price by purchasing a unit portfolio of contracts from the exchange for \$1 and selling the other contract at its best bid.[17] If a market maker submits a new best ask, we classify it as a violation if the ask submitted forces $\sum_{i=1}^n A_t^i < \$1$. This ask creates a violation at the market level because it creates a no-arbitrage violation.[18] This ask is also a violation at the individual level. The trader could have effectively sold the contract *immediately* at a *higher* price by purchasing the other contract at its best ask and selling a unit portfolio of contracts to the exchange

[15] Here n refers to the number of different security types in a market, not the total outstanding shares in a market. In our case, n will equal 2 in each submarket.

[16] Any trader noticing this situation could purchase a unit portfolio of contracts from the exchange for \$1, sell each contract at its best bid and make an immediate risk-free profit.

[17] To see this, consider the two contract case and identify the new bid to be in contract 1: $\sum_{i=1}^n B_t^i > \$1$ implies $B_t^1 > \$1 - B_t^2$. The latter equation says that the trader is offering to purchase a contract at a price higher than the price at which he or she could have attained the same position through a unit portfolio purchase and immediate sale of the other contract at its bid. The latter transactions give a lower price *and* eliminate execution risk. For example, suppose a trader puts in a bid of \$0.60 for R.BU when a bid for D.CL of \$0.50 is outstanding. The trader is offering to pay \$0.60 for R.BU when the trader could have purchased the unit portfolio (for \$1) and sold the D.CL at the outstanding bid of \$0.50 for an immediate net price of \$0.50 for one additional R.BU. Since most traders held unit portfolios, the portfolio purchase half of this transaction is generally not even necessary. Further, order would not matter in this case, so they could have traded the D.CL contract first, eliminating any execution risk. Similar arguments hold for larger numbers of contracts.

[18] Any trader noticing this situation could purchase each contract in a unit portfolio at its best ask, sell the unit portfolio to the exchange for \$1 and make an immediate risk-free profit.

Table 2 Violations of No-Arbitrage Restrictions and Individual Trader Rationality

Action type	Data	Obs.	Violation frequency[I] (%)	Average size ($)	Max ($)	Total ($)
All	All bids,[II] asks,[II] and trades[III]	5,713	16.93	0.097	4.26	89.24
Market making	All bids and asks[II]	3,674	5.39	0.071	1.345	14.06
	Bids[II]	1,900	3.26	0.061	0.720	3.77
	Asks[II]	1,774	7.67	0.076	1.345	10.29
Price taking	All trades[III]	2,039	37.7	0.104	4.260	75.18
	Purchases[III]	1,074	34.7	0.115	4.260	40.21
	Sales[III]	965	41.0	0.095	1.580	34.97

[I]Percentage of actions that result in a violation of individual rationality (the percentage of total bids, asks, and/or sales that force arbitrage violations (for market-making bids or asks) or are not at the best available price (for price-taking purchase or sales)).
[II]Excluding bids and asks that crossed the opposite queue and traded immediately.
[III]Including bids and asks that crossed the opposite queue and traded immediately.

for $1.[19] These violations are irrational in the sense that traders who submit these orders give up certain immediate profits. These violations create market inefficiencies because traders who observe only current market information can make immediate profits.

Table 2 shows the frequency and size of market-making violations. Out of 3,674 new best bids or asks submitted to the market, 5.39% resulted in no-arbitrage violations.[20] The average total size of the violation (dollar size of the violation × quantity available) was 7.1 cents. This is within the range of payoffs conventionally regarded as salient for inducing trade (see Smith 1976). Since the average bid or ask was for a total value of $8.82 and there were no transactions costs, this represented an average violation of 0.8% of the order's total value. While 7.1 cents does not seem large, if errors of the same order of magnitude occurred in the Treasury Bill or Eurodollar futures markets, they would result in arbitrage opportunities worth thousands of dollars each.

The frequency and size of market-making violations surprised us given the size, duration, and apparent efficiency of the market, however, they paled in comparison to price-taking violations. We turn to these violations now.

4.1.2. Price-Taking Violations. A price-taking violation occurs when a trader accepts a price for a trade that is not the best available price. This violates individual rationality assumptions underlying the law of one price. In this market, there are always two ways to make a trade. We classify a trade as a violation only if the other means of making the trade was both at a better price and immediately available.

Assume that a trader has decided to buy contract i. The trader could purchase the share from the contract's ask queue (by accepting the best ask or crossing the queue with a bid higher than the best ask). The trader could also effectively purchase the contract by purchasing a whole unit portfolio and selling the complementary contracts in the market at their respective bids. Since there are no transactions costs, a trader should buy contract i at A_t^i if $A_t^i < 1 - \sum_{j\neq i} B_t^j$. Similarly, a trader should buy a whole unit portfolio and sell the complementary contracts at their bids if $A_t^i > 1 - \sum_{j\neq i} B_t^j$.[21] A similar argument holds for a trader who wants to sell contract i immediately.

[19] To see this, consider the two contract case and identify the new ask to be in contract 1: $\sum_{i=1}^{n} A_t^i < \$1$ implies $A_t^1 < \$1 - A_t^2$. The latter equation says that the trader is offering to sell a contract at a price lower than the price at which he or she could have attained the same position through a combination transaction. For example, suppose a trader puts in an ask of $0.40 for an R.BU when an ask for D.CL of $0.50 is outstanding. The trader is offering to sell an R.BU for $0.40 when the trader could have bought a D.CL at the outstanding ask for $0.50 and sold the unit portfolio (for $1) for an immediate net price of $0.50 for one less R.BU. Since most traders hold cash, the portfolio half of this transaction is generally not even necessary. Again, similar arguments hold for larger numbers of contracts.

