

Longshots, Overconfidence and Efficiency on the Iowa Electronic Market*

Joyce E. Berg and Thomas A. Rietz

Tippie College of Business

University of Iowa

Iowa City, Iowa 52242-1000

August 2012

* We thank the faculty and staff who run the Iowa Electronic Markets. For helpful comments and suggestions, we thank Matthew Billett, Robert Forsythe, Forrest Nelson, George Neumann, Toni Whited and seminar participants at Tulane University.

Longshots, Overconfidence and Efficiency

on the Iowa Electronic Market

Abstract

A basic proposition of behavioral finance is that individual decision-making biases affect financial markets. We examine this proposition using data from the Iowa Electronic Market (IEM), a real money market conducted for teaching and research purposes. Using this data, we ask whether the longshot bias or an overconfidence bias affects market prices. The IEM is ideal for addressing this question because it mixes many desirable research features from betting markets (where the longshot bias is observed) with a closer parallel to naturally occurring financial markets (where researchers look for evidence of overconfidence), and the two biases yield opposing predictions. We find that while the longshot bias affects many sports betting markets robustly, no such bias appears in IEM markets. Nor does overconfidence influence prices at short horizons. If there is a bias, it results from overconfident traders at long horizons. While the markets incorporate information efficiently at short horizons, non-market data indicate some long-horizon inefficiency. When markets appear inefficient, we calculate Sharpe ratios for static trading strategies and document returns for dynamic trading strategies to show the economic content of the inefficiencies.

Longshots, Overconfidence and Efficiency

on the Iowa Electronic Market

Using price data from the Iowa Electronic Markets (IEM), we evaluate the predictive ability of market efficiency and two propositions from behavioral finance: the longshot bias and the overconfidence bias. These biases make opposing predictions about price behavior in our setting. Due to the contract design in these markets, prices reveal trader expectations about the probabilities of particular events occurring. If the markets are efficient, a contract's price should equal the probability of a \$1 liquidation (that is, the probability of the associated event happening). We test this hypothesis against behavioral predictions of: (1) a longshot bias which would result in over-pricing low probability contracts and under-pricing high probability contracts and (2) an overconfidence bias, which would result in under-pricing low probability contracts and over-pricing high probability contracts.

The "Longshot Bias" is commonly observed in betting markets.¹ At the racetrack, bettors are willing to pay more for bets that are unlikely to pay off ("longshots") than those bets are actually worth. This results in relatively large negative average returns for longshot bets. At the same time, bettors are not willing to pay the full value for bets that are most likely to pay off ("favorites"). This results in relatively small negative (because of the track "take") or sometimes positive, average returns for favorite bets. These biases seem quite robust (see Ziemba and Hausch, 1986). Behaviorally, bettors seem attracted to the large potential returns of an unlikely winner. In the context of the IEM, traders who behave similarly would drive up the price of contracts that are unlikely to payoff. At the same time, high probability contracts (i.e., likely outcomes) would be underpriced.

Overconfidence can take several forms depending on how it is defined and interpreted. Two interpretations are of interest here. Overconfidence can affect perceived probabilities of outcomes (e.g., Lichtenstein, Fischhoff and Phillips, 1982). It can also cause over-reactions to information in financial markets (e.g., Daniel, Hirshleifer and Subrahmanyam's, 1998, misestimation of the precision of information). In Daniel, Hirshleifer and Subrahmanyam (1998), asset prices overreact to private

information signals and slowly return to efficient levels. If traders behave according to this model in the IEM, prices for contracts that information reveals are more likely to pay off will be driven up in price beyond the true likelihood of payoff. Similarly, prices for contracts that information reveals are less likely to pay off will be driven down in price beyond the true likelihood of payoff. As a result, high priced contracts are over-priced and low priced contracts are under-priced.²

While there is ample evidence that individuals have biases and traders may behave irrationally (in the context of the IEM, see Oliven and Rietz, 2004, and Berg and Rietz, 2006), there is also reason to believe that financial markets may be efficient in spite of these biases. First, conventional wisdom suggests that a few rational traders can arbitrage biases out of existence in markets. This is a debatable issue because biases that arise from a pervasive cause may be impervious to arbitrage. For example, if all traders were affected by prospect theory's over-weighting of low probability events (Kahneman and Tversky, 1979), none of the traders would recognize the bias in prices and, hence, none would drive prices to efficient levels. Even if a few do recognize the bias, whether they would have the ability or incentive to drive markets to efficient levels is an open question. Grossman and Stiglitz (1980) suggest that traders will never want to drive out all arbitrage opportunities because to do so would eliminate all profits for information collection and arbitrage activities.

Second, context may influence whether anomalies generalize to financial markets. Again, the issue is debatable. Slovic (1972) argues that context may affect observed risk preferences. Cosmides (1989) argues that context may influence the incidence of a different bias (the "confirmation bias"). Because a financial market is a different context than a betting market and financial information comes in a variety of contexts, context effects may reinforce, weaken or eliminate longshot or overconfidence biases.

Third, the market structure may reduce or reverse predicted biases. While it is not unambiguously apparent what the effect will be, some evidence suggests that market structure matters. With respect to the longshot bias, Woodland and Woodland (1994) suggest that a difference in structure (the ability to bet against teams) reverses the bias in baseball betting. Further, while Gandar, Zuber,

O'Brien and Russo (1988) get mixed results, Gray and Gray (1997) show that betting home-team underdogs in football betting markets can have positive profits. This could be interpreted as a partial reversal of the longshot bias in these differently structured markets (point-spreads, instead of odds, adjust to guarantee the bookie's take). Because financial markets differ in structure as well, one might expect structural differences to affect how biases manifest themselves. For example, Thaler and Ziemba (1988) cite the inability to short-sell racing bets as a possible contributing factor to the longshot bias. Since traders can short sell in typical financial markets, this structural difference may mitigate or eliminate the manifestation of the bias in financial markets. In the context of financial markets, Reny and Perry (2006) show that, under the right conditions, double auction markets are not prone to the winner's curse that plague one-sided auction markets. Again, structure matters.

Fourth, evidence suggests that the existence of a two-sided dynamic market itself will mitigate the effects of irrational or biased traders. Gode and Sunder (1993) show that markets do not necessarily need to be populated by perfectly rational, optimizing traders to result in prices and allocations consistent with efficient markets. Instead, they show that the market itself can drive efficient outcomes even when populated with "zero-intelligence" robot traders. Forsythe, Nelson, Neuman and Wright (1992) and Forsythe, Ross and Rietz (1999) show that while biases commonly observed in political science research (the false consensus effect and the assimilation contrast effect) affect traders at an individual level in IEM political markets, those biases do not appear to affect market prices.³ Instead, more rational traders set market prices. Thus, the existence of irrational or biased traders does not necessarily imply a market-level effect.

Here, we ask whether longshot or overconfidence biases carry over to financial markets using data from the IEM. Do traders appear to over-value assets that will pay off with low probability and under-value assets that will pay off with high probability? This would mirror the longshot bias and would be consistent with over-weighting low-probability events. Alternatively, do traders appear to under-value assets that will pay off with low probability and over-value assets that will pay off with high probability? This would mirror the overconfidence bias and would be consistent with over-weighting high-probability

events. Thus, our analysis is a test between unbiased pricing and two potential biases, one reflecting over-weighting of low probability events that is consistent with a longshot bias and one reflecting over-weighting of high probability events that is consistent with an overconfidence bias.

Whether longshot or overconfidence biases affect prices is very difficult to answer using naturally occurring financial market data.⁴ Unlike a betting market, there is no definitive point in time when the value of a stock is known with certainty. This makes it difficult to determine whether prices at any point in time are actually biased. In addition, stocks, options and futures are all very complex bets. The value of a stock should depend on the distribution of dividends over the infinite horizon. The value of an index depends on the values of its component stocks. The values of options and futures, in turn, depend on the future distributions of the values for underlying stocks or indices. The complexity of these “bets” would make it difficult to detect the effects of simple biases like the longshot bias even if we knew the “true” values of the underlying assets. Further, repeated observations under similar conditions give betting markets a definite edge in testing. But, betting markets have their own significant drawbacks and results may not generalize to financial markets. Thus, neither betting markets nor naturally occurring financial markets are the ideal testing grounds to determine whether these biases affect financial markets.

We use the Iowa Electronic Markets (IEM) to test for longshot and overconfidence bias effects in real-money, real-time, repeated and simple financial markets. There are many advantages of using the IEM. It is a real financial market in which traders buy and sell assets on their own account, bearing the real-money risks and reaping the real-money rewards of their activity. IEM participants trade simple contingent claim assets (called “contracts”) that payoff \$1 in one state of the world and \$0 otherwise. These contracts parallel the simplicity of racetrack bets (which either pay off or not depending on the outcome of the race). IEM contracts liquidate at predetermined dates with no uncertainty about their value at that time and the market repeats itself under essentially identical conditions monthly. Finally, the structure of the market and the contracts traded imply that the price of a particular contract at any point in time should equal the market consensus probability that the contract will pay off \$1 on the liquidation date. Thus, prices can be compared to objective estimates of payoff probabilities (from multinomial logit

models) to determine whether biases exist. At a given point in time, prices should be sufficient statistics for estimating payoff probabilities. This makes testing for efficiency simple. If adding past price information or outside information to IEM price information increases the explanatory power of logit models in predicting payoffs, then the markets are not efficient.

These features make using the IEM well suited for studying the longshot and overconfidence biases because of similarities between IEM contracts and racetrack bets. In addition, many structural features of the IEM more closely resemble naturally occurring financial markets than racetracks. The IEM markets we examine are two-sided markets with an intermediate horizon (up to 5 weeks). News about relative payoff probabilities is revealed during trading on the IEM. Traders can synthetically short sell in the IEM.⁵ Because IEM traders know the price at the time they trade, they effectively know the “odds” at the time they place their bets (in contrast to U.S. racetracks). IEM traders can reverse their positions at any time by undertaking opposite transactions at market prices. All of this trading can be done without paying the track’s take with each position change. Instead, traders pay the bid/ask spread as they would in any other financial market. Because these features parallel naturally occurring financial markets, examining IEM data will help determine whether we might expect the longshot bias to generalize to financial markets in a straightforward manner or whether the evolution of market prices is consistent with overconfident traders.

Our results accord well with a transitory overconfidence effect on prices. The prices move exactly as predicted by Daniel, Hirshleifer and Subrahmanyam, 1998. Initial trading prices in the IEM markets appear relatively noisy, but unbiased as one would expect from a market with little information. Prices display an overconfidence bias at intermediate horizons when some time has passed and some information has come in, but considerable uncertainty about outcomes remains. The bias disappears as more information arrives and the liquidation date approaches.

In the next section, we discuss the IEM in more detail and show how biases would manifest themselves in these markets. Then, we describe the testing procedure in detail, give the results and end with conclusions and discussion.

I. Biases, Efficiency and Market Price Predictions

A. The Iowa Electronic Markets and Contracts

The Iowa Electronic Markets (IEM) are real-money, real-time, futures markets operated as a not-for-profit teaching and research tool by the Tippie College of Business at the University of Iowa.

Relative to typical experimental markets, they are large scale and long duration. The markets discussed here have thousands of registered traders and, often, hundreds of active traders. They run from four to five weeks and have been liquidated and reinitialized monthly for more than five years. The IEM also differs from typical experimental markets in that the traders invest their own money to trade, bearing the real-money risks and reaping the real-money rewards of their activity.⁶

The IEM market structure closely parallels naturally occurring financial markets. It operates as a continuous electronic double auction that traders access through the Internet. Traders can place both limit and market orders. Outstanding bids and asks are maintained in price and time ordered queues which function as continuous electronic limit order books with the best price in each queue made public.

The markets we examine are the Microsoft Price Level and Computer Industry Returns Markets.⁷ Prospectuses for these markets appear in the appendices and we describe them briefly here. Each market runs either four or five weeks (from the Monday after the third Friday of one month to the Monday after the third Friday of the next month).⁸ Each market contains a complete “slate” of contracts, one of which will liquidate at \$1 depending on the state of the world on the liquidation date. All other contracts in the market liquidate at a value of \$0. Thus, all contracts are simple state-contingent claims, similar to bets that either pay off or not, depending on the state of the world.

In the Microsoft Price Level Market, a slate consists of two contracts. One of these contracts (the “H” contract, denoted generically as “MSH”) will liquidate at \$1 if Microsoft’s actual stock price closes above a pre-determined “cutoff” value on the third Friday of the month. The other contract (the “L” contract, denoted “MSL”) will liquidate at \$1 if Microsoft’s price closes less than or equal to the cutoff.⁹ Thus, these contracts are simple binary options. (See the prospectus in Appendix I for more details.)

In the Computer Industry Returns Market, a slate consists of four contracts. Each contract corresponds to an underlying security (company stock or index). The underlying securities are Apple Computer (AAPL), IBM (IBM), Microsoft (MSFT) and the Standard and Poor's 500 Index (SP500). A contract liquidates for \$1 when its underlying security has the highest return from third Friday to third Friday.¹⁰ Thus, these contracts are also simple state contingent claims, though somewhat more complex than the contracts in the Microsoft Price Level Market. (See the prospectus in Appendix II for more details.)

B. Prices, Predictions and Models

In each market, the complete slate of contracts is a risk-free portfolio. One of the contracts will always liquidate at \$1 while the others expire worthless. Cash holdings are also risk-free. Both of these risk-free assets earn a zero return. There are no transaction fees. Complete slates can be purchased from or sold to the exchange at any time for a fixed price of \$1. This implies that the numbers of contracts of each type in a market are always the same. Thus, there is no aggregate uncertainty in this market. Because of these features, contract prices should always equal expected values of the contracts regardless of risk preferences and time remaining to liquidation.¹¹ Let p_t^j be the price of contract j on date t and $q(V_T^j = 1 | I_t) = q_t^j$ be the probability that contract j will liquidate at \$1 ($V_T^j = 1$) on date T given information available on date t (I_t). In theory, we should observe:

$$\forall j \in \{1, 2, \dots, J\}, \forall t \leq T : p_t^j = q(V_T^j = 1 | I_t) \times \$1 + [1 - q(V_T^j = 1 | I_t)] \times 0 = q_t^j \quad (1)$$

where j indexes the slate of J contracts that pay off in different states of the world and t represents any date up to the date T when liquidation values are determined.¹²

Testing will be very simple. First, following the existing racetrack literature, we use a simple frequency analysis. Then, we refine and expand the analysis using logit models¹³ to estimate the true probabilities of \$1 liquidations (i.e., the q 's) and see whether they deviate systematically from market prices (i.e., the p 's). To see how this works, consider the multinomial logit model:

$$\forall j \in \{1, 2, \dots, J\}, \forall t \leq T : q_t^j = \frac{e^{X_t b_t^j}}{\sum_{i=1}^J e^{X_t b_t^i}} \quad (2)$$

where, in addition to the variables defined above, X_t is a vector of independent variables at date t and b_t^j is a coefficient vector for contract j at date t .

