

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

February 2018

Forthcoming: *International Journal of Forecasting*

Abstract

We study the forecast accuracy and efficiency of popular “binary” prediction markets. Such markets forecast probabilities for future states of the world (e.g., election winners) by paying off \$0 or \$1 depending on the realized state (e.g., who actually won). To assess accuracy, forecast probabilities must be compared to realization frequencies, not individual realizations. We use Iowa Electronic Market (IEM) data to test efficiency against two alternative propositions from behavioral finance: the longshot bias and the overconfidence bias (which yield opposing predictions). No longshot bias appears in IEM markets. Nor does overconfidence influence prices at short horizons. However, overconfident traders may bias prices at intermediate horizons. While the markets are efficient at short horizons, non-market data indicate some intermediate-horizon inefficiency. We calculate Sharpe ratios for static trading strategies and document returns for dynamic trading strategies to assess the economic content of the inefficiencies.

Keywords: Prediction Markets, Market Efficiency, Longshot Bias, Overconfidence

* 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.

[†] Corresponding Author: Thomas Rietz, e-mail: Thomas-Rietz@uiowa.edu, phone: 319-335-0856.

Longshots, Overconfidence and Efficiency

on the Iowa Electronic Market

Prediction markets have become a popular means of forecasting (Berg, Forsythe, Nelson and Rietz, 2008; Tziralis and Tatsiopoulos, 2007), largely because of accuracy evidence from “linear” markets (Berg, Nelson and Rietz, 2008). The Iowa Electronic Markets (IEM) “vote-share” markets typify linear markets: contracts payoff as a linear function of a continuous variable, the vote-shares taken by candidates or parties in upcoming elections. Accuracy is measured directly by comparing actual vote shares to forecasts. Berg, Nelson and Rietz (2008) show a 1.33 percentage point average absolute difference between election-eve forecasts and outcomes across 14 US Presidential Election contracts. Further, in advance of elections, vote-share market prices outperform polls as forecasts 74% of the time.

However, the most popular prediction markets today are “binary” markets, where prices forecast discrete outcome probabilities. The IEM “winner-takes-all” markets typify binary markets: contracts associated with a candidate who “wins” the election pay off \$1 and other contracts pay off \$0. Prices should reflect probabilities of \$1 payoffs which then would equal probabilities that associated candidates will “win.” In addition to being the most popular IEM markets, binary contracts form most markets on other prediction markets such as PredictIt, InTrade (now defunct) and others.

In binary contract markets, people often have the mistaken belief that the market is “right” if the outcome forecast as most likely actually occurs. If not, the market is “wrong.” For example, in the 2004 US Presidential Election, InTrade ran 51 markets on the election winner in each individual State and the District of Columbia. They famously claimed to have predicted every race correctly (e.g., Authors, 2013). Authors goes on, illustrating the interpretation of “right” and “wrong” stating, “Last year, [InTrade] showed it was not flawless by getting Florida wrong. Mr Obama held the state, but Intrade had put a 69 per cent chance on a Republican victory there.” This does not indicate that the market was “wrong” or inefficient. In fact, if the market was perfectly efficient, we should observe Obama taking Florida nearly 1/3 of the time. Unless all forecasts were at or near 100% in the 2004 election (they

weren't), getting every state "right" in this sense actually indicates that the markets are strikingly inefficient: the winning frequencies do not match up with the forecast probabilities.

Herein lies the challenge of assessing efficiency in binary prediction markets: one has to compare forecast probabilities to outcome frequencies. True, pre-election prediction markets only gave Trump an 18% chance of winning (Silver, 2016) and pre-Brexit prediction markets forecast about an 85% probability that Great Britain would stay in the Economic Union (The Economist, 2016).¹ But, that means that, if we could observe 100 Brexit votes under identical circumstances, 15 of them should result in Britain leaving the EU. Trump should win 18 in 100 elections. Were the outcomes we saw simply one of those 15 or 18? Or, were there biases in prices? If so, how large were the biases and what caused them? The question isn't whether prediction markets "miss" from time to time, but whether they "miss" at the right rate. One Trump election or one Brexit vote cannot answer these questions. We can only answer them by aggregating across many, essentially similar, markets.