[20] The existence of arbitrage opportunities in financial markets is not atypical. Research shows that naturally occurring financial markets often violate no-arbitrage restrictions and weak form market efficiency (e.g., Galai 1978, Bhattacharya 1983, Evnine and Rudd 1985, Halpern and Turnbull 1985, Whaley 1986, Figlewski 1989, Followill and Helms 1990, Chan and Chung 1993, and Sternberg 1994, among others). Using the IEM to study arbitrage and efficiency has several distinct advantages over this research. For example, research using naturally occurring markets cannot test price efficiency relative to nonpublic information since the econometric data set does not include such information. It is also difficult to identify or analyze the sources of inefficiencies, because the data is not trade by trade and attributable to specific traders with known characteristics.

[21] For example, suppose a trader wishes to buy an R.BU, the outstanding ask for an R.BU is $0.30 and the outstanding bid for a D.CL is $0.60. The trader should make the direct purchase. The direct purchase price of the R.BU is $0.30 at the ask, while the indirect purchase price is $1 (from the portfolio purchase) minus $0.60 (from the D.CL sale at the bid) for a net price of $0.40 for a net transaction of 1 additional R.BU. Alternatively, suppose a trader wishes to buy an R.BU, the outstanding ask for an R.BU is $0.50 and the outstanding bid for a D.CL is $0.60. The trader should make the indirect purchase. The direct purchase price of the R.BU is $0.50 at the ask, while the indirect purchase price remains $0.40. Since most traders hold unit portfolios, the portfolio half of the indirect purchase is generally not even necessary.

In this case the trader could either sell to the bid queue (by accepting the bid or crossing the queue with an ask lower than this bid) or purchase the complementary contracts at $\sum_{j\neq i} A_t^j$ and sell a unit portfolio for \$1.[22] In all these cases, the trader is trading at preexisting bids and asks. Thus, the trader is a price taker.

We classify a trade as a price-taking violation if the trader purchases a contract at the ask when this is not the lowest available price or the trader sells a contract at the bid when this is not the highest available price.[23] These are violations of individual rationality in the sense that traders who submit these orders give up certain immediate better prices. Traders who observe differences between the offer prices for securities via the two available means can offer to sell at relatively high prices and buy at relatively low prices. Should their offers be accepted, they could make immediate profits by trading in the complementary contract(s) and executing a unit portfolio transaction. Thus, knowing current bids and asks and knowing that such violations occur, traders who strategically place bids and asks can profit on average, resulting in profits for market making that equal the amounts of the losses incurred by the price takers.

Table 2 shows the frequency and size of price-taking violations. Out of 2,039 accepted bids and asks, 37.7% were not at the best price, violating the law of one price. The average violation size (dollar size of the violation $\times$ the actual traded quantity) was 10.4 cents.[24] Since the average accepted bid or ask was for a total value of \$8.03 and there were no transactions costs, this represented an average violation of 1.3% of a typical transaction's value. Again, violations of the same relative size in naturally occurring futures markets would be large in absolute terms.

4.2. Trader Role, Market Structure, and Violation Frequencies

The idea that traders take on different roles in markets is not new. Working (1958, p. 193) postulated that traders will self-select into different roles, stating "different traders [will] seek out and use different sorts of available information; and if at any time some sort of available and useful information is being generally neglected, someone is likely to soon discover that that neglect offers him a profitable field to exploit." Modern market microstructure literature (e.g., Kyle 1985) recognizes the importance of the objectives, information, and behavior of market makers in determining prices. Here, market makers perform an obvious critical role in determining prices. Instead of being rigidly segregated into roles (as in Kyle), traders self-select into roles (à la Working). Nevertheless, market makers may differ from average traders in the market and, as a result, prices set by them may differ from those that would be predicted by assuming that prices are simply set by the average or median trader in the market.

The observations above suggest a difference between market-making behavior and price-taking behavior. Traders in the role of market makers (setting the best outstanding bids and asks) are much less likely to produce violations than traders acting as price takers (accepting outstanding bids and asks as trading prices). Table 3 contains univariate analyses of several categorical variables in the data set including the role of the trader in the transaction. It shows the clear significance of market making versus price taking in determining violation frequencies ($\chi^2 = 974$ with 1 degree of freedom). We will discuss the other independent variables in Table 3 later.

Market microstructure literature also suggests that nuances of market structure can affect how prices are formed. Here, the market rules suggest a possible difference between the frequencies of violations on the buy side versus the sell side of the market. All transactions must be cash-covered in the IEM. This means that a trader must have sufficient cash on hand to place a best bid for a contract or to purchase it at the ask. A trader must also have enough contracts

[22] For example, suppose a trader wishes to sell an R.BU, the outstanding bid for an R.BU is \$0.50 and the outstanding ask for a D.CL is \$0.60. The trader should make the direct sale. The direct sale price of the R.BU is \$0.50 at the bid, while the indirect sale price is \$1 (from the portfolio sale) minus \$0.60 (from the D.CL purchase at the ask) for a net price of \$0.40 for a net transaction of 1 less R.BU. Alternatively, suppose a trader wishes to sell an R.BU, the outstanding bid for an R.BU is \$0.30 and the outstanding ask for a D.CL is \$0.60. The trader should make the indirect sale. The direct sale price of the R.BU is \$0.30 at the bid, while the indirect sale price remains \$0.40.