To identify the model, b_t^J is arbitrarily set to 0 and contract J becomes the base contract.¹⁴ Then, the probabilities of all other contracts are computed relative to the base contract as:

$$\forall j \in \{1, 2, \dots, J-1\}, \forall t \leq T : \frac{q_t^j}{q_t^J} = e^{X_t b_t^j} \quad (3)$$

Now consider using log ratios of contract prices as independent variables.¹⁵ Since the contract price should represent the probability that a contract will liquidate at \$1 conditional on all available information, price ratios should be sufficient statistics for forecasting probabilities of payoffs. In particular, consider using as independent variables:

$$X_t = \left[1, \ln\left(\frac{p_t^1}{p_t^J}\right), \dots, \ln\left(\frac{p_t^{J-1}}{p_t^J}\right) \right] \Rightarrow \frac{q_t^j}{q_t^J} = e^{\alpha_t^j + \beta_t^{1,j} \ln\left(\frac{p_t^1}{p_t^J}\right) + \dots + \beta_t^{J-1,j} \ln\left(\frac{p_t^{J-1}}{p_t^J}\right)} \quad (4)$$

If prices do indeed equal probabilities, then it should be the case that $\alpha_t^j = 0$ and $\beta_t^{j,j} = 1$ for all j and $\beta_t^{i,j} = 0$ for all $i \neq j$. Only in this case does $q_t^j = p_t^j$ for all j . Further, if prices incorporate all available information about the probabilities of payoffs (as efficient markets would suggest), then, apart from co-linearity issues, coefficients on any other independent variables will be zero. Even if co-linearity creates the appearance of significance, adding other variables should not increase explanatory power of the model. This will serve as the basis for detecting any biases and inefficiencies that exist.

C. Biases, Efficiency and the Logit Model Estimates

Testing for the longshot bias and overconfidence will take a particularly simple form: run logit models with logged price ratios as the independent variables and check the resulting coefficients. If

traders are unbiased and the market is efficient, we should observe estimates consistent with $\alpha_t^j = 0$ and $\beta_t^{j,j} = 1$ for all j and $\beta_t^{i,j} = 0$ for all $i \neq j$ in the logit model equation (4). If we graph the logit model estimated probabilities against prices in this case, we would get a 45-degree line labeled “beta=1” in Figure 1.

In the two-state case, a longshot bias would show up if $\beta_t^{j,j} > 1$. This leads to the “beta>1” mapping in Figure 1 where prices exceed the probability of payoff for low payoff probabilities and fall short of the probability of payoff for high payoff probabilities.^{16,17} Paying too much for a low-payoff-probability contract here corresponds to betting too much on the longshot in a horse race. This is consistent with Kahneman and Tversky's (1979) proposition from Prospect Theory that people overweight the likelihood of low probability events.

In contrast, an overconfidence bias would show up if $\beta_t^{j,j} < 1$. This leads to the “beta<1” mapping in Figure 1 where prices fall below the probability of payoff for low payoff probabilities and exceed the probability of payoff for high payoff probabilities. This effect, which is the opposite of a longshot bias, can easily result from traders being overconfident in their forecasts (e.g., the overconfidence bias of Lichtenstein, Fischhoff and Phillips, 1982). For example, when the rational probability of payoff is actually 90%, overconfident traders may assess the probability at an even higher level and be willing to pay, say, 95 cents for the contract.¹⁸

Testing for efficiency is also simple. If a contract's log price ratio is sufficient for explaining its own probability of payoff, there should be no difference in explanatory power between an unrestricted logit model and a logit model with the restrictions that cross-price ratio coefficients are zero ($\beta_t^{i,j} = 0 \forall i \neq j$). If markets are weak-form efficient, adding recent prices or price changes should not increase explanatory power. Finally, if markets are semi-strong-form efficient, adding additional information (e.g., stock market returns or prices) should not increase explanatory power. On the other hand, if traders over- or under-react to the information embodied in market prices of the underlying securities, then stock market returns and prices will add significant explanatory power. These

propositions can be tested using likelihood ratio tests in a series of nested logit models (though the quantity of data available restricts the amount of testing that can be done).

We also ask whether the time horizon affects any biases that exist. Rubinstein (1985) shows a “time-to-expiration” bias in out-of-the money call options. The racetrack literature suggests that the “smart” money is bet late and that this mitigates the longshot bias somewhat late in the betting process (see Thaler and Ziemba, 1988 and Asch and Quandt, 1986). Daniel, Hirshleifer and Subrahmanyam 1998 suggest that impact of overconfidence on prices is transitory. They argue that overconfident traders will over-estimate the precision of *private* information. Prices are unaffected by the bias before information arrives, over-react to it temporarily and move to (new) unbiased values as the information is fully incorporated into prices. In context, if private information implies that a particular contract is more likely to pay off in the IEM, then the price should rise *more* than the information justifies and return to efficient levels as uncertainty is resolved. We argue that this dynamic can occur in response to the *public* information embodied in the market prices of the securities underlying IEM contracts. This is consistent with Daniel, Hirshleifer and Subrahmanyam’s (1998) idea if either (1) marginal traders over-estimate the precision of public signals or (2) private signals are correlated with the public signals.¹⁹ These ideas would create horizon effects in any biases or inefficiencies that exist. By looking at different horizons (numbers of days to liquidation value determination), we can study whether horizon effects exist. We use 1 to 21 day horizons to study these issues.²⁰

D. The Advantages of using the Iowa Electronic Markets

For many reasons the IEM is an ideal place to study expectations in financial markets and the longshot and overconfidence biases in particular. Some of these reasons mirror those cited to justify using racetrack data. For example, similar to betting markets, but in contrast to most naturally occurring financial markets, there exist well defined points in time at which the values of IEM contracts are realized with certainty. The frequency of this realization can be compared to the projected probability of

occurrence given by market prices at any prior point in time. This is a clear advantage over using stock market returns because stock values always depend on unknown future outcomes.

Several distinguishing features make the IEM better than betting markets for studying these issues. First, it is a real-money, two-sided financial futures market. Structurally, it much more closely resembles naturally occurring financial markets and financial markets as modeled in theory. Traders can take both sides of bets and unwind bets at any time as information is revealed. Traders self-select in to or out of markets at any time or can be driven out of the markets through bad trading. This resemblance to naturally occurring markets may matter. Some explanations for the longshot bias cite particular characteristics of the betting market. For example, Shin (1992) argues that the longshot bias arises from bookies setting prices to maintain their take while recognizing that knowledgeable “insiders” place some bets. On the IEM, there are no transactions costs (no “take” for the track or bookie). Of course, traders who place limit orders in the IEM face adverse selection as well, but it should be symmetric and should more closely resemble typical market microstructure theory than racetrack betting markets.

Two additional differences between the IEM and racetrack betting markets lie in trading/betting horizon and information flow. Asch and Quandt (1986) argue that late changes in odds help predict racing outcomes. Thus, the horizon and information flow into the market may affect the efficiency of the market and profitability of various strategies. But, analysis from the racetrack is limited by the relatively short betting horizon. Similarly, traditional experimental asset markets last for a few hours at most, with trading periods that typically last a few minutes. The IEM's intermediate market horizons (four or five weeks for the markets discussed here) address the gap between the short-term prospects of racetrack betting or traditional experiments and the infinite horizon of stock markets. Avoiding the valuation problems caused by the stock market's infinite horizon, this allows for analysis at various horizons while keeping repeated observations. As a result, we can test for dynamic predictions such as Daniel, Hirshleifer and Subrahmanyam's (1998) assertion that overconfidence will only manifest itself in price biases after payoff relevant information exists and before it is fully incorporated in prices.

Betting markets are evaluated for efficiency using odds during the betting period before the race is run. While the betting process aggregates information during the betting period, little news is produced during this time. Things might look considerably different if bettors could continue placing and withdrawing bets during the race, as news about actual performance is being produced, evaluated and incorporated into odds. In the stock market, news is constantly being generated about companies and their prospects. The market both incorporates news and aggregates information across traders. Paralleling the stock market, news relevant to likely outcomes arrives during IEM trading. Much of this news is available to all traders, but subject to interpretation. Some may only be known by one or a few traders. All of it should be incorporated in IEM prices if the market is efficient. This allows for testing of a much more dynamic type of information efficiency, and one that more closely resembles naturally occurring financial markets, than betting odds allow.

Finally, two other features make testing in the IEM particularly informative. The bets here are much simpler than many bets in gambling markets and certainly simpler than the bets implied in most naturally occurring financial markets. There are no particularly difficult terminologies or conventions to learn. In fact, the IEM is quite transparent. In addition, IEM markets are designed so that risk preferences should not matter. This is often a difficulty for traditional experimental research.

Overall, in a relatively simple environment, the IEM keeps many of the desirable features of the racetrack betting markets for analysis (simple asset structure and repetition) while more closely paralleling naturally occurring financial markets (market structure and trading mechanics), making for fewer worries about external validity.

II. Tests for Biases and Efficiency

A. Frequency Tests

To compare to previous research, we sort the data into cells according to observed prices and compare the average price in each cell to the objective probability determined by the actual payoff rate within the cell. This is analogous to the analysis in much of the racetrack literature where researchers

aggregate bets (contracts) across odds (prices) and compare average bets within each cell (average prices) to average payoffs (liquidation values).

Table 1 shows the average price of contracts in \$0.20 ranges aggregating across all contracts and markets at different horizons (1, 2, 4, 7, 14 and 21 days to liquidation determination). This is similar to aggregating ranges of odds across all horses and races in a data set. The table also shows average payoffs for contracts in each range and the differences between average payoffs and prices. Finally, the table shows the number of observations in each range and simple t-tests for differences between payoffs and prices. A longshot bias would result in significantly higher prices than payoffs for low-priced (longshot) contracts and significantly lower prices than payoffs for high-priced (favorite) contracts. An overconfidence bias would push prices in the opposite direction. The data is consistent with overconfidence. Prices fall significantly below payoffs for contracts falling in the \$0.0 to \$0.2 range at most horizons longer than one day. Prices are significantly above payoffs for contracts falling in the \$0.8 to \$1.0 range at 4- and 14-day horizons. Intermediate contracts appear to be priced efficiently. Table 2, which aggregates across price quintiles, shows similar, though slightly weaker, results.

There are several limitations inherent in this simple frequency analysis (and similar racetrack betting analyses). It raises aggregation issues and is not a particularly efficient use of the data. Frequency analysis ignores the interdependent nature of the contract payoffs. It fails to take into account the fact that only one contract in a multiple-contract market can payoff at \$1. Analogously, only one horse can win a race, only two can show and only three can place. This creates a negative correlation across outcomes that can be addressed by logit models (for the Microsoft Price Level Market with two contracts) and multinomial logit models (for the Computer Industry Returns Market with four contracts). In addition, analysis of the possible sources of inefficiency and the search for data that may provide better predictions both require a more detailed analysis than simple frequency analysis. Nevertheless, the frequency analysis serves as a quick summary of the data. In spite of weaknesses, its results mirror the logit models that follow, though the logit models allow much more extensive investigation while overcoming the weaknesses discussed here

B. Logit Models for the Microsoft Price Level Market

For the two-contract Microsoft Price Level Market, the model becomes a simple logit model. It accounts for the perfect negative correlation between the MSH and MSL contracts by recognizing that the normalized price of the MSL contract is one minus the price of the MSH contract and using the logged price ratio for analysis.

The “Model I” rows in Table 3 give the results of a simple logit model with the log price ratio as the only explanatory variable. Table 3 shows several sample horizons while Figure 2 summarizes estimated β coefficients at all horizons. Recall that, if prices are unbiased estimates of true payoff probabilities, then the coefficient on the log price variable should be 1 and the intercept should be zero. As Figure 2 shows, estimated slope coefficients generally fall below one. At one and two day horizons, this null is not rejected. However, at all horizons of three days and longer (including the 4, 7, 14 and 21 day sample horizons shown in Table 3), the coefficients on the logged price ratio fall significantly below one (at the 95% level of confidence), indicating an overconfidence bias. Thus, the results from simple logit models mirror the results from the simple frequency analysis.

Logit models can go beyond simple frequency analysis in testing for efficient markets. If markets are efficient in incorporating all relevant information into prices, then adding additional information should not improve the ability to predict payoffs. Model II asks whether adding the most recent change in log price ratios adds to the explanatory power of prices. Asch and Quandt (1986) suggest that late changes in odds help predict winning frequencies at racetracks.²¹ In addition, short-horizon “momentum” effects are observed in stock market data (e.g., Jegadeesh and Titman, 1993). This might carry over to financial markets if traders over- or under-react to news, creating serial correlations in prices. Likelihood ratio tests for Model II against the (restricted) Model I show that adding the change in log price ratios seldom adds explanatory power. As shown in Figure 2, the only significant likelihood ratio statistics (at the 95% level of confidence) are at the 2-, 6- and 20-day horizons.

Finally, if markets are efficient, then adding information about the current Microsoft stock price relative to the cutoff should not increase explanatory power. On the other hand, if the market is over- or

under-reacting to this particular information, then adding it should significantly increase explanatory power. The ratio of the current Microsoft stock price to the cutoff is added as an independent variable in Model III. This represents the percentage change in Microsoft's stock price required to change the payoff outcome in the IEM market. Table 3 (at sample horizons) and Figure 2 (at all horizons) give the results of likelihood ratio tests for Model III against (the restricted) Model I. The results are mixed. At several intermediate horizons (5, 7, 8, 10, 12, 14, 15, 16, 17 days), it appears that adding this information increases power. That is to say, a trader at these horizons would do better at predicting the chances of payoff by knowing the current relative level of Microsoft's price to the cutoff instead of knowing only the "H" contract price on the IEM.²² Thus, there is limited evidence against market efficiency using Model III. In Section III, we discuss the risks and return to investment strategies designed to exploit these inefficiencies.