Fortunately, the IEM ran a unique long series of essentially similar markets based on monthly stock prices and returns. While these are not exact replications, they are similar enough to credibly test binary market efficiency by comparing forecasts to outcome frequencies. This allows us to study both when and why biases might arise in binary market prices. Specifically, we use frequency and logistic analysis on these markets to test whether market prices appear efficient against two competing alternative biases postulated by behavioral finance: the longshot bias and the overconfidence bias. In binary prediction markets, a longshot bias would result in over-pricing low probability contracts and under-pricing high probability contracts. In contrast, an overconfidence bias would result in under-pricing low probability contracts and over-pricing high probability contracts.²

¹ We note that the IEM differs from most markets in that "winning" is based on who takes the majority of the popular vote. The election eve forecast was that Clinton would win the popular vote with a 0.775 probability and an expected vote share of 53.7%. She actually did win the popular vote with 51.1%.

² Over-pricing can arise from traders in search of large, but unlikely, returns driving up prices of low priced contracts, similar to "over-betting" on longshots documented in the racetrack betting literature (e.g., Ziemba and Hausch, 1986) or by traders overestimating the probability of unlikely outcomes (as in prospect theory, Kahneman and Tversky, 1979). Under-pricing can be driven by traders either over estimating probabilities for events they deem likely (e.g., Lichtenstein, Fischhoff and Phillips, 1982) or by overreacting to information that makes an event more likely (e.g., Daniel, Hirshleifer and Subrahmanyam, 1998). We note that there are other effects of overconfidence

Testing for longshot or overconfidence biases is very difficult using prices from naturally occurring financial markets. There is no definitive point in time when the value of a stock is known with certainty making it challenging to assess whether prices at any point in time are actually biased. In addition, stocks, indices, “regular” options and futures are much more complex than simple binary options. Stock values depend on the distribution of dividends over the infinite horizon. Index values depend on values of underlying component stocks. Future and option values depend on future distributions of values for underlying stocks or indices. This complexity makes it difficult to detect the effects of simple biases even if we knew the “true” values of the underlying assets.

In contrast, the IEM is a good testing ground to assess whether behavioral biases affect market prices. The IEM structure closely resembles naturally occurring financial markets and increasing evidence shows that IEM market prices behave like any other financial market.³ Traders buy and sell contracts in a double auction, can reverse their positions any time and can synthetically short sell. The traded contracts are particularly simple and, at the close of each market, their values are known with certainty. The markets we examine are intermediate horizon markets (running up to 5 weeks each) and are repeated under essentially identical conditions. The market structure and 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 frequency or logit analysis) to determine whether biases exist. At any given point in time, prices should be sufficient statistics for estimating payoff probabilities. This allows us to test for price efficiency by adding past IEM prices and outside information to our analyses.

While there is ample evidence that IEM traders have biases and sometimes behave irrationally (e.g., Forsythe, Rietz and Ross, 1999; Oliven and Rietz, 2004; and Berg and Rietz, 2006), there is also

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 For example, Bonderanko and Bossaerts (2000) show that IEM prices evolve consistently with rational learning and updating. Majumder, Diermeier, Rietz and Amaral (2009) show that the distribution of returns in IEM markets mirror closely return distributions in other financial and derivatives markets.

reason to believe that financial markets may be efficient in spite of these biases. First, a few rational traders can arbitrage biases out of existence in markets. However, if all traders are affected by the same pervasive bias (e.g., overweighting low probability events), no traders may recognize or act to eliminate the bias. Even if they do, Grossman and Stiglitz (1980) argue 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, the market structure may reduce biases. In particular, there are account limits in the IEM so no trader is large relative to the market, there are no trading fees and there is a particularly simple arbitrage structure that allows traders to easily exploit arbitrage opportunities and allows synthetic short-selling. These features are intentionally designed to promote price efficiency. Third, 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 can price efficiently even when populated with “zero-intelligence” robot traders. Primarily using linear markets, Forsythe, Nelson, Neuman and Wright (1992) and Forsythe, Rietz and Ross (1999) document biases in many traders yet show that prices are driven to efficient levels by more rational, price-setting traders.

Our frequency analysis results are consistent with an overconfidence bias in which the highest priced contracts pay off with a lower frequency than the prices would indicate