[23] At first, it may appear that we are ignoring some violations: those that may occur when a trader combines a unit portfolio transaction with a single-contract market transaction. In reality, we *do* identify such violations when they occur by identifying them as violations in the single contract portions of the transactions. To see why this is the case, consider a trader who effectively buys contract 1 by buying a unit portfolio and selling contract 2 at the best bid. This is a violation if $A_t^1 < 1 - B_t^2$. We do not consider the unit portfolio transactions directly. However, we will look at the second half of this transaction—the sale of contract 2 at the bid. When will we tag this as a violation? When the sale price is less than the price the trader could have received by buying contract 1 and selling the unit portfolio to the exchange. We will identify this as a violation if $B_t^2 < 1 - A_t^1$. Rearranging gives $A_t^1 < 1 - B_t^2$. This is exactly the same condition and, therefore, we will indeed tag this transaction as a violation. Similar arguments hold for other combinations of transactions. What we do avoids double counting of violations.

[24] Again, this may understate the true size of the violations. If the trader put in a quantity larger than the quantity available at the best bid or ask, he or she was willing to make a larger violation than the violation we account for here.

Table 3 Univariate Analysis of Violation Rates

Panel A: Categorical variables

Independent variable	Value	Violation rate: Percentage	Violation rate: Frequency	Pearson χ^2 tests: Statistic	Degrees of freedom	p-value
Role	Market making	37.71	769/2,039	974.4346[I]	1	0.0000
	Price taking	5.39	198/3,674			
Side	Ask or sale	19.42	532/2,739	23.3276[I]	1	0.0000
	Bid or purchase	14.63	435/2,974			
Education level[II]	High school	21.74	195/897	55.0790[I]	3	0.0000
	Bachelor's deg.	22.81	360/1,578			
	Master's deg.	13.49	183/1,357			
	Doctorate	15.33	151/985			
Family income level	Lowest 1/3	14.03	86/613	10.0054[I]	2	0.0070
	Middle 1/3	19.65	308/1,568			
	Highest 1/3	17.34	513/2,958			
Financial knowledge level[III]	Beginner	18.70	235/1,257	42.2261[I]	2	0.0000
	Intermediate	17.73	325/1,833			
	Advanced	11.34	219/1,931			

Panel B: Continuous variables

Independent variable	Statistics conditional on no violation: Mean	Std. err.	Obs.	Statistics conditional on violation: Mean	Std. err.	Obs.	t-test for difference in means: Statistic	p-value
ln(Order size)	0.9786	0.0224	4,746	0.6879	0.0483	967	−5.3662[I]	0.0000
ln(Order number)	5.0791	0.0341	4,746	3.7615	0.0645	967	−16.2881[I]	0.0000
Daily dollar volume	138.48	1.6561	4,746	159.73	3.7921	967	5.2476[I]	0.0000
ln(Relative spread)	−0.2551	0.0169	4,746	−0.0497	0.0428	967	4.8654[I]	0.0000

[I]Significant at the 95% level of confidence.
[II]A category for "other" was dropped due to unclear interpretation.
[III]Financial market experience and participation give similar results.

on hand in order to place a best ask for a contract or to sell it at the bid. Thus, to avoid violations as discussed above, a trader may need to have either \$1 in cash or a unit portfolio of contracts. If either budget constraint binds more often, we may see violations on one side of the market or the other more often.[25] Table 3 shows that this is indeed the case. There is a significant relationship between violation frequencies and market side ($\chi^2 = 23.33$ with 1 degree of freedom).

Table 4 gives the results of a logistic regression, with violations as the dependent variable. The independent variables "role" and "side" are dummy variables. Role takes on the value of 1 for market-making actions. Side takes on the value of 1 for buying transactions. Again, the results show that market makers are much less likely to produce violations than price takers and the buy side of the market is much less likely to produce violations, than the sell side (with $z = -22.043$ and $z = -6.716$, respectively). We will discuss the other independent variables in Table 4 later.

Last, we ask whether the size of the commitment that the trader is making affects the probability of violations. Another independent variable in Table 3 and Table 4 is ln(Order size), which gives the natural log of the total dollar size of a bid, ask, or trade (dollar value × quantity ordered or traded). Both the univariate and multivariate analyses give the same result. Traders submitting or executing larger orders are much less likely to produce violations ($t = -5.3662$ in the univariate difference in means test and $z = -4.430$ for significance of the coefficient in the regression).[26]

[25] Note that \$1 can be converted to and from a unit portfolio of contracts very easily. Thus, in reality, these are only truly independent constraints for traders who have very low cash and contract balances. However, traders may perceive them as independent in the same way as they apparently fail to perceive the sale of a single contract as equivalent to the sale of a unit portfolio and the purchase of the other contracts in the unit portfolio.

[26] Raw order size gives similar results. However using the log of order size reduces the effects of a few outliers and may better reflect the decreasing marginal utility of increments to wealth.