C. Multinomial Logit Models for the Computer Industry Returns Market

Since there are four contracts in the Computer Industry Returns Market, we use a four-state multinomial logit model for estimation. Table 4 lists three models for this market and shows results for sample horizons. Figure 3 summarizes this information for all horizons. In each model, the "base" contract is the S&P500 and log price ratios are relative to the price of this contract.²³ Each model has three component equations: one to estimate the probability of an AAPL payoff (relative to S&P500), one for IBM and one for MSFT. The simplest version (Model I) restricts the coefficients to zero on all terms except the contract's own log price ratio. This model is the multinomial analog of Model I for the Microsoft Price Level Market. It recognizes the multinomial nature of the payoffs, but assumes that each contract's own log price ratio is sufficient for forecasting its probability ratio. The table lists individual t-tests for the null that each individual own price ratio coefficient equals one. It also lists joint χ^2 test statistics for the null that all three own price coefficients equal one for each horizon. The upper part of Figure 3 shows the own log price ratio coefficients for all contracts. The results mirror the Microsoft market results. At horizons of 2 or more days, all joint tests for all own-price coefficients equaling one

are rejected at the 95% level of confidence with one exception: the 10-day horizon (with a 94% level of confidence). Individual coefficients often fall significantly below one. The coefficients generally fall as the horizon increases. As one would expect from Table 1, inspection of the mappings between prices and probabilities overwhelmingly show overconfidence biases. Thus, the evidence for an overconfidence bias from the Microsoft market is paralleled in the more complex Computer Industry Returns Market.

Model II asks if the cross-price restrictions imposed in Model I are reasonable or if removing them can significantly increase explanatory power. Likelihood ratio tests of Model II against (restricted) Model I show no significant effects of relaxing these restrictions at any horizon. (Table 4 shows sample horizons and Figure 3 shows the likelihood ratio test statistics across all horizons relative to the critical value.)

Model III asks whether the markets fully incorporate the information about the relative returns of the underlying stocks by adding relative return standings to the Model I regressions. The additional variable is the return on the stock underlying a given contract minus the maximum of the other underlying returns.²⁴ Cross-price restrictions are kept in light of the Model II results. Likelihood ratio tests of Model III against (the restricted) Model I show that at all horizons adding relative return information increases explanatory power. (Again, Table 4 and Figure 3 show these results.) Thus, the markets are not making fully efficient use of all of the available information.

In summary, the IEM markets appear quite efficient at short horizons, but prices seem affected by an overconfidence biases at longer horizons. Further, typically at intermediate or long horizons, additional information can help improve explanatory power in logit models for predicting payoffs. In the next section, we discuss the risks and return to investment strategies designed to exploit these inefficiencies.

III. Investment Strategies

Given the overconfidence bias evident in the IEM, one might conjecture that there are profitable investment opportunities based on price alone. This is analogous to the observation that betting heavy favorites at the horse track can generate positive expected returns. In fact, profitability should be easier to

attain in the IEM because there are no explicit transactions costs (i.e., there is no track “take” or trading commission). To evaluate risk and returns for different contracts, we assume that the logit models give the best prediction of the actual probability of payoff. This is analogous to the usual assumption in the racetrack literature of assuming the ex-post frequencies of winning are the best estimate of the ex-ante actual frequencies. While the racetrack work typically calculates win frequencies relative to odds or the expected returns to bets at various odds levels, we extend results here to the risk-return tradeoff and to the returns to dynamic trading strategies.

A. Sharpe Ratios for Buy and Hold Strategies

The relationships between price, expected return and risk for the Microsoft Price Level Market under Model I are particularly simple and can be graphed easily. Figure 4 shows the predicted probabilities versus the prices for the MSH contracts at various horizons. The under-pricing of low-priced contracts and over-pricing of high-priced contracts is apparent from this graph. This results in extremely high-return profit opportunities at low prices, but also high risk. Figure 5 shows Sharpe ratios for contract purchases at available prices and various horizons given the payoff probability predictions of Model I. It shows how much a trader could profit relative to the risk undertaken by buying contracts at available prices and holding until liquidation. The highest ratios are attained by purchasing contracts in the \$0.10 range at two-week or four-day horizons.

Graphs for the Computer Industry Returns Market are somewhat more complicated because of the inter-dependent, multinomial nature of the contracts. Nevertheless, similar patterns can be seen at similar horizons. Figure 6 shows the Sharpe Ratios for contracts in the Computer Industry Returns Market for a 14-day horizon. Across most prices, the highest ratios are for AAPL. For any given contract, the highest ratios are in the \$0.15 and under range. At this 14-day horizon, the average AAPL Sharpe ratio was 0.267. Its median was 0.342 and the 75th percentile was 0.389. The same statistics for IBM were -0.142, -0.317, and 1.29, respectively. For MSFT they were -0.028, 0.055 and 0.181. Finally, for SP500, they were -0.015, 0.059 and 0.127. For comparison purposes, the Sharpe ratios computed on

third-Friday to third-Friday returns for AAPL stock, IBM stock, MSFT stock and the S&P500 index from January 1996 to December 2000 were 0.079, 0.276, 0.295 and 0.305, respectively.

In summary, Sharp ratios show a positive reward and, often, a high reward relative to risk for static buy-and-hold trading strategies for low-priced contracts across intermediate to long-term horizons.

B. Dynamic Trading Strategies

Suppose a hypothetical investor knew how to forecast payoff probabilities according to the logit model estimates before the markets started. How would such an investor trade and how much would she earn if she could trade each night at closing market prices conditional on these payoff probability assessments? This calculation provides a measure of the economic significance of the bias documented here. To determine this, consider an investor with a mean-variance utility function given by $U(x) = E(x) - \gamma \text{Var}(x)$, where x is the level of payoffs at the date of liquidation and γ measures the investor's risk aversion. Suppose that the investor maximized her utility at each date from 21 days to one day before liquidation value determination given current prices and forecast payoff probabilities. Given an initial budget of, say \$100, it is easy to determine the holdings and return for such an investor at any given point in time given the constraints that cash and contract holdings may never fall below zero and contract holdings must be in integer values. Because this strategy requires the investor to know more than she possibly could have known at the time of trading and it allows the trader to trade any desired quantities at closing market prices without adverse price effects, it could not actually have been implemented.. Nevertheless, it provides a measure of the potential for profitability of dynamic trading strategies in these markets.

Table 5 summarizes the holdings and returns for such an investor in the Microsoft Price Level Market. Using Model I, the trader would have generated returns that average 0.31% to 1.2% a month for risk aversion parameters ranging from 0.5 to 0.1. These returns are generated with surprisingly low volumes of 0.21 to 1.22 contracts traded per day on average and a maximum of 10 contracts traded. Over all of the markets, the trader's portfolio would have increased in value from an initial \$100 to final values

in the \$120 to \$207 range depending on risk aversion. If the trader used the current price of Microsoft stock relative to the cutoff as extra-market information in Model III, returns would increase dramatically. For the most part, volumes remain small, ranging from 1 to 5 contracts daily on average. However, on one trading day, the trader would have wanted volumes ranging from 366 to 1,826 contracts depending on the risk aversion parameter. Much of the return to this strategy, which ranges from 4.76% to 11.32% per month on average and gives final portfolio values ranging from \$411 to \$1,655, result from this day's trading. If the trader were not able to trade on this day, returns still would have averaged 0.80%, 1.31% and 2.31% per month for risk aversion parameters of 0.5, 0.25 and 0.1, respectively. Final portfolio values still would have ranged up to \$404.

Table 6 summarizes the holdings and returns for such an investor in the Computer Industry Returns Market. Using Model I, the trader would have generated returns that average 0.16% to 0.84% a month for risk aversion parameters ranging from 0.5 to 0.1. These, somewhat lower, returns require higher volumes than the Microsoft Price Level Market strategy (averaging 0.94 to 4.77 contracts per day). Over all the markets, the trader's initial \$100 portfolio would have increased in value to \$109 to \$166 depending on risk aversion. If the trader used current relative returns as extra-market information in Model III, returns would have increased dramatically as in the Microsoft Price Level Market. The trader would have frequently hit the budget constraint and volumes would often be high. Hypothetical returns range from 2.09% to 3.96% per month on average and result in final portfolio values ranging from \$430 to \$1,299. To attain these returns, the trader would have had to trade from 24.64 to 73.52 contracts per day on average. Because the trader would frequently hit short sale and budget constraints, returns could have been even higher with a larger initial endowment. In contrast to the Microsoft Price Level Market, high volumes and returns occurred on numerous days.

Overall, the overconfidence bias in the market had the potential for generating considerable excess returns for traders using dynamic trading strategies designed to exploit these biases.

IV. Summary and Discussion

The longshot bias is well documented for racetrack betting markets (see Thaler and Ziemba, 1988). In other markets, it may be mitigated (e.g., Gray and Gray, 1997, for football betting) or even reversed (e.g., Woodland and Woodland, 1994, for baseball betting). These differences may arise because of differences in context or market structure. The overconfidence bias has been documented for individual assessments of probability and choices (Lichtenstein, Fischhoff and Phillips, 1982) and in overreactions to private information (Daniel, Hirshleifer and Subrahmanyam, 1998). Because such biases could affect significantly financial markets and because financial markets differ both in context and structure from betting markets and individual choice tasks, it is important to determine whether such biases carry over to and actually do affect financial markets.

We ask whether the longshot bias or the overconfidence bias appears to generalize to and affect prices in the Iowa Electronic Markets (IEM). These are active, two-sided real-money markets in a financial context. IEM traders have the ability to synthetically create short positions. There are no explicit transactions costs. These features may be important in allowing traders to drive out mispricing caused by biases. In fact, the longshot bias appears in reverse in a series of IEM markets. Prices seem biased by overconfidence on the part of traders in their ability to forecast what will happen in these markets several days to weeks in advance. Inefficiencies arising from this bias could result in considerable profits for astute traders.

Can the existence of the longshot bias in some betting markets, its absence in others and reversal in the IEM be explained in a consistent manner? Perhaps. Most of the explanations for the existence of the longshot bias depend on a combination of bettor biases and structural characteristics of the betting markets. For example, Thaler and Ziemba (1988) suggest a possible explanation as follows: “Some bettors may choose horses for essentially irrational reasons, like the horse’s name. Since there is no possibility of short sales, such bettors can drive the odds down on the worst horses, with the “smart money” simply taking the better bets on the favorites.” Others have developed this idea more formally.²⁵ Most of the models use “irrational reasons” which each could easily be interpreted as an overconfident

reaction to some information, or even just a hunch, that a particular horse might win. Each explanation relies on some additional feature of the betting market that creates an asymmetric effect to make its case (e.g., no shorting, the pari-mutuel pricing mechanism, etc.). Exceptions to the longshot bias occur when these other features change. For example, Woodland and Woodland (1994) speculate that the reason that the bias reverses for baseball betting is that betters can bet against teams. Gray and Gray (1997) cite the fact that, in football betting, the point spreads (instead of odds) adjust to keep the bookie's take constant. Hurley and McDonough (2000) cite the track "take."

This leaves the debate over whether biases will extend to financial markets completely unsettled. For the IEM and other financial markets, two major structural factors change and these may mitigate longshot and overconfidence bias effects. First, there is no "take" for the track (set in response to an asymmetric betting structure) in either the IEM or naturally occurring markets. Instead, there is a bid/ask spread, presumably set by market makers who face symmetric adverse selection issues. Second, often traders in naturally occurring markets can short sell. While traders cannot directly short-sell on the IEM, they can set up synthetic shorts and these positions are feasible because of the absence of explicit transactions costs. Thus, traders are free to sell against prices that they believe are too high. The differences in these features may remove the asymmetric effects of overconfidence. This would leave only an overall tendency for overconfidence on the parts of traders that would drive up prices of contracts that traders are relatively "sure" will pay off and drive down prices on the contracts that traders are relatively "sure" will not pay off.²⁶ The evidence from the IEM suggests the overconfidence bias is the one likely to appear in two-sided financial markets.

Finally, it is interesting to note that the resulting price impact of the overconfidence bias is transitory. As Daniel, Hirshleifer and Subrahmanyam (1998) predict, it disappears and the markets become efficient as the liquidation date for contracts approaches. Nevertheless, traders who were aware of the biases could have reaped substantial profits by exploiting the resulting mispricing. Mirroring the typical pattern of open interest in futures markets, mean-variance utility maximizing trading strategies generally purchase or sell mispriced contracts at intermediate horizons and then close out these positions

as the liquidation date approaches. This is the strategy one would expect arbitrageurs to follow when exploiting mispricing. Traders acting in this manner could be the reason for prices converging to efficient levels near liquidation. Thus, trading dynamics, not just the market structure, of the IEM parallels naturally occurring markets. This instills even greater confidence in the external validity of the IEM and its usefulness for studying factors driving convergence, price dynamics and efficiency in naturally occurring markets.

Footnotes

- ¹Snyder (1978) cites eight early studies. See Thaler and Ziemba (1988) for a more recent and complete list of papers.
- ² There are other effects of overconfidence that we do not investigate. As examples, traders may trade more than they should (Barber and Odean, 2001), they may think they know more than they actually do (Russo and Schoemaker, 1992), or they may over- or under-react information based on its type (Odean, 1998).
- ³ The false consensus effect is the belief that ones own views are more representative of the general population than they are in reality (see Ross, Greene and House, 1977). The assimilation-contrast effect occurs when people interpret news more favorably with respect to maintained positions than warranted (see Parducci and Marshall, 1962). The IEM political markets in the referenced research were vote-share markets where contract prices represented the associated candidate's share of the popular vote (not probabilities of outcomes).
- ⁴ Many of the following points are also used to justify the racetrack literature. See Thaler and Ziemba (1988).
- ⁵ Traders can obtain a short position by buying a complete set of contracts from the market for its fixed known payoff of \$1 and selling the contract(s) they want to short. This creates the same cash flows and net positions as a short, but protects the IEM exchange from losses. It constitutes a short that is cash covered to the worst outcome.
- ⁶ In typical experimental markets, traders are "staked" by the experimenter and therefore cannot lose their own money. This differs in the IEM where traders must invest their own money. Because IEM contracts are real futures contracts, the IEM is under the regulatory purview of the Commodity Futures Trading Commission (CFTC). The CFTC has issued a "no-action" letter to the IEM stating that as long as the IEM conforms to certain restrictions (related to limiting risk and conflict of interest), the CFTC will take no action against it. Under this no-action letter, IEM does not file reports that are required by regulation and therefore it is not formally regulated by, nor are its operators registered with, the CFTC.
- ⁷ These are the two series of IEM markets that trade only "binary" (\$1 or \$0 payoff) contracts and that are repeated under essentially identical conditions monthly. We do not use IEM contracts with other types of payoffs because we cannot interpret prices as probabilities for other types of contracts. We focus only on repeated markets to insure the stationarity needed for econometric analysis across markets.