Table 4 Logistic Regression on Violations

Independent variable[II]	Estimated coefficient	Std. err.	z-statistic
Role	−2.2832	0.10358	−22.043[I]
Side	−0.6586	0.09807	−6.716[I]
ln(Order size)	−0.1398	0.03156	−4.430[I]
Education level	−0.0737	0.04956	−1.488
Family income class	0.1981	0.08071	2.455[I]
Financial market knowledge	−0.2238	0.07044	−3.177[I]
ln(Order number)	−0.1864	0.02570	−7.255[I]
Daily dollar volume	0.0009	0.00042	2.266[I]
ln(Relative spread)	0.3141	0.04250	7.391[I]
Constant	0.8588	0.2317	3.707[I]

Model classification table

Predicted	Observed negative (nonviolation)	Observed positive (violation)	Total
Negative	3,276 (86.94%)	492 (13.06%)	3,768 (100%)
Positive	115 (34.33%)	220 (65.67%)	335 (100%)
Total	3,391 (82.65%)	712 (17.35%)	4,103 (100%)

(Dependent variable = 1 if violation occurs)

Log likelihood = −1,463.887	$\chi^2(10) = 858.82$
Number of obs. = 4,103	Prob. > $\chi^2 = 0.0000$
Model sensitivity: 23.31%	Pseudo-$R^2 = 0.2268$
Model specificity: 96.64%	Area under ROC curve = 0.8138

[I]Significant at the 95% level in two sided tests.
[II]Independent variables are defined as follows: Role = 1 if market making, 0 if price taking; Side = 1 if buying, 0 if selling; ln(Order size) = natural log of the dollar size of the order (price × quantity); Education level = 1 if high school, 2 if bachelor's degree, 3 if master's degree, 4 if doctorate; Family income class = 1 if lowest 1/3, 2 if middle 1/3, 3 if highest 1/3; Financial market knowledge = 1 if beginner, 2 if intermediate, 3 if advanced; ln(Order number) = natural log of the trader specific order number; Daily dollar volume = total dollar volume on the day of the trade; ln(Relative spread) = the log of the spread in the traded contract divided by the spread in the complementary contract. See Table 1 for more detailed descriptions.

In summary, larger, buying-side and market-making traders are less likely to produce violations than smaller, selling-side, and price-taking traders. Overall, the market-making traders, who are setting prices, appear more rational than average traders in the market. Later we will show that market makers tend to be more experienced and educated.

4.3. Trader Demographics and Violation Frequencies

IEM traders were asked to complete surveys periodically. An initial survey asked traders about demographic information including sex, race, religion, age, income, education level, academic affiliation, and major. A later survey asked about financial market knowledge, prior financial market experience, and current financial market participation levels. Table 1 describes these variables in more detail. We investigated each to detect whether it significantly affected violation frequencies. We found several of these (self-reported) variables that did across both univariate and multivariate analyses: education level (though weakly at best in the multivariate analysis), family income class level, and financial market knowledge.

Table 3 shows the significance of education level alone in determining violation frequencies ($\chi^2 = 55.079$ with 3 degrees of freedom). While the overall effect of education is strong, higher education levels do not always result in lower violation rates. Recall that Table 4 gives the results of a logistic regression with violation as the dependent variable. The independent variable *education level* represents the four levels of education described in Table 1. The results show that the effect of education is not strong in the multivariate regression ($z = -1.448$).

Table 3 shows the significance of family income class alone in determining violation frequencies ($\chi^2 = 10.0054$ with 2 degrees of freedom). Taken alone, the direction of the effect is ambiguous. However, regressions that consider the interactions of more demographic variables can help clarify this. The independent variable *family income class* in Table 4 represents the three levels of family income as described in Table 1. The results show that higher reported family income class significantly increases violations ($z = 2.455$). One interpretation is that, in terms of marginal utilities, the cost of making violations is lower for higher-income individuals. We also wonder whether individuals, particularly students, all held the same definitions of "family" when estimating family income and/or were unable to assess accurately their family's income levels relative to national standards.

Table 3 shows the significance of financial market knowledge alone in determining violation frequencies ($\chi^2 = 42.2261$ with 2 degrees of freedom).[27] Higher levels of financial market knowledge reduce violation frequencies. The independent variable *financial market knowledge* in Table 4 represents the three levels of financial market knowledge as described above. The results show that more knowledge significantly reduces violations ($z = -3.177$).

We also ask whether market-specific experience affects violation frequencies. The exchange identified each order by a trader-specific order number. (A trader's first order was numbered 1, the second order numbered 2, etc.) Orders could be any bid, ask, trade, or unit portfolio transaction (with other traders

[27] The self-reported levels of financial market experience and participation give essentially identical results. They are highly co-linear and, as a result, cannot all be included in the regression. Financial knowledge had (marginally) higher explanatory power in the regressions, so we chose to focus on it here.

or directly with the exchange). As a proxy for market-specific experience, we included the log of the trader-specific order number in the logistic regression with violations as the dependent variable. Again, Table 3 and Table 4 give results with ln(Order number) representing the log of the trader-specific order number. The results show more experience with the market, as given by higher order numbers, significantly reduces the probability of a violation ($t = -16.2881$ in the univariate difference in means test and $z = -7.255$ for significance of the coefficient in the regression).

In summary, more experienced and knowledgeable traders produced fewer violations. More education reduced violation rates through a master's degree, but apparently not beyond. In contrast, higher-income classes increased violation rates.

4.4. Market Characteristics and Violation Frequencies

Next, we ask whether the state of the market itself affects the chances of observing violations. We find that market activity (as measured by total daily dollar volume in the market) and relative market liquidity (as measured by relative bid-ask spreads) both effect violation frequencies. The market-level results here reinforce the individual level results from above. Again, Table 3 gives univariate results and the logistic regression results are contained in Table 4.

The variable *daily dollar volume* gives the total dollar trading volume on the date on which each order was submitted. Higher volume implies that more traders are accepting the outstanding bids and asks. The results above suggest that more activity from price takers should increase violations. Indeed, the results here coincide with this. Greater dollar volume increases the violation rate ($t = 5.2476$ in the univariate difference in means test and $z = 2.266$ for significance of the coefficient in the regression).

The variable ln(Relative spread) gives the log of the spread for the ordered contract relative to the complementary contract in the market. It is simply the bid-ask spread for the ordered contract divided by the bid-ask spread in the complementary contract in the same market.[28] More market-making activity should reduce the spread. The results above suggest that market makers produce relatively few violations. Again, the results coincide with this. They show that a higher relative spread (less market-maker competition) increases the frequency of violations ($t = 4.8654$ in the univariate difference in means test and $z = 7.391$ for significance of the coefficient in the regression).[29]

In summary, markets with more competitive market-making (given by lower relative bid-ask spreads) and markets with less relative price-taking behavior (given by lower total dollar trading volumes) produce fewer violations.

4.5. Predictability of Violation Frequencies

Overall, violation frequencies are relatively predictable. Table 4 shows a pseudo-R^2 of 0.2268. Using a predicted probability of 0.5 as a cutoff for prediction, the classification table shows that 65.67% of actions that were predicted to produce violations actual did so and only 13.06% of actions predicted not to produce violations resulted in actual violations. Figure 1 illustrates the predictive power in more detail. The figure breaks the data into twenty five-percentile ranges for predicted violation probabilities. It shows the close relationship between the predicted violation rate and actual violation rate in each range. It also shows a great deal of variance across ranges, indicating a high power to discriminate between actions that are likely to result in violations versus those that are unlikely to do so.

4.6. Characteristics of Market Makers and Price Takers

Knowing that market makers produce fewer violations, we ask who chooses to become (or remain) a market maker. Can we characterize market makers versus price takers? In this section, we analyze the choice between market-making and price-taking behaviors—independent of violations.

Table 5 gives the results of a logistic regression with the dependent variable taking on the value of 1 if the act represented market-making behavior and 0 if the act represented price-taking behavior. The demographic information that proved important in the regression includes *education level*, *family income class*, *sex*, and *religious affiliation dummy* as described in Table 1. The demographic variable *race* was not included due to insufficient representation in several categories. *College major* was not included due to

[28] To avoid what would appear to be a locked market when bid and ask queues cross, the spread is actually defined as the minimum of the ask minus the bid and 0.001 (the smallest pricing increment in the market). The relative spread is the spread in the contract divided by the spread in the complementary contract in the unit portfolio.

[29] Absolute spreads give similar results. Note, however, that the two transactions being compared here are an acceptance in the market for one contract versus an acceptance in the complementary contract's market combined with a unit portfolio transaction with the exchange. If a trader randomly chose between these two transactions, an acceptance in the market with the higher relative spread is more likely to produce a violation. On the other hand, rational traders will try to avoid the higher level of transactions costs by trading in the side of the market with the lower relative spread. Thus, since the probability of a violation resulting from an irrational trader and the side of the market chosen by a rational trader both depend on the relative spread, the relative spread is the appropriate variable to focus on here.

Figure 1 Predicted vs. Actual Violation Rates According to the Logistic Regression Model in Table 4