⁸These dates were chosen because the contract values are linked to underlying values of stocks on option expiration dates (the third Fridays of each month) as discussed later.

⁹Cutoffs are chosen to be the strike price of the closest-to-the-money traded option for the stock (i.e., the closest \$5 increment to the current stock price). This insures that, at least at the outset, both contracts have intermediate values. Contracts can be split (see Appendix I) when stock prices deviate significantly from the cutoff and contract prices reach extreme levels (close to 0 or 1). Only one split has occurred in the data set used here. This split created a “middle” range contract in addition to the “L” contract (that pays off if Microsoft closes below a cutoff) and the “H” contract (that pays off if Microsoft closes above a cutoff). For consistency, only the “L” and “H” contracts are used in the data analysis here. Making a different choice or omitting data from this month does not change any of our results.

¹⁰We use dividend-adjusted returns for the stocks and the capital gains returns for the index. If two or more securities tie, the payoffs are evenly split among the tied contracts. This has never happened.

¹¹See Malinvaud (1974) for a general equilibrium proof of this proposition. It also results from CAPM or APT along the following lines of argument: According to each theory, the equilibrium return of any security (including our contracts) should equal the risk-free rate plus a risk premium associated with each aggregate risk factor. Since there is zero aggregate risk, the risk premiums will be zero. Since the risk-free rate is also zero, the expected return for each contract will be zero. This will only be true if, at every point in time, the price of each contract is equal to its expected future value. Technically, these will be risk-neutral probabilities and hedging demand may drive them away from true probabilities. However, the markets are restricted to students and set the maximum investment amount at \$500 per trader. This should minimize any such effects. Further, work from the political markets on the IEM suggests that traders do not hedge against their own political preferences (see Forsythe, Ross and Rietz, 1999).

¹²Some readers may have difficulty with a zero risk-free rate. One might speculate that a positive risk free rate

would result in $p_t^j = \frac{q_t^j}{1 + r_t}$, where r is the (positive) risk-free rate for t days. (Since there is still no aggregate

risk, there will still be no risk premium.) However, this would violate arbitrage restrictions since this implies

that $\sum_{j=1}^J p_t^j < \sum_{j=1}^J q_t^j = 1$. If this were so, traders could buy the portfolio from other traders at a combined price

of less than \$1 and sell it back to the exchange for \$1, making a sure profit. This activity would drive the discount rate on each contract to zero. Even if traders ignored this arbitrage opportunity, the discount rates would fall out of the analysis discussed later because normalized prices and price ratios, not raw prices, are used.

¹³ We use standard logit models for the two-contract Microsoft Price Level Market and four-state multinomial logit models for the four-contract Computer Industry Returns Markets.

¹⁴ The selection of the particular base contract is irrelevant. It has no effect on the estimates or predictions.

¹⁵ We will actually use “closing normalized” contract prices. Closing prices are the last trade before midnight because the IEM operates 24 hours a day. To control for non-synchronous trading, normalized contract prices are computed as the closing price of a contract divided by the sum of closing prices for all contracts in a complete slate. This insures that prices can be interpreted as probabilities because they will sum to one.

¹⁶ To see how the graph works, consider first the $\beta > 1$, two-contract case for small and large probabilities of \$1 payoffs. If $p < 1-p$, then $p/(1-p) < 1$ and $\ln(p/(1-p)) < 0$. Taking logs of both sides of the logit relationship, $\ln(q/(1-q)) = \beta \ln(p/(1-p)) < \ln(p/(1-p)) \Rightarrow q/(1-q) < p/(1-p) \Rightarrow q < p$. Similarly, if $p > 1-p$, then $q > p$. The opposite relationships hold for $\beta < 1$. The functions “wrap back” to 0 and 1 because of the nature of log price ratios and the “crossover” point is affected by α (crossing at 0.5 if α is zero).

¹⁷ One can easily show that using the other contract as the base contract in the two contract case results in no change in the estimated price ratio coefficient. Consider the logit specification:

$$\frac{q}{1-q} = e^{\alpha + \beta \ln\left(\frac{p}{1-p}\right)} \Rightarrow \ln q - \ln(1-q) = \alpha + \beta \ln p - \beta \ln(1-p)$$

$$\Rightarrow \ln(1-q) - \ln q = -\alpha + \beta \ln(1-p) - \beta \ln p \Rightarrow \frac{1-q}{q} = e^{-\alpha + \beta \ln\left(\frac{1-p}{p}\right)}$$

Thus, the price ratio coefficient is invariant to changing the base contract. Typically, the multinomial case mirrors this result. If the coefficients on all price ratios relative to the base contract equal 1, there is no bias. Generally, ratios greater than one give price to probability maps that reflect a longshot bias. However, for some cases, a longshot bias can arise even though not all of the coefficients exceed one and vice versa for the overconfidence bias. In the analysis below, we will rely on joint tests that all coefficients equal one and direct inspections of the price to probability maps to determine what kind of bias exists.

¹⁸This effect could also result from overreaction. For instance, when the rational forecast of the payoff probability from the current information state is actually 90%, this is “good news” for the contract. If traders overreact to this information in assessing the probability of payoff, they may be willing to pay, say, 95 cents for this contract. We do not attempt to distinguish between these two possible causes for this effect and simply refer to it as an overconfidence effect or reverse longshot bias.

¹⁹This remains consistent with overconfidence in probability assessments documented in Lichtenstein, Fischhoff and Phillips (1982) but adds a dynamic element. We observe a dynamic consistent with this. Shortly after each monthly IEM market opens, prices are noisy and relatively unbiased. As information about the underlying security prices comes in over the month, a price bias develops. As liquidation approaches and uncertainty is resolved, the price bias is mitigated.

²⁰While the markets run either 28 or 35 days, the liquidation values are determined with three days left in the market. We will calculate the horizon from this date (the third Friday of each month). Thus, markets run with either a 25- or a 32-day initial horizon. However, volume in the initial days is often low or absent, spreads are high and the amount of data for analysis drops off significantly.

²¹They use simple logit models to predict the probability of winning for all horses without considering the multinomial nature of each individual race.

²²This is information that the traders should have had before closing prices are determined because the closing prices in the IEM are measured at midnight.

²³The selection of the particular base contract is irrelevant. It has no effect on the estimates or predictions.

²⁴Again, this is information that the traders should have had before closing prices are determined because the closing prices in the IEM are measured at midnight.

²⁵Hurley and McDonough (1995) develop this idea with a simple model in which “uninformed” bettors place bets on their own beliefs while “informed” bettors place optimal bets in response. Because of the track take and costly information, this could drive the bias. Hurley and McDonough’s (2000) later model of the bias rests on the assumption that each bettor receives a subjective probability signal and then believes that “only he/she is really capable of processing the available information in the correct way.” They do not update based on observed odds. This, combined with features of the pari-mutuel mechanism (no shorts and final odds that are unknown at the time of the bet), produces a bias.

²⁶ Whether this is the actual reason will be the subject of future research directly contrasting laboratory markets in which prices are set according to an asymmetric parimutuel mechanism and those in which prices are set by a full-fledged market mechanism.

References

- Asch, P, and RE Quandt, 1986, Racetrack Betting: The Professors' Guide to Strategies, Dover, MA: Auburn House.
- Barber, BM and T Odean, 2001, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," Quarterly Journal of Economics, 261-292.
- Berg, JE and TR Rietz, 2006, "The Iowa Electronic Markets: Stylized Facts and Open Issues," in Information Markets: A New Way of Making Decisions, edited by R. Hahn and P. Tetlock, AEI/Brookings Center for Regulatory Studies, 142-169.
- Cosmides, L, 1989, "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Task," Cognition, 31, 187-276
- Daniel, K, D Hirshleifer and A Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Overreactions," The Journal of Finance, 58(6): 1839-1885.
- Forsythe, R, FD Nelson, GR Neumann and J Wright, 1992, "Anatomy of an Experimental Political Stock Market," American Economic Review, 82, 1142-1161.
- Forsythe, R, TW Ross and TA Rietz, 1999, "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets," Journal of Economic Behavior and Organization, 39, 83-110.
- Gandar, J, R Zuber, T O'Brien and B Russo, 1988, "Testing Rationality in the Point Spread Betting Market," Journal of Finance, 1988, 995-1008.
- Gode, DK, and S Sunder, 1993, "Allocative Efficiency of Markets with Zero Intelligence (ZI) Traders: Market as a Partial Substitute for Individual Rationality," Journal of Political Economy, 101, 119-137.
- Gray, PK, and SF Gray, 1997, "Testing Market Efficiency: Evidence from the NFL Sports Betting Market," Journal of Finance, 52, 1725-1737.
- Grossman, SJ, and JE Stiglitz, 1980, "On the Impossibility of Informationally Efficient Markets," The American Economic Review, 70, 393-408.

- Hurley, W, and L McDonough, 1995, "A Note on the Hayek Hypothesis and the Favorite-Longshot Bias in Parimutuel Betting," The American Economic Review, 85, 949-955.
- Hurley, W, and L McDonough, 2000, "Boundedly Rational Agents and the Favourite-Longshot bias in Parimutuel Betting," Working Paper, Department of Economics, Queens University, Kingston, Ontario, Canada.
- Jegadeesh, N, and S Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,": Journal of Finance 48, 65–91.
- Kahneman, D, and A Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk," Econometrica, 47, 263-91.
- Lichtenstein, S, B Fischhoff and LD Phillips, 1982, "Calibration of Probabilities: The State of the Art to 1980," in Kahneman, D, P Slovic and A Tversky (Eds.), Judgment Under Uncertainty: Heuristics and Biases, Cambridge: Cambridge University Press, 306-334.
- Malinvaud, E, 1974, "The Allocation of Individual Risks in large Markets," in J. H. Dréze, Ed., Allocation under Uncertainty: Equilibrium and Optimality, London: MacMillan Press, 110-125.
- Parducci, A, and L Marshall, 1962, "Assimilation vs. Contrast in the Anchoring of Perceptual Judgments of Weight," Journal of Experimental Psychology, 63, 426-437.
- Odean, T, 1998, "Volume, Volatility, Price, and Profit When All Traders Are Above Average," Journal of Finance, 53(6), 1887-1934.
- Oliven, K, and TA Rietz, 2004, "Suckers are Born but Markets are Made: Individual Rationality, Arbitrage and Market Efficiency on an Electronic Futures Market," Management Science, 50, 336-351.
- Reny, PJ, and M Perry, 2006, "Toward a Strategic Foundation for Rational Expectations Equilibrium," Econometrica, 74(5): 1231–1269.
- Ross, L, D Greene and P House, 1977, "The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes," Journal of Experimental Psychology, 13, 279-301.

- Rubinstein, M, 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23,1976 Through August 31,1978," Journal of Finance, 40, 455-480.
- Russo, JE, and PH Schoemaker, 1992, "Managing Overconfidence," Sloan Management Review, 33(2), 7-17.
- Shin, HS, 1992, "Prices of State Contingent Claims with Insider Trades, and the Favorite-Longshot Bias," The Economic Journal, 102, 426-435.
- Slovic, P, 1972, "Psychological Study of Human Judgment: Implications for Investment Decision Making," Journal of Finance, 27, 779-99.
- Snyder, WW, 1978, "Horse Racing: Testing the Efficient Markets Model," Journal of Finance, 33, 1109-1118.
- Thaler, RH, and WT Ziemba, 1988, "Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries," The Journal of Economic Perspectives, 2, 161-174.
- Woodland, LM, and BM Woodland, 1994, "Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market," The Journal of Finance, 49, 269-279.
- Ziemba, WT, and DB Hausch, 1986, Betting at the Racetrack, Vancouver and Los Angeles: Dr. Z Investments, Inc.

Appendix I: Microsoft Price Level Market Prospectus

At noon (central time), Monday, January 18, 1999, the Iowa Electronic Markets (IEM) opened trade in a series of contracts based on price levels of securities in the computer industry. This document describes these contracts. Except as specified in this prospectus, trading rules for the computer these contracts are the same as those specified in the Trader's Manual for the Iowa Electronic Markets.

CONTRACTS

Each month, a new set of winner-takes-all contracts will be offered in this market. Contract liquidation values will be determined by closing stock price levels on the third Friday of the month after contracts are created (see note 1 below).

The liquidation values for these contracts are determined solely by closing prices of Microsoft Corp. Common Stock (MSFT). Each month, an initial pair of contracts will consist of "MSxxxmH" and "MSxxxmL," where "xxx" corresponds to a "cutoff" price of \$xxx and "m" corresponds to the liquidation month as given in the following table:

<u>Month</u>	<u>Code</u>	<u>Month</u>	<u>Code</u>	<u>Month</u>	<u>Code</u>
January	a	May	e	September	i
February	b	June	f	October	j
March	c	July	g	November	k
April	d	August	h	December	l

The payoff for the "H" contract will equal \$1.00 if the Wall Street Journal closing price for Microsoft Common Stock on the third Friday month "m" exceeds \$xxx. It will equal \$0.00 otherwise. The payoff for the "L" contract will equal \$1.00 if the Wall Street Journal closing price for Microsoft Common Stock on the third Friday month "m" is less than or equal to \$xxx. It will equal \$0.00 otherwise.

We will choose \$xxx to correspond to the strike price of the exchange traded option that lies closest to the price of Microsoft Common Stock on the date we create the contracts.