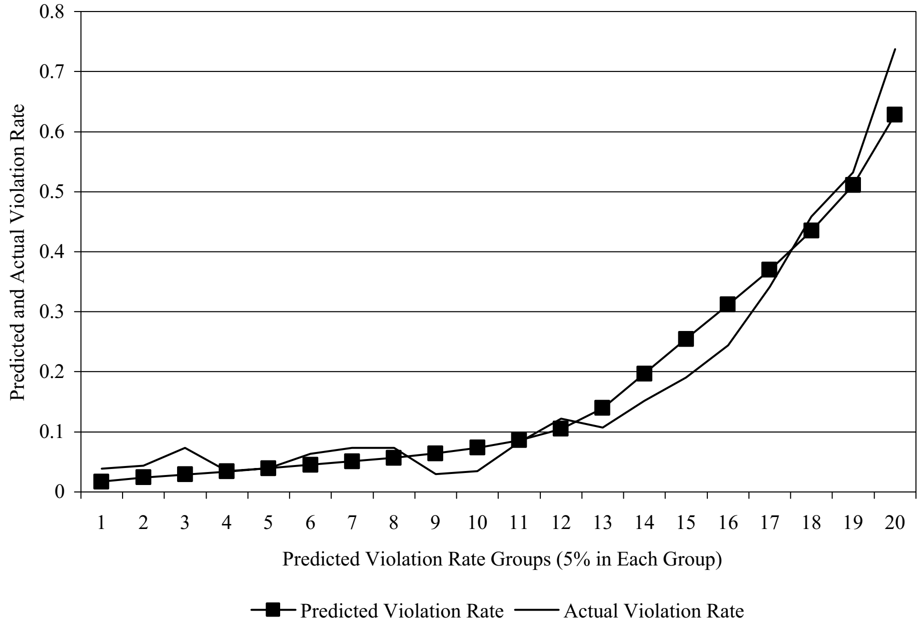

low response rates, colinearity with other variables (most notably age, income, and education), and biases it may introduce by eliminating a large, nonrepresentative portion of the sample. Finally, we include ln(Order number) as a proxy for the trader's market-specific experience level. Other variables discussed above were not significant in explaining the choice between market making and price taking. As Table 5 shows, those who chose to make markets tended to be male traders with more market-specific experience and education and who were less likely to report a religious affiliation. On the other hand, traders with more general financial market knowledge tended not to become market makers and, hence, trade through price-taking behavior. This may be similar to the effect documented by Barber and Odean (2001) that more overconfident traders are more likely to take actions that harm themselves.

Table 5 Logistic Regression on Trader Role

Independent variable[II]	Estimated coefficient	Std. err.	z-statistic
Sex	−0.9770	0.1252	−7.803[I]
ln(Order number)	0.1198	0.0171	7.009[I]
Religious Affiliation dummy	−0.3481	0.0751	−4.635[I]
Education level	0.1083	0.0338	3.202[I]
Financial market knowledge	−0.1176	0.0476	−2.473[I]
Constant	0.2706	0.1399	1.933

Model classification table

Classified	Observed negative (price taking)	Observed positive (market making)	Total
Negative	170 (55.37%)	137 (44.63%)	307 (100%)
Positive	1,381 (35.15%)	2,548 (64.85%)	3,929 (100%)
Total	1,551 (36.61%)	2,685 (63.39%)	4,236 (100%)

(Dependent variable = 1 if market making)

Log likelihood = −2,697.30	$\chi^2(5) = 170.42$
Number of obs. = 4,236	Prob. > $\chi^2 = 0.0000$
Model sensitivity: 94.90%	Pseudo-$R^2 = 0.0306$
Model specificity: 10.96%	Area under ROC curve = 0.6138

[I]Significant at the 95% level in two sided tests.

[II]Independent variables are defined as follows: Sex = 1 if female, 0 if male; ln(Order number) = natural log of the trader specific order number; Religious affiliation dummy = 1 if reported a religious affiliation, 0 otherwise; Education level = 1 if high school, 2 if bachelor's degree, 3 if master's degree, 4 if doctorate; Financial market knowledge = 1 if beginner, 2 if intermediate, 3 if advanced. See Table 1 for more detailed descriptions.

While Table 5 shows several significant factors, the choice of role remains largely unexplained, in contrast to the predictability of violation rates. The pseudo-R^2 of the regression is only 0.0306. Using a predicted probability of 0.5 as a cutoff for prediction, the classification table in Table 5 shows that 64.85% of times traders were predicted to take market-making actions they actually did so. However, they also took market-making actions 44.63% of the time when they were predicted not to do so.

Figure 2 breaks the data into twenty five-percentile ranges for predicted probabilities of market-making behavior. It shows a close relationship between the predicted violation rate and actual violation rate in each range. However, it shows little variance in predicted rates across ranges, indicating a low power to discriminate between times when traders are likely to engage in market-making behavior versus price-taking behavior.

Figure 2 Predicted vs. Actual Role Choice According to the Logistic Regression Model in Table 5

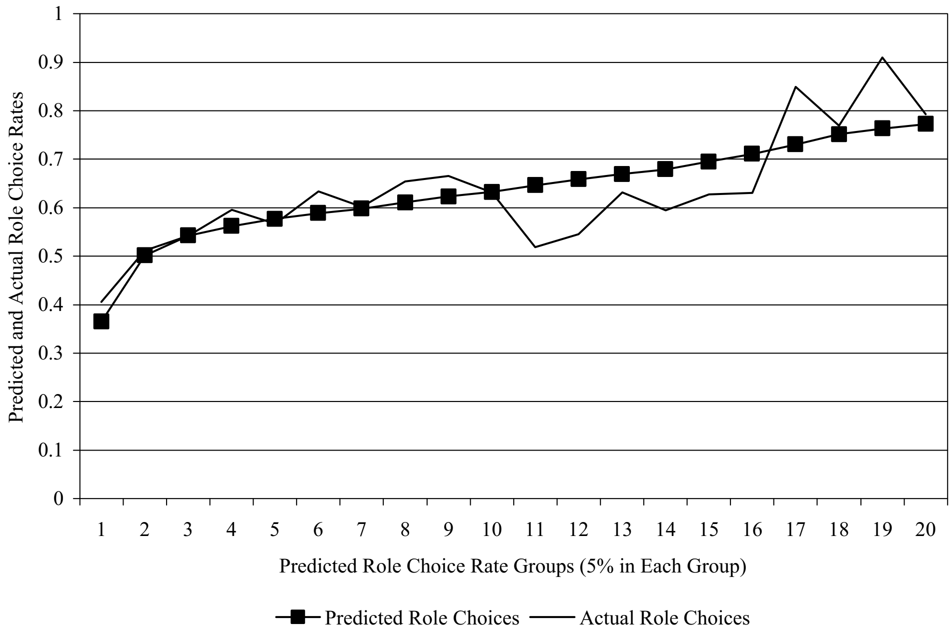

4.7. A Note on the Interpretation of the Results

We present the main analysis of this paper on a trade-by-trade basis. This allows us to account for things like the increasing experience of the trader and order-specific variables (e.g., the size of the order, the state of the market, the relative spread, etc.) on the frequency of errors. Because the same traders can act as market makers and price takers, strictly speaking, one should qualify the main results: When acting as market makers, traders appear more rational. However, we show here that the main results also hold on a trader-by-trader basis. Thus, on a trader-by-trader basis, market makers appear more rational.

Our data set consists of 5,713 price-taking or market-making orders placed by 385 unique traders. Of these traders, 76 submitted only price-taking orders and 110 submitted only market-making orders. More than half submitted more than 90% of their trades in a single role. Trader-by-trader analyses on those traders who chose only one role or that incorporate the frequency of role choice mirror the main results of the paper. In particular, traders who always chose to be price takers made errors an average of 47% of the time (whether averaging across traders or orders). In contrast, traders who always chose to be market makers made errors an average of 8% of the time (first averaging across orders for a given trader and then averaging across traders) to 10% of the time (averaging across all orders). For other traders, the correlation between the percent of time they choose to be market makers and their error rate was −0.2882 ($p = 0.0000$) treating each order as an observation and −0.5064 ($p = 0.0000$) treating each trader as an observation. As for role choice, gender remains the primary explanatory variable with a trader-by-trader correlation with average role choice of −0.2320 ($p = 0.0000$).

As a result, either interpretation is correct. When acting as market makers, traders make fewer mistakes and, hence, appear more rational. Alternatively, market-making traders make fewer mistakes and, hence, appear more rational.

5. Summary and Discussion

We have documented a high frequency of apparently irrational trader behavior in markets that are nevertheless remarkably efficient. What drives the market efficiency in spite of this irrational behavior? In this market, trade quantities are determined by price-taking traders who accept outstanding limit orders. Prices are determined by the limit orders submitted by market-making traders. Traders self-select into these roles. The evidence here shows that market-making traders, who determine market prices, are far less mistake prone than price-taking traders. These traders also had more market-specific experience and general education on average. Other evidence (Forsythe et al. 1992, 1999) shows that traders who regularly submit orders at prices near the market are less prone to particular psychological biases. The prices set by these apparently more "rational" traders are efficient. While there are mistake-prone "suckers" in this market, market makers set prices that do not appear systematically biased by them.

This paper contributes to the market microstructure literature by providing direct evidence of trader roles with abilities that differ across the roles. While some traders appear "irrational," other, more rational, traders set market prices. One can interpret price-taking traders who make mistakes as noise traders. However, the roles are not rigidly defined as they often are in the literature (e.g., Kyle 1985). Roles here are self-selected and fluid. The dynamics are consistent with Working's (1958, p. 193) theory of anticipatory prices where he allows for different traders to take on different roles including "a small group of other traders with a low level of trading competence." In his model, the result is a level of "ill-informed and inept trading," that turns out to be "without substantial price effect." Here, we show that prices can be efficient even with a relatively large amount of "inept" trading. This highlights the importance of recognizing the different roles taken on by traders in markets and the ways traders interact through markets to determine trade quantities and prices.

This paper contributes to the efficient markets literature by showing how a market can be efficient in spite of theoretical arguments against the possibility of efficiency and in spite of being populated by traders who violate usual assumptions of rational maximizing behavior. The results suggest a source of profit for informed market-making traders. Because the market is a zero sum game, for every dollar a price taker loses as a result of a mistake, a market maker gains the dollar.[30] While profiting from the mistakes of others, market makers have an incentive to set efficient prices. If they do not bracket the best guess of the future value of the contracts, market makers face adverse selection losses as other, better-informed traders trade against them. Minimizing this risk provides the incentive to gather information and impound it in bids, asks, and, hence, prices.

This paper adds to the growing literature on prediction markets (the IEM and markets like it designed to forecast future events).[31] The evidence here and elsewhere (e.g., Forsythe et al. 1992, 1999; Berg et al. 2003) shows that efficient forecasts from prediction markets do not rely on a representative sample of traders. The evidence here suggests that the market simply needs a more rational subset of traders who self-select into market-making roles and are good forecasters. Other, mistake-prone and likely biased traders provide market liquidity and profits for the market makers.

This paper also contributes important evidence for the debate on behavioral finance. In a market, it documents consistent systematic mistakes on the parts of traders that conflict with traditional rational agent based theories. The results provide support for the conjecture that biases or mistakes observed in individual choice settings will also show up in market settings. There are some welfare implications. Mistakes result in transfers from price takers to market makers. However, evidence does not support the conjecture that biases and mistakes affect prices. Overall, prices appear efficient, so price-taking traders do not incur substantial losses resulting from trades at prices that deviate significantly from fundamental values. While it is tempting to conjecture that behavioral anomalies affect market prices, the evidence here does not support this conjecture.

An online appendix to this paper is available at mansci.pubs.informs.org/ecompanion.html.

Acknowledgments
The authors thank Joyce Berg, Carol Fethke, Robert Forsythe, Ravi Jaganathan, Forrest Nelson, and several anonymous referees for helpful discussions, comments, and suggestions, and Drew Lechter for programming assistance.

[30] This differs from the common assumption that their profits are driven solely by fundamentally mis-priced assets (e.g., Grossman and Stiglitz 1980).