Thus, the initial contracts are:

<u>Contract</u>	<u>Underlying Fundamental</u>	<u>Liquidation Value</u>
MsxxxmH	Microsoft Common Stock	\$1.00 if MSFT closing price > \$xxx
MsxxxmL	Microsoft Common Stock	\$1.00 if MSFT closing price ≤ \$xxx

CONTRACT SPLITS

If the trading price of a particular contract becomes unusually high, the Directors of the IEM may authorize a contract split. The decision to split a contract will be announced at least two days in advance of the split, and the new contract names and the timing of the split will be included in the announcement. This announcement will appear as a News Bulletin on your screen.

When a split occurs, the original contract will be split into two contracts. If the MSxxxmH contract is split, all traders holding an MSxxxmH contract will receive in its place two “new” contracts: An MSxxx-yyym contract and an MSyyymH contract where yyy is a new, higher cutoff price level. After the split, MSxxx-yyym contracts will pay \$1.00 if the MSFT closing price on the third Friday of the liquidation month is higher than \$xxx and lower than or equal to \$yyy. MSyyymH contracts will pay \$1.00 if the MSFT closing price on the third Friday of the liquidation month is higher than \$yyy. Thus, splits determine mutually exclusive ranges of prices over which each contract pays. Since the value of the two new contracts differ, outstanding bids and asks for MSxxxmH will be canceled at the time of the split. Since the payoffs to MSxxxmL are unaffected by the split, bids and offers for this contract will remain.

If the MSxxxmL contract is split, all traders holding an MSxxxmL contract will receive in its place two “new” contracts: MSzzz-xxxm contract and a MSzzzmL contract where zzz is a new, lower cutoff price level. Similar splits of any other contracts may also occur. All other aspects of these splits and the payoffs from the resulting contracts are analogous to those described above. Again, splits determine mutually exclusive ranges of prices over which each contract pays.

NOTE: On April 27, 2000 the naming convention was updated to make the meaning of contract names clearer after splits. All other aspects of splits remain unchanged.

CONTRACT LIQUIDATION

Existing contracts will be liquidated by the IEM on the Monday after the third Friday of each month (see note 1). The Midwest Edition of the Wall Street Journal will be the official source of closing prices.

If Microsoft stock is de-listed, the last available closing price will be used as the closing price for determining liquidation values.

If Microsoft stock undergoes a stock split during the trading period, the closing price of its stock used to calculate payoffs will be adjusted to take account of this split. Specifically if each existing share is split into M

shares, then the closing price used to calculate payoffs will be multiplied by M since this represents the value of one pre-split share in the company. Stock dividends will be treated in the same manner.

LISTING NEW CONTRACTS

New contracts will be created by the IEM on the Monday after the third Friday of each month (see note 1 below). Contracts may be moved across and within market display windows to facilitate access. However, once trading commences in any contract, it will remain listed until the liquidation value is determined.

UNIT PORTFOLIOS

For each month's contracts, unit portfolios consisting of bundles of contracts whose payoff is guaranteed to be \$1.00 and can be purchased from or sold to the IEM system at any time. The price of each unit portfolio is \$1.00. To buy and sell bundles, select the appropriate bundle from the "Market Order" drop down menu on the market trading screen. Unit portfolio bundle names are Msft_1\$m for month "m" liquidation.

ACCESS

Current and newly enrolled IEM traders with academic affiliation will automatically be given access rights to the MSFT (Microsoft) Price Level Market. Access to the contracts is achieved via the "Market Selection" pull down menu. Funds in a trader's cash account are fungible across all contracts so new investment deposits are not required. Additional investments up to the maximum of \$500 can be made at any time. With five days' advance notice, funds may be withdrawn on the 15th of any month.

Note 1: Generally, exchange traded options for the underlying stocks expire on the Saturday following the third Friday of each month. In the event that the options' expiration dates change for any reason, we will change the dates used to determine contract creations, liquidations, returns and payoffs accordingly.

Appendix II: Computer Industry Returns Market Prospectus

At noon (central time), Monday, August 28, 1995, the Iowa Electronic Markets (IEM) will open trade in a series of contracts based on the returns of securities in the computer industry. This document describes these contracts. Except as specified in this prospectus, trading rules for these contracts are the same as those specified in the Trader's Manual for the Iowa Electronic Markets.

CONTRACTS

Each month, a new set of winner-takes-all contracts will be offered in this market. Contract liquidation values will be determined by rates of return measured from the third Friday of one month to the third Friday of the next month (see note 1 below).

The liquidation values for these contracts are determined solely by the dividend adjusted rates of return of Apple Computer, Inc. Common Stock (AAPL, listed on NASDAQ), International Business Machines Corporation Common Stock (IBM, listed on the NYSE) and Microsoft Corporation Common Stock (MSFT, listed on NASDAQ); and the capital gains rate of return on the Standard and Poor's 500 Index. Whichever of these has the highest rate of return as specified below will pay off \$1.00 per contract. All other contracts will pay off zero (see note 2 below).

Contracts will be designated by a ticker symbol and a letter denoting the month of contract liquidation.

Thus, the contracts traded in this market for liquidation in month "m" are:

<u>Code</u>	<u>Contract Description</u>	<u>Liquidation Value</u>
AAPLm	Apple Computer	\$1.00 if AAPL return is highest
IBMm	IBM	\$1.00 if IBM return is highest
MSFTm	Microsoft	\$1.00 if MSFT return is highest
SP500m	S&P 500 Market Index	\$1.00 if SP500 return is highest

The month code, "m," refers to the month of liquidation as given by the following table:

<u>Month</u>	<u>Code</u>	<u>Month</u>	<u>Code</u>	<u>Month</u>	<u>Code</u>
January	a	May	e	September	i
February	b	June	f	October	j
March	c	July	g	November	k
April	d	August	h	December	l

COMPUTING RETURNS

For AAPLm, IBMm and MSFTm, we will compute the dividend adjusted rate of return based on closing stock prices of the underlying listed firm between the third Friday in the liquidation month and the third Friday in the previous month. For these purposes, we will use closing prices as reported in the Midwest edition of the Wall Street Journal.

The Dividend Adjusted Rate of Return is calculated as follows: First, we compute the raw return on the underlying stock (the closing price on the third Friday of the liquidation month, minus the closing price from the third Friday of the previous month, plus any dividends on ex-dividend dates). Then, we divide the raw return by the closing stock price from the previous month to arrive at the dividend-adjusted rate of return.

For the SP500 contract, we compute the capital gains rate of return by subtracting the closing index value on the third Friday of the previous month from the closing index value on the third Friday of the liquidation month and then divide by the previous month's closing index value.

CONTRACT LIQUIDATION

Existing contracts will be liquidated by the IEM on the Monday after the third Friday of each month (see note 1 below). The Midwest Edition of the Wall Street Journal will be the official source of closing prices.

If one of the companies is de-listed, the last available closing price will be used as the closing price for determining liquidation values.

If one of the companies undergoes a stock split during the trading period, the closing price of its stock used to calculate payoffs will be adjusted to take account of this split. Specifically if each existing share is split into M shares, then the closing price used to calculate payoffs will be multiplied by M since this represents the value of one pre-split share in the company. Stock dividends will be treated in the same manner.

LISTING NEW CONTRACTS

New contracts will be created by the IEM on the Monday after the third Friday of each month (see note 1 below).

Contracts may be moved across and within market display windows to facilitate access. However, once trading commences in any contract, it will remain listed until the liquidation value is determined.

UNIT PORTFOLIOS

For each month's contracts, unit portfolios consisting of bundles of contracts whose payoff is guaranteed to be \$1.00 and can be purchased from or sold to the IEM system at any time. The price of each unit portfolio is \$1.00. To buy and sell bundles, select the appropriate bundle from the "Market Order" drop down menu on the market trading screen. Unit portfolio bundle names are Comp_1\$m for month "m" liquidation.

ACCESS

Current and newly enrolled IEM traders with academic affiliations will automatically be given access rights to the Computer Industry Returns Market. Access to the contracts is achieved via the "Market Selection" pull down menu. Funds in a trader's cash account are fungible across all contracts so new investment deposits are not required. Additional investments up to the maximum of \$500 can be made at any time. With five days' advance notice, funds may be withdrawn on the 15th of any month.

Note 1: Generally, exchange traded options for the underlying stocks expire on the Saturday following the third Friday of each month. In the event that the options' expiration dates change for any reason, we will change the dates used to determine contract creations, liquidations, returns and payoffs accordingly.

Note 2: If two or more contracts tie for the highest return, the \$1.00 will be divided as evenly as possible among the tied contracts with any residual \$0.001's allocated in order of the highest to lowest final values.

Figures




Figure 1: Mapping between IEM contract payoff probabilities (q) and IEM contract prices (p) under a longshot bias ($\beta > 1$), unbiased pricing ($\beta = 1$) and overconfidence ($\beta < 1$). Betas correspond to own log price ratio coefficients in the logit regressions designed to predict the actual payoff probability for a contract from its log price ratio.

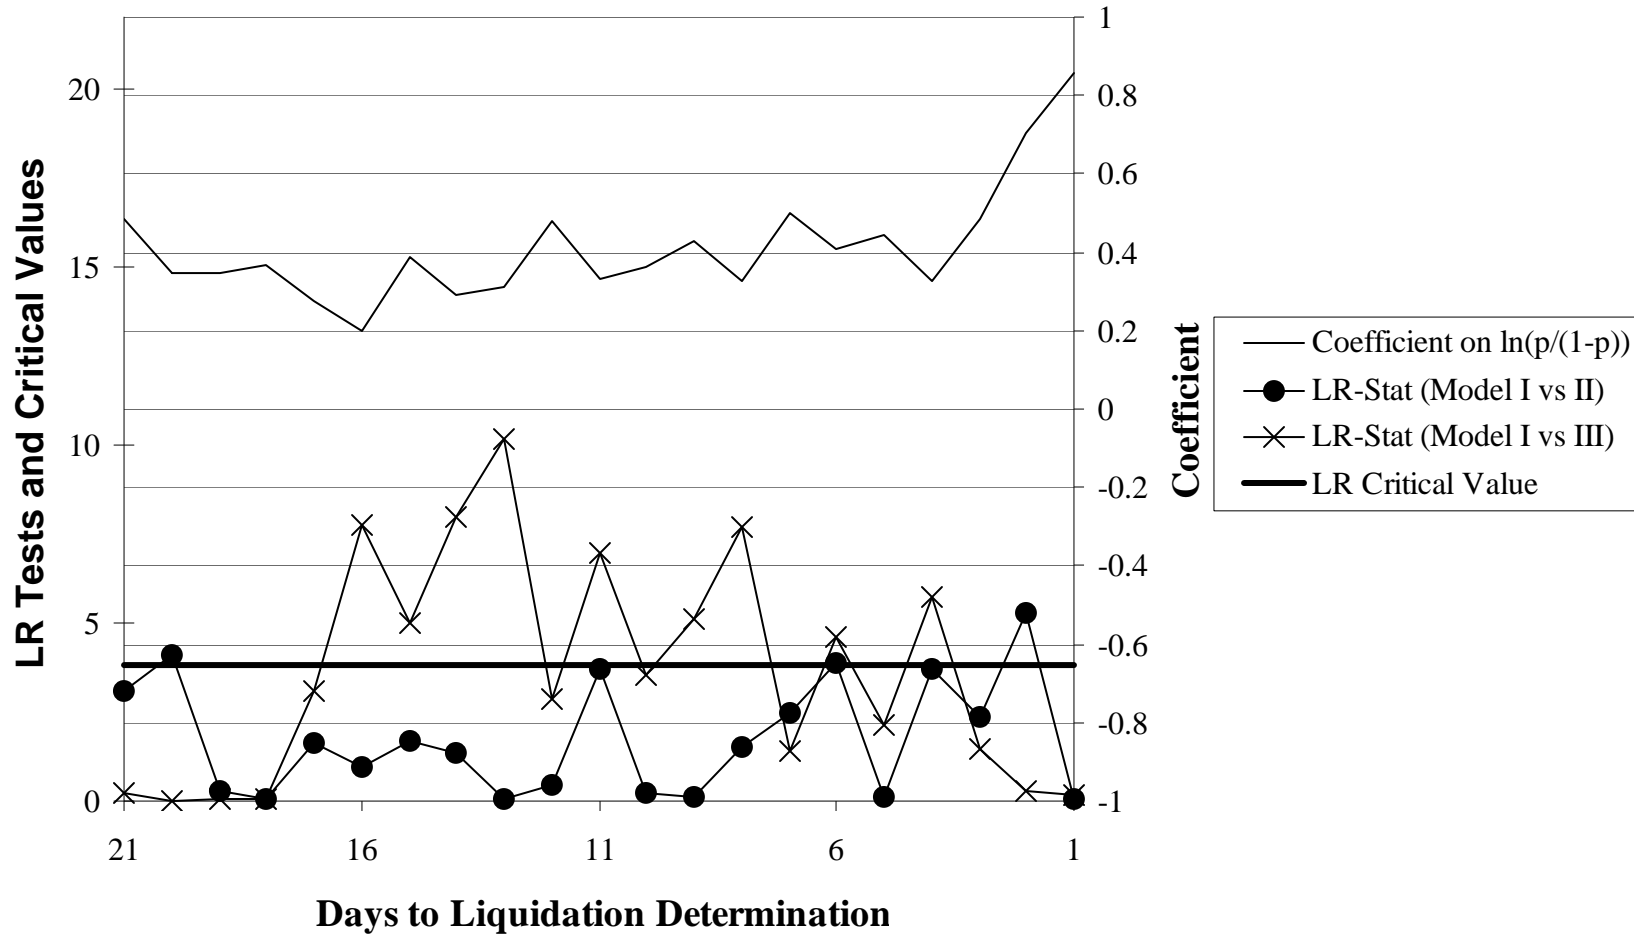


Figure 2: Microsoft Price Level Market Logit Model I coefficients and Likelihood Ratio Test Statistics to test differences between Model I and Model II and between Model I and Model III. The dependent variable in all models is the likelihood of a contract payoff. All models use the contract's log price ratio as an independent variable. Model II adds the one-day change in the log price ratio to the log price ratio. Model III adds the current ratio of Microsoft's stock price to the payoff-determining cutoff to the log price ratio.

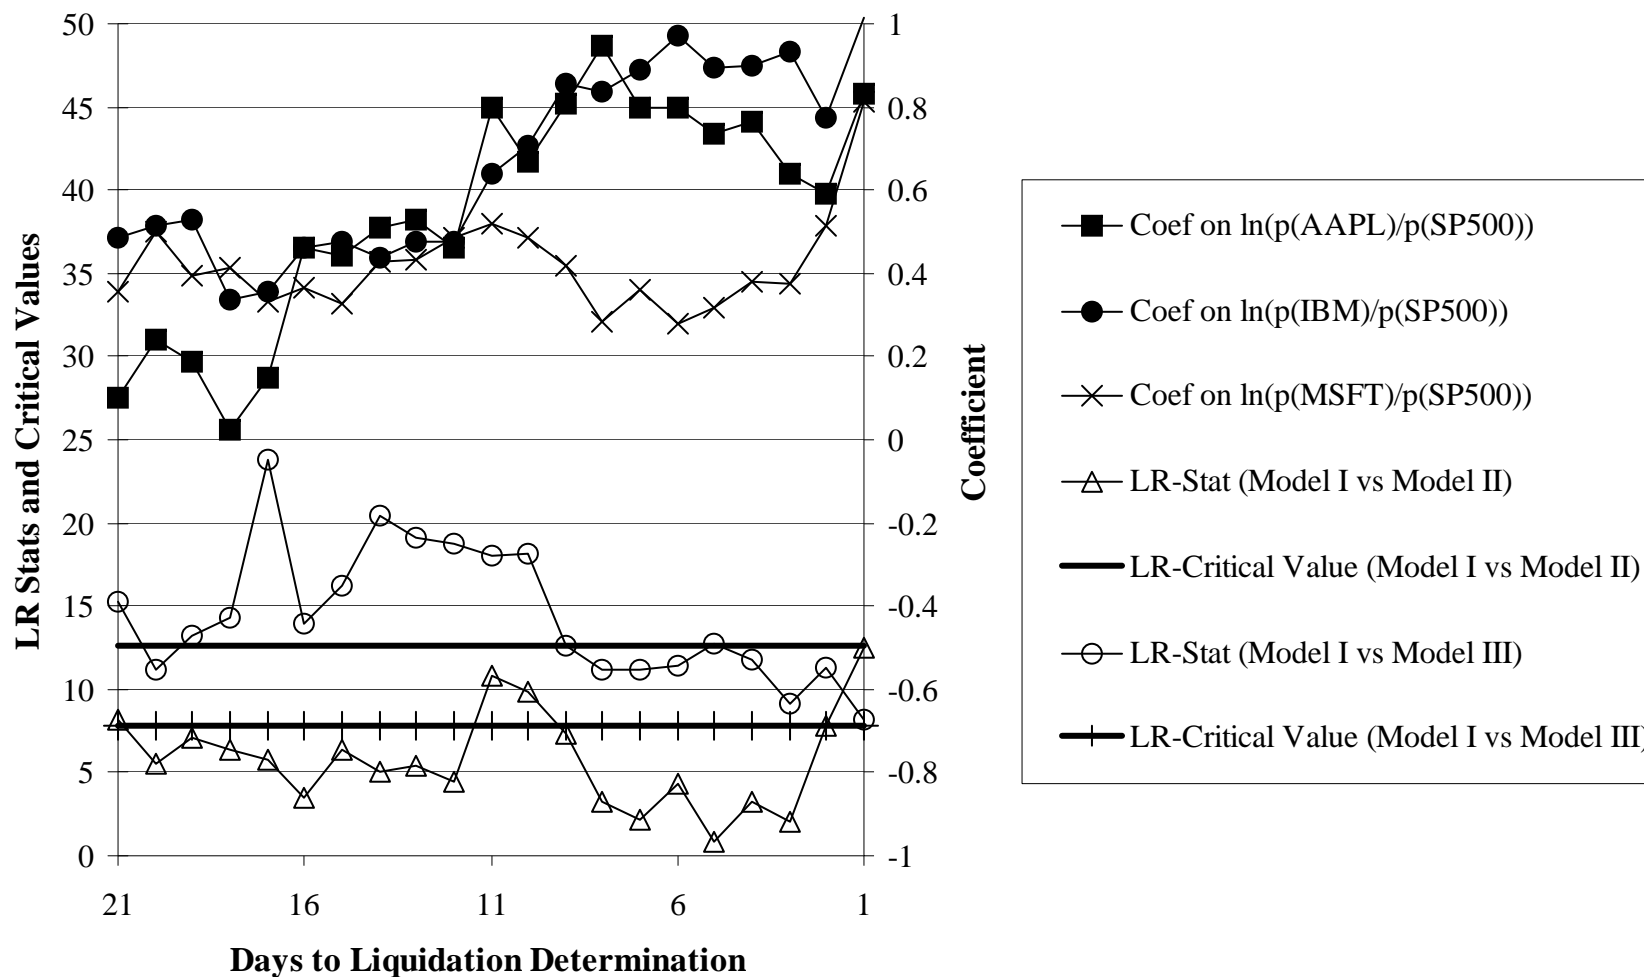


Figure 3: Computer Industry Returns Market Logit Model I coefficients and Likelihood Ratio Test Statistics to test differences between Model I and Model II and between Model I and Model III. The dependent variable in all models is the likelihood of a contract payoff. All models use the contracts' own log price ratios to the S&P500 contract as independent variables. Model I restricts cross log price ratio coefficients to zero. Model II allows cross log price ratio coefficients different from zero. Model III adds the current return lead/lag for the contract (return to date on the contract's underlying stock minus the maximum return on the other three stocks) to Model I while restricting cross log price ratio and lead/lag coefficients to zero.

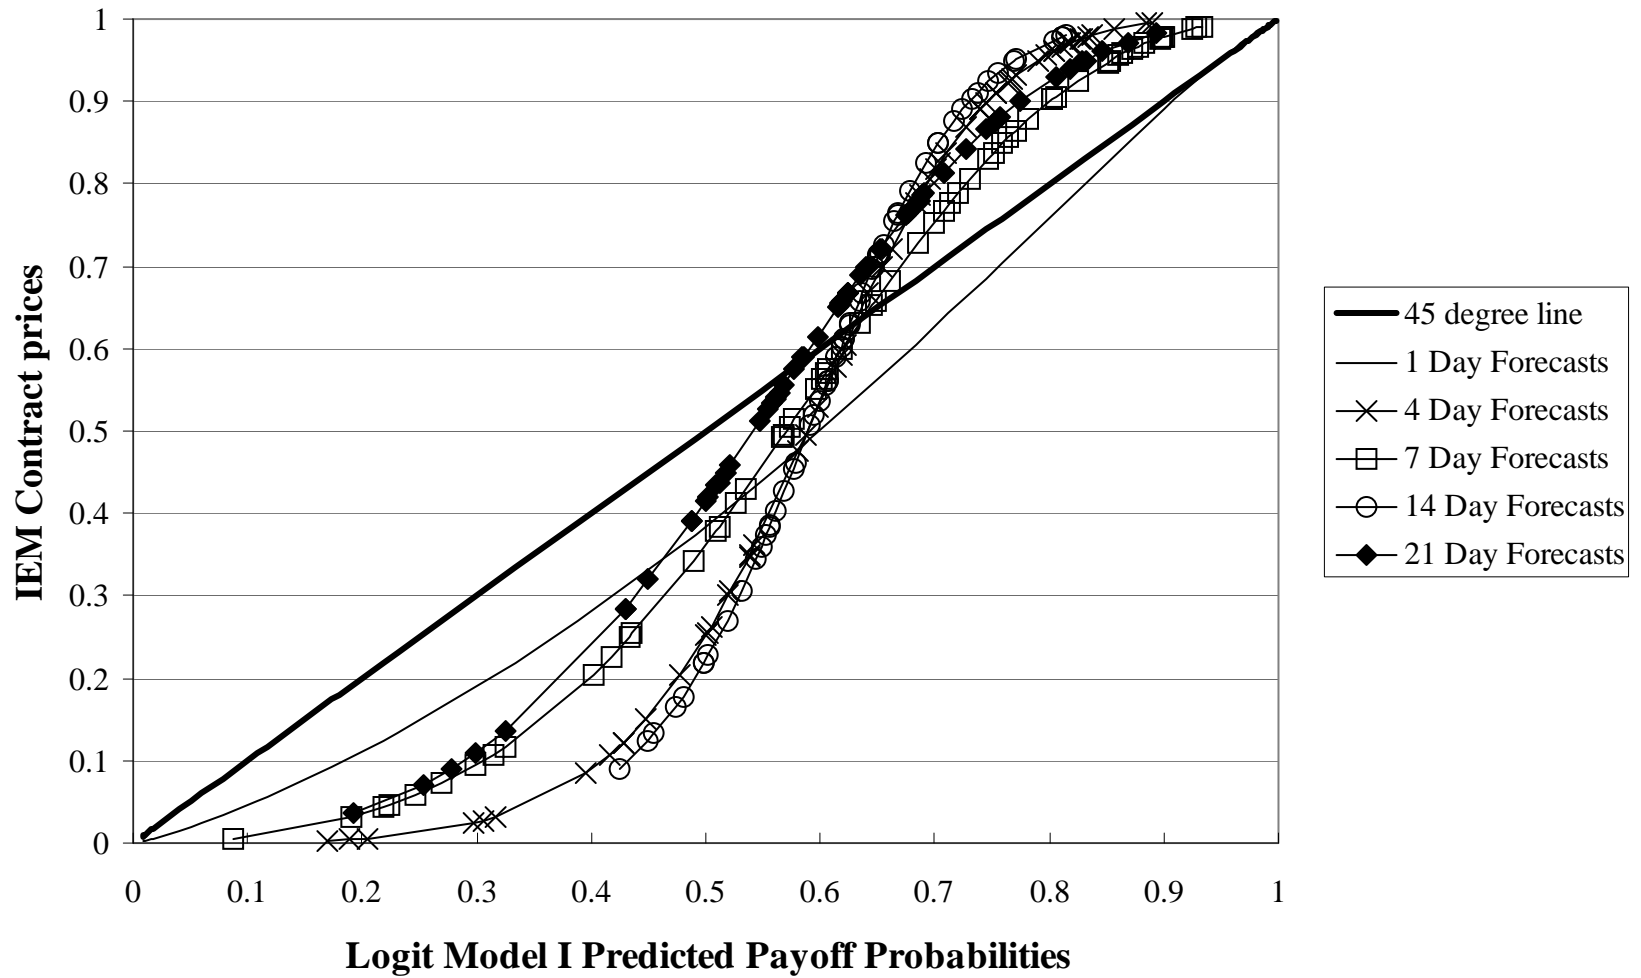


Figure 4: Logit Model I predicted payoff probabilities versus contract prices for the MSH contract in the Microsoft Price Level Market. Logit Model I uses the log price ratio as the only independent variable to explain the payoff probability.

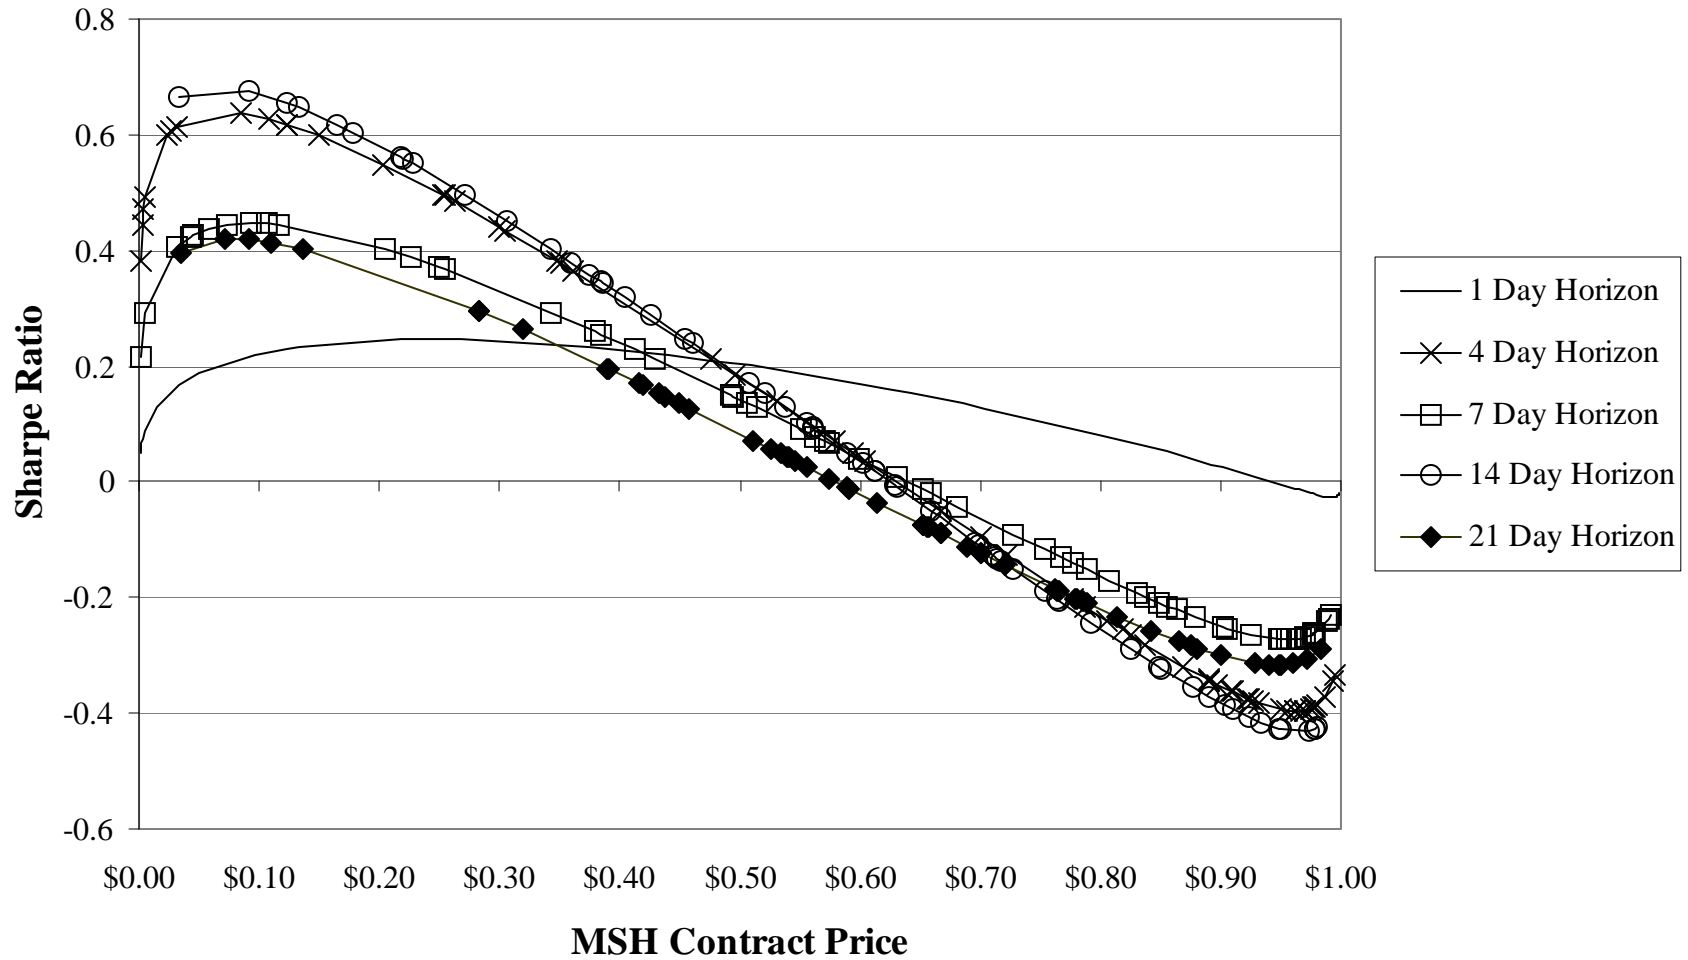


Figure 5: Logit Model I predicted Sharpe ratios versus contract prices for the MSH contract in the Microsoft Price Level Market. Logit Model I uses the log price ratio as the only independent variable to explain the payoff probability.




Figure 6: Logit Model I predicted Sharpe ratios versus contract prices for contracts in the Computer Industry Returns Market at a 14-day horizon with 4th order polynomial trend lines. Logit Model I uses the own log price ratios as the only independent variables (with cross-price coefficients restricted to zero) to explain the payoff probability.

Tables

Table 1: Average Prices, Payoffs and Dollar Returns to Purchases
(Aggregated by Price Range across all Contracts and Markets)

Price Range	Item	Days to Liquidation Determination					
		1	2	4	7	14	21
\$0.0-\$0.2	Avg. Payoff	\$0.0372	\$0.0825	\$0.1256	\$0.1170	\$0.1622	\$0.1504
	Avg. Price	\$0.0300	\$0.0342	\$0.0510	\$0.0588	\$0.0804	\$0.0896
	Diff.	\$0.0072	\$0.0482	\$0.0746	\$0.0582	\$0.0818	\$0.0608
	n	215	206	199	188	148	133
	tstat	0.5884	2.5921*	3.2098*	2.529*	2.6724*	1.9294
\$0.2-\$0.4	Avg. Payoff	\$0.3000	\$0.3667	\$0.2432	\$0.3333	\$0.3023	\$0.3837
	Avg. Price	\$0.3063	\$0.2973	\$0.2874	\$0.2952	\$0.2909	\$0.3012
	Diff.	\$(0.0063)	\$0.0694	\$(0.0442)	\$0.0381	\$0.0115	\$0.0825
	n	20	30	37	51	86	86
	tstat	-0.0606	0.8092	-0.6153	0.5746	0.2316	1.5994
\$0.4-\$0.6	Avg. Payoff	\$0.4400	\$0.4000	\$0.4848	\$0.3939	\$0.4259	\$0.3857
	Avg. Price	\$0.4945	\$0.5013	\$0.5036	\$0.5077	\$0.4860	\$0.4830
	Diff.	\$(0.0545)	\$(0.1013)	\$(0.0188)	\$(0.1137)	\$(0.0601)	\$(0.0973)
	n	25	25	33	33	54	70
	tstat	-0.5059	-1.057	-0.2198	-1.3591	-0.9045	-1.7201
\$0.6-\$0.8	Avg. Payoff	\$0.6667	\$0.5217	\$0.5926	\$0.5697	\$0.6667	\$0.5556
	Avg. Price	\$0.6924	\$0.6767	\$0.7142	\$0.7062	\$0.6952	\$0.6953
	Diff.	\$(0.0258)	\$(0.1549)	\$(0.1216)	\$(0.1165)	\$(0.0285)	\$(0.1397)
	n	18	23	27	39	48	36
	tstat	-0.2245	-1.4786	-1.2406	-1.4513	-0.4322	-1.6889
\$0.8-\$1.0	Avg. Payoff	\$0.9551	\$0.8675	\$0.7887	\$0.8393	\$0.5926	\$0.7500
	Avg. Price	\$0.9506	\$0.9388	\$0.9199	\$0.9213	\$0.9059	\$0.9167
	Diff.	\$0.0044	\$(0.0742)	\$(0.1311)	\$(0.0820)	\$(0.3133)	\$(0.1667)
	n	89	83	71	56	27	20
	tstat	0.2004	-1.9524	-2.7083*	-1.6461	-3.2092*	-1.6118

*Significant at the 95% level of confidence.

Table 2: Average Prices, Payoffs and Dollar Returns to Purchases
(Aggregated by Price Quintiles across all Contracts and Markets)

Price Quintile	Item	Days to Liquidation Determination					
		1	2	4	7	14	21
1 (Lowest 20%)	Avg. Payoff	\$0.0135	\$0.0133	\$0.0405	\$0.0541	\$0.1370	\$0.1449
	Avg. Price	\$0.0016	\$0.0033	\$0.0076	\$0.0133	\$0.0334	\$0.0473
	Diff.	\$0.0119	\$0.0101	\$0.0329	\$0.0408	\$0.1036	\$0.0977
	n	74	75	74	74	73	69
	tstat	0.8857	0.7578	1.4316	1.5533	2.5383*	2.2631*
2	Avg. Payoff	\$0.0121	\$0.0694	\$0.1644	\$0.1233	\$0.1918	\$0.1739
	Avg. Price	\$0.0137	\$0.0210	\$0.0418	\$0.0563	\$0.1241	\$0.1407
	Diff.	\$(0.0016)	\$0.0484	\$0.1226	\$0.0670	\$0.0677	\$0.0332
	n	73	72	73	73	73	69
	tstat	-0.1146	1.6063	2.7980*	1.7281	1.4523	0.7269
3	Avg. Payoff	\$0.0912	\$0.2027	\$0.2000	\$0.2568	\$0.2639	\$0.3333
	Avg. Price	\$0.0946	\$0.1215	\$0.1630	\$0.1966	\$0.2680	\$0.2920
	Diff.	\$(0.0034)	\$0.0813	\$0.0370	\$0.0601	\$(0.0041)	\$0.0414
	n	74	74	75	74	72	69
	tstat	-0.1042	1.7458	0.7933	1.1846	-0.0785	0.7329
4	Avg. Payoff	\$0.5753	\$0.4795	\$0.4861	\$0.4658	\$0.4384	\$0.3768
	Avg. Price	\$0.5950	\$0.5701	\$0.5452	\$0.5177	\$0.4682	\$0.4517
	Diff.	\$(0.0197)	\$(0.0907)	\$(0.0591)	\$(0.0520)	\$(0.0298)	\$(0.0749)
	n	73	73	72	73	73	69
	tstat	-0.3558	-1.5487	-1.0259	-0.9173	-0.5178	-1.2680
5 (Highest 20%)	Avg. Payoff	\$0.9726	\$0.9041	\$0.7808	\$0.7671	\$0.6389	\$0.6377
	Avg. Price	\$0.9701	\$0.9538	\$0.9165	\$0.8844	\$0.7780	\$0.7351
	Diff.	\$0.0025	\$(0.0497)	\$(0.1357)	\$(0.1173)	\$(0.1391)	\$(0.0974)
	n	73	73	73	73	72	69
	tstat	0.1256	-1.4410	-2.8156*	-2.4419*	-2.4017*	-1.6555

*Significant at the 95% level of confidence.

Table 3: Logit Models for the Microsoft Price Level Market

MODEL I		Days to Liquidation Determination					
	Null	1	2	4	7	14	21
Obs		61	61	61	61	61	58
ln(p/(1-p))	1	0.8549	0.7019	0.3285*	0.5022*	0.2888*	0.4859*
(Std. Dev.)		(0.2193)	(0.1926)	(0.1249)	(0.1516)	(0.2027)	(0.2390)
Intercept	0	0.4040	0.4186	0.3583	(0.2842)	0.3613	0.1674
(Std. Dev.)		(0.4561)	(0.3600)	(0.2922)	(0.3115)	(0.2855)	(0.2776)
Log Likelihood		-16.9353	-23.8667	-35.4848	-33.2567	-39.0865	-36.0661
Pseudo R2		0.5754	0.4016	0.1221	0.1772	0.0330	0.0741
MODEL II		Days to Liquidation Determination					
	Null	1	2	4	7	14	21
Obs		61	61	61	61	61	56
ln(p/(1-p))	1	0.8319	0.6651	0.4493*	0.5065*	0.3879*	0.2830*
(Std. Dev.)		(0.1711)	(0.2183)	(0.1412)	(0.1750)	(0.2182)	(0.2920)
D(ln(p/(1-p)))	0	0.0687	0.7635	-0.5664*	0.4341	-0.4791	0.7057
(Std. Dev.)		(0.3615)	(0.4096)	(0.2668)	(0.2462)	(0.4850)	(0.5650)
Intercept	0	0.3812	0.6639*	0.2813	0.3466	0.3475	0.1992
(Std. Dev.)		(0.4123)	(0.3217)	(0.3008)	(0.3293)	(0.3024)	(0.2987)
Log Likelihood		-16.9101	-21.2334	-33.6187	-32.0100	-38.4254	-34.5151
Pseudo R2		0.5760	0.4676	0.1682	0.2080	0.0493	0.0898
LR Test vs. Model I (Chi2(1))		0.05	5.27*	3.73	2.49	1.32	1.74
MODEL III		Days to Liquidation Determination					
	Null	1	2	4	7	14	21
Obs		61	61	61	61	61	58
ln(p/(1-p))	1	0.7171	0.5733	-0.1604*	0.1862*	-0.2863*	0.6062
(Std. Dev.)		(0.3804)	(0.4012)	(0.2240)	(0.2587)	(0.3757)	(0.3208)
P(MSFT)/Cut	0	6.4810	4.6603	18.7116*	9.1976	16.5609*	-3.4020
(Std. Dev.)		(12.5475)	(13.42055)	(9.3916)	(7.5309)	(8.1080)	(5.7318)
Intercept	0	0.4264	0.4352	0.5234	0.3943	0.5968	0.1200
(Std. Dev.)		(0.4419)	(0.3566)	(0.2839)	(0.3166)	(0.3760)	(0.2835)
Log Likelihood		-16.8460	-23.7374	-32.1695	-32.5394	-35.0914	-35.9416
Pseudo R2		0.5776	0.4048	0.1931	0.1949	0.1318	0.0745
LR Test vs. Model I (Chi2(1))		0.18	0.26	5.75*	1.43	7.99*	0.25

The dependent variable in all models is the probability of a \$1 payoff for the MSH contract. In Model I, the only independent variable is ln(p/(1-p)). In Model II, the one-day change in ln(p/(1-p)) is added to Model I as an independent variable. In Model III, the current Microsoft stock price divided by the payoff-determining cutoff is added to Model I as an independent variable. All likelihood ratio tests are relative to (the restricted) Model I.

*Significantly different from the Null at the 95% level of confidence.

Table 4: Models for the Computer Industry Returns Market							
MODEL I: Own Price Ratios Only							
Days to Liquidation Determination							
	Null Value	1	2	4	7	14	21
AAPL Component							
ln(p(AAPL)/p(SP500))	1	0.8312	0.5915*	0.7644	0.7972	0.5073*	0.1026*
(Std. Dev.)		(0.1807)	(0.1496)	(0.1776)	(0.1974)	(0.1848)	(0.2402)
Intercept	0	0.1052	0.2782	0.1372	0.2095	0.8911	1.0408*
(Std. Dev.)		(0.6277)	(0.7011)	(0.7208)	(0.6771)	(0.5039)	(0.4767)
IBM Component							
ln(p(IBM)/p(SP500))	1	1.0152	0.7719	0.8991	0.8876	0.4359*	0.4869*
(Std. Dev.)		(0.2732)	(0.1690)	(0.1791)	(0.1953)	(0.2117)	(0.2535)
Intercept	0	-0.4205	-0.2224	-0.3739	-0.2973	0.4275	0.2030
(Std. Dev.)		(0.8627)	(0.6653)	(0.7480)	(0.6729)	(0.5268)	(0.5978)
MSFT Component							
ln(p(MSFT)/p(SP500))	1	0.8102	0.5125*	0.3819*	0.3609*	0.4287*	0.3564*
(Std. Dev.)		(0.2228)	(0.1446)	(0.1312)	(0.1637)	(0.1872)	(0.2700)
Intercept	0	0.2596	0.6075	0.8076	0.7867	0.5806	0.5633
(Std. Dev.)		(0.5609)	(0.5452)	(0.5288)	(0.5194)	(0.5399)	(0.5194)
Joint χ^2 Test Statistic (3 dof)		2.81	19.38*	22.82*	16.94*	14.26*	15.32*
Obs		55	57	60	60	60	57
Log Likelihood		-28.5956	-44.8707	-50.7949	-54.0141	-72.6081	-72.3809
Pseudo R2		0.6057	0.4000	0.3552	0.3143	0.0763	0.0353

Table 4 (Continued): Models for the Computer Industry Returns Market

MODEL II: Own and Cross Price Ratios							
Days to Liquidation Determination							
	Null Value	1	2	4	7	14	21
AAPL Component							
ln(p(AAPL)/p(SP500))	1	2.1332*	0.5980*	0.7929	0.9287	0.3944*	0.2693
(Std. Dev.)		(0.5708)	(0.1771)	(0.2420)	(0.2310)	(0.2319)	(0.4741)
ln(p(IBM)/p(SP500))	0	-0.6549	-0.2728	-0.3529	-0.0754	-0.3094	-0.1442
(Std. Dev.)		(0.4185)	(0.2717)	(0.3473)	(0.3535)	(0.5288)	(0.4058)
ln(p(MSFT)/p(SP500))	0	0.1248	0.2404	0.2293	-0.1268	-0.0257	-0.7013
(Std. Dev.)		(0.3157)	(0.3882)	(0.2611)	(0.3024)	(0.5463)	(0.5155)
Intercept	0	1.0792	0.2610	0.0473	0.3647	1.2838*	2.2697*
(Std. Dev.)		(1.3216)	(0.8082)	(0.7002)	(0.7162)	(0.5799)	(0.9504)
IBM Component							
ln(p(AAPL)/p(SP500))	0	0.2671	-0.4044	-0.1073	0.2609	0.0427	0.3144
(Std. Dev.)		(0.4419)	(0.2124)	(0.1997)	(0.1991)	(0.2444)	(0.4248)
ln(p(IBM)/p(SP500))	1	1.0397	0.9525	0.8842	0.8682	0.3990	0.4814
(Std. Dev.)		(0.2732)	(0.2308)	(0.3835)	(0.3644)	(0.4887)	(0.4181)
ln(p(MSFT)/p(SP500))	0	1.1556*	0.5194	-0.1022	-0.3732	-0.4486	-1.0078
(Std. Dev.)		(0.4396)	(0.4165)	(0.3367)	(0.3734)	(0.5541)	(0.5395)
Intercept	0	0.4702	-0.2686	-0.5033	-0.1758	0.8004	1.5183
(Std. Dev.)		(1.4449)	(0.8481)	(0.7610)	(0.7479)	(0.6247)	(1.0323)
MSFT Component							
ln(p(AAPL)/p(SP500))	0	0.5879	-0.2271	0.0792	0.1985	-0.2681	-0.1298
(Std. Dev.)		(0.3789)	(0.21362)	(0.2193)	(0.2219)	(0.3022)	(0.4693)
ln(p(IBM)/p(SP500))	0	0.1220	0.3306	-0.2546	-0.2452	-0.2755	-0.2327
(Std. Dev.)		(0.2850)	(0.2293)	(0.3013)	(0.2893)	(0.4985)	(0.3760)
ln(p(MSFT)/p(SP500))	1	1.4731	0.7811	0.4458*	0.2223*	0.4456	-0.0221*
(Std. Dev.)		(0.3657)	(0.4096)	(0.2595)	(0.2822)	(0.5488)	(0.4751)
Intercept	0	1.9103	0.7464	0.8447	1.0078	0.8906	1.5859
(Std. Dev.)		(1.2418)	(0.6965)	(0.5089)	(0.5196)	(0.6243)	(0.9350)
Joint χ^2 Test Statistic (3 dof)		6.09	6.76	8.96*	11.68*	15.40*	18.00**
Obs		55	57	60	60	60	57
Log Likelihood		-22.3661	-40.9907	-49.18.6	-52.9354	-70.0943	-68.2869
Pseudo R2		0.6916	0.4519	0.3757	0.3280	0.1083	0.0857
LR Test vs. Model I (Chi2(6))		12.46	7.76	3.23	2.16	5.03	8.19

Table 4 (Continued): Models for the Computer Industry Returns Market

MODEL III: Own Price Ratios and Own Relative Return							
	Null Value	Days to Liquidation Determination					
		1	2	4	7	14	21
AAPL Component							
ln(p(AAPL)/p(SP500))	1	0.0425*	-0.0107*	0.1615*	0.1548*	-0.3153*	-0.8051*
(Std. Dev.)		(0.2954)	(0.2617)	(0.1953)	(0.2295)	(0.2325)	(0.3776)
R(AAPL)-Max(R(.))	0	37.2185	15.0964*	10.5711*	10.6456*	10.0743*	14.2847
(Std. Dev.)		(23.8098)	(6.7977)	(4.0316)	(4.2570)	(3.8167)	(8.5370)
Intercept	0	0.9570	1.1438	0.8620	0.9169	1.7133*	2.3030*
(Std. Dev.)		(0.7856)	(0.8577)	(0.7248)	(0.6870)	(0.6006)	(0.7327)
IBM Component							
ln(p(IBM)/p(SP500))	1	0.5684	0.2575*	0.3869*	0.1881*	-0.3729*	-0.1174*
(Std. Dev.)		(0.3159)	(0.2646)	(0.2072)	(0.2593)	(0.3010)	(0.2961)
R(IBM)-Max(R(.))	0	14.6792	24.4116	16.2386	22.8276*	21.7421*	12.9387
(Std. Dev.)		(10.0518)	(13.0333)	(9.0827)	(10.8520)	(8.8743)	(7.8880)
Intercept	0	0.01597	0.8854	0.2687	.6842	1.7609*	1.3155
(Std. Dev.)		(0.7673)	(0.9077)	(0.8404)	(0.8773)	(0.8169)	(0.7810)
MSFT Component							
ln(p(MSFT)/p(SP500))	1	0.4076*	0.1164*	-0.1278*	-0.1173*	-0.2458*	-0.2797*
(Std. Dev.)		(0.2978)	(0.2095)	(0.2004)	(0.2120)	(0.2548)	(0.3772)
R(IBM)-Max(R(.))	0	8.3428	9.9092	14.3312	9.3323	9.8242	10.7831
(Std. Dev.)		(13.3637)	(7.7181)	(7.8889)	(5.3186)	(5.3854)	(6.7212)
Intercept	0	0.3886	1.0217	1.6158*	1.4573*	1.8364*	1.7765*
(Std. Dev.)		(0.6402)	(0.7055)	(0.7475)	(0.6549)	(0.7867)	(0.7972)
Joint χ^2 Test Statistic (3 dof)		11.19	26.51*	37.87*	31.37*	41.61*	31.13*
Obs		55	57	60	60	60	57
Log Likelihood		-24.4788	-39.2301	-44.9220	-48.4262	-62.3990	-64.7645
Pseudo R2		0.6624	0.4754	0.4297	0.3853	0.2062	0.1368
LR Test vs. Model I (Chi2(3))		8.23*	11.28*	11.75*	11.18*	20.42*	15.23*

The dependent variables in all models are the probabilities of \$1 payoffs for each contract. In all models, the base contract is the S&P500 contract. In Model I, the independent variables are $\ln(p_{AAPL}/p_{SP500})$, $\ln(p_{IBM}/p_{SP500})$ and $\ln(p_{MSFT}/p_{SP500})$ and coefficients on cross-log-price-ratio terms are restricted to zero. In Model II, the independent variables are the same but cross-log-price-ratio terms are unrestricted. In Model III, cross-log-price-ratio restrictions are restored and a relative standing variable (the contracts return minus the maximum of the other three returns) is added as an independent variable in each component regression. All cross-standing terms are restricted to be zero. The null in the joint χ^2 tests is for all own price ratio coefficients to equal 1. All likelihood ratio tests are relative to (the restricted) Model I.

*Significant from the Null at the 95% level of confidence.

Table 5: Portfolio Characteristics of a Mean-Variance Utility Maximizing Trader in the Microsoft Price Level Market

Panel A: Trader Uses Model I to Predict Payoff Probabilities						
	Quantity of MSH Held	Quantity of MSL Held	Quantity Volume	Expected Value of Portfolio	Variance in Portfolio Value	Monthly Return
Portfolio statistics for $\gamma = 0.50$				Final Value: \$120.35		
Average	0.23	0.28	0.21	107.70	0.21	0.31%
Std. Dev.	0.42	0.45	0.43	8.42	0.22	0.81%
Max	2.00	1.00	2.00	120.39	0.99	4.98%
Min	0.00	0.00	0.00	99.87	0.00	-0.33%
Geo. Avg.						0.30%
Portfolio statistics for $\gamma = 0.25$				Final Value: \$146.65		
Average	0.52	0.59	0.48	118.57	0.47	0.64%
Std. Dev.	0.83	0.79	0.69	18.46	0.37	1.61%
Max	3.00	2.00	4.00	146.85	1.50	9.31%
Min	0.00	0.00	0.00	100.93	0.00	-0.48%
Geo. Avg.						0.63%
Portfolio statistics for $\gamma = 0.10$				Final Value: \$207.18		
Average	1.32	1.46	1.22	142.67	1.19	1.24%
Std. Dev.	2.03	1.87	1.43	42.23	0.86	3.08%
Max	8.00	6.00	10.00	207.51	3.97	18.33%
Min	0.00	0.00	0.00	102.26	0.00	-1.22%
Geo. Avg.						1.20%
Panel B: Trader Uses Model III to Predict Payoff Probabilities						
	Quantity of MSH Held	Quantity of MSL Held	Quantity Volume	Expected Value of Portfolio	Variance in Portfolio Value	Monthly Return
Portfolio statistics for $\gamma = 0.50$				Final Value: \$410.83		
Average	0.60	0.30	1.02	199.50	0.40	4.76%
Std. Dev.	10.10	0.55	13.88	139.33	4.77	28.46%
Max	366.00	7.00	366.00	410.97	173.02	219.41%
Min	0.00	0.00	0.00	100.55	0.00	-5.47%
Geo. Avg.						2.34%
Portfolio statistics for $\gamma = 0.25$				Final Value: \$722.83		
Average	1.24	0.61	2.05	300.96	0.82	7.23%
Std. Dev.	20.17	1.03	27.70	278.45	9.53	43.13%
Max	731.00	13.00	730.00	723.25	345.57	353.85%
Min	0.00	0.00	0.00	100.97	0.00	-9.26%
Geo. Avg.						3.30%
Portfolio statistics for $\gamma = 0.10$				Final Value: \$1,655.28		
Average	3.11	1.53	5.17	601.64	2.06	11.32%
Std. Dev.	50.43	2.50	69.28	695.43	23.83	65.77%
Max	1,828.00	33.00	1,826.00	1,656.10	864.16	581.22%
Min	0.00	0.00	0.00	102.73	0.00	-14.02%
Geo. Avg.						4.71%

The trader starts with \$100 in cash. At each date, the trader estimates the probability of payoff for each contract using logit Model I or Model III (each described in Table 3). Then, the trader maximizes the utility function $U(x)=E(x)-\gamma\text{Var}(x)$ where x is the payoff for the current month's contracts. The trader re-estimates and re-optimizes daily from 21 days to 1 day before liquidation value determination.

Table 6: Portfolio Characteristics of a Mean-Variance Utility Maximizing Trader in the Computer Industry Returns Markets

Panel A: Trader Uses Model I to Predict Payoff Probabilities								
	Quantity of AAPL Held	Quantity of IBM Held	Quantity of MSFT Held	Quantity of SP500 Held	Quantity Volume	Expected Value of Portfolio	Variance in Portfolio Value	Monthly Return
Portfolio statistics for $\gamma = 0.50$					Final Value: \$109.41			
Average	0.88	0.37	0.59	0.79	0.94	\$104.28	0.24	0.16%
Std. Dev.	0.81	0.75	0.62	0.79	1.63	\$2.92	0.26	0.48%
Max	13.00	12.00	3.00	13.00	26.00	\$109.58	3.67	1.17%
Min	0.00	0.00	0.00	0.00	0.00	\$98.89	0.00	-1.17%
Geo. Avg.								0.15%
Portfolio statistics for $\gamma = 0.25$					Final Value: \$127.25			
Average	1.81	0.80	1.13	1.49	1.93	\$112.73	0.89	0.39%
Std. Dev.	1.58	1.49	1.13	1.46	3.10	\$8.24	1.00	0.88%
Max	25.00	24.00	6.00	25.00	52.00	\$127.61	13.49	2.46%
Min	0.00	0.00	0.00	0.00	0.00	\$98.31	0.00	-1.54%
Geo. Avg.								0.39%
Portfolio statistics for $\gamma = 0.10$					Final Value: \$166.06			
Average	4.50	2.00	2.90	3.80	4.77	\$129.89	5.36	0.84%
Std. Dev.	3.90	3.63	2.81	3.65	7.60	\$20.19	6.02	1.88%
Max	63.00	59.00	16.00	64.00	130.00	\$166.60	85.92	4.67%
Min	0.00	0.00	0.00	0.00	0.00	\$95.45	0.00	-3.02%
Geo. Avg.								0.83%
Panel B: Trader Uses Model III to Predict Payoff Probabilities								
	Quantity of AAPL Held	Quantity of IBM Held	Quantity of MSFT Held	Quantity of SP500 Held	Quantity Volume	Expected Value of Portfolio	Variance in Portfolio Value	Monthly Return
Portfolio statistics for $\gamma = 0.50$					Final Value: \$430.46			
Average	12.48	3.45	13.38	14.26	24.64	\$224.83	0.96	2.09%
Std. Dev.	43.45	23.12	46.42	47.82	98.39	\$97.37	3.60	3.48%
Max	355.00	420.00	420.00	420.00	1,228.00	\$457.62	97.38	18.06%
Min	0.00	0.00	0.00	0.00	0.00	\$100.17	0.00	0.03%
Geo. Avg.								2.42%
Portfolio statistics for $\gamma = 0.25$					Final Value: \$648.97			
Average	20.64	5.70	21.45	23.10	37.18	\$294.91	3.58	2.75%
Std. Dev.	66.90	34.02	69.54	72.32	142.71	\$157.29	13.96	4.12%
Max	497.00	628.00	628.00	628.00	1,819.00	\$704.15	389.52	22.19%
Min	0.00	0.00	0.00	0.00	0.00	\$99.95	0.00	-0.28%
Geo. Avg.								3.11%
Portfolio statistics for $\gamma = 0.10$					Final Value: \$1,299.19			
Average	42.11	11.86	43.09	46.43	73.52	\$500.81	20.90	3.96%
Std. Dev.	128.36	62.85	132.68	137.84	275.36	\$335.37	86.08	5.15%
Max	1,186.00	1,247.00	1,247.00	1,247.00	3,579.00	\$1,439.09	2,440.43	29.83%
Min	0.00	0.00	0.00	0.00	0.00	\$100.06	0.00	-0.19%
Geo. Avg.								4.29%

The trader starts with \$100 in cash. At each date, the trader estimates the probability of payoff for each contract using logit Model I or Model III (each described in Table 4). Then, the trader maximizes the utility function $U(x)=E(x)-\gamma\text{Var}(x)$, where x is the payoff for the current month's contracts. The trader re-estimates and re-optimizes daily from 21 days to 1 day before liquidation value determination.