[31] See Berg et al. (2004) for a description of the wide range of markets run on the IEM, including those designed to forecast other political events, movie box office receipts, corporate earnings, returns, stock prices, etc. Many similar markets have been run in other countries to predict election outcomes and political events. The IEM website (www.biz.uiowa.edu/iem) lists many of them. The fictitious currency Foresight Exchange (http://www.ideosphere.com) attempts to predict a wide range of social, political, and scientific events/issues. The Hollywood Stock Exchange (http://www.hsx.com) is a fictitious currency version of the IEM's markets designed to predict movie box office takes. Several markets have obvious applications as decision support tools. For example, Ortner (1997, 1998) and Plott (2000) ran markets designed to predict internal metrics at corporations that are important for corporate decisions. Berg and Rietz (2003) discuss how markets with contracts designed to predict conditional events can support decisions.

References

Barber, B. M., T. Odean. 2001. Boys will be boys: Gender, overconfidence and common stock investment. *Quart. J. Econom.* **116** 261–292.

Berg, J. E., T. A. Rietz. 2003. Prediction markets as decision support systems. *Inform. Systems Frontiers* **5** 79–93.

Berg, J. E., F. D. Nelson, T. A. Rietz. 2003. Accuracy and forecast standard error of prediction markets. Working paper, Department of Finance, University of Iowa, Iowa City, IA.

Berg, J. E., R. Forsythe, F. D. Nelson, T. A. Rietz. 2004. Results from a dozen years of election futures markets research. C. A. Plott, V. Smith, eds. *Handbook of Experimental Economic Results*. Elsevier Science, Amsterdam, The Netherlands.

Bhattacharya, M. 1983. Transactions data tests of efficiency of the Chicago Board Options Exchange. *J. Financial Econom.* **12** 161–185.

Borch, K. 1960. The safety loading of reinsurance premiums. *Skand. Aktuar.* **43** 163–184.

Caspi, Y. 1974. Optimum allocation of risk in a market with many traders. J. H. Dréze, ed. *Allocation Under Uncertainty: Equilibrium and Optimality*. MacMillan Press, London, England, 89–97.

Chan, K., Y. P. Chung. 1993. Intraday relationships among index arbitrage, spot and futures price volatility, and spot market volume: A transactions data test. *J. Banking Finance* **17** 663–687.

Daniel, K. D., D. Hirshleifer, A. Subrahmanyam. 2001. Overconfidence, arbitrage, and equilibrium asset pricing. *J. Finance* **56** 921–965.

Daniel, K. D., D. Hirshleifer, S. H. Teohc. 2002. Investor psychology in capital markets: Evidence and policy implications. *J. Monetary Econom.* **49** 139–209.

Fama, E. F. 1970. Efficient capital markets: A review of theory and empirical work. *J. Finance* **25** 383–417.

Fama, E. F. 1991. Efficient capital markets II. *J. Finance* **46** 1575–1617.

Evnine, J., A. Rudd. 1985. Index options: The early evidence. *J. Finance* **40** 743–756.

Figlewski, S. 1989. Options arbitrage in imperfect markets. *J. Finance* **44** 1289–1311.

Followill, R. A., B. P. Helms. 1990. Put-call-futures parity and arbitrage opportunity in the market for options on gold futures contracts. *J. Futures Markets* **10** 339–352.

Forsythe, R., T. A. Rietz, T. W. Ross. 1999. Wishes, expectations and actions: Price formation in election stock markets. *J. Econom. Behavior Organ.* **39** 83–110.

Forsythe, R., F. D. Nelson, G. R. Neumann, J. Wright. 1992. Anatomy of an experimental political stock market. *Amer. Econom. Rev.* **82** 1142–1161.

Galai, D. 1978. Empirical tests of boundary conditions for CBOE options. *J. Financial Econom.* **6** 187–211.

Grossman, S. J., J. E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *Amer. Econom. Rev.* **70** 393–408.

Halpern, P. J., S. M. Turnbull. 1985. Empirical tests of boundary conditions for Toronto stock exchange options. *J. Finance* **40** 481–500.

Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* **53** 1315–1335.

Malinvaud, E. 1974. The allocation of individual risks in large markets. J. H. Dréze, ed. *Allocation Under Uncertainty: Equilibrium and Optimality*. MacMillan Press, London, England, 110–125.

Nofsinger, J. R. 2002. *The Psychology of Investing*. Prentice Hall, Upper Saddle River, NJ.

Ortner, G. 1997. Forecasting markets: An industrial application, Part I. Working paper, Department of Managerial Economics and Industrial Organization, University of Technology, Vienna, Austria.

Ortner, G. 1998. Forecasting markets: An industrial application, Part II. Working paper, Department of Managerial Economics and Industrial Organization, University of Technology, Vienna, Austria.

Plott, C. R. 2000. Markets as information gathering tools. *Southern Econom. J.* **67** 1–15.

Smith, A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Reprinted 1982, Penguin English Library, New York.

Smith, V. L. 1976. Experimental economics: Induced value theory. *Amer. Econom. Rev.* **66** 274–279.

Sternberg, J. S. 1994. A reexamination of put-call parity on index futures. *J. Futures Markets* **14** 79–101.

Whaley, R. E. 1986. Valuation of American futures options: Theory and empirical tests. *J. Finance* **41** 127–150.

Working, H. 1958. A theory of anticipatory prices. *Amer. Econom. Rev.* **48** 188–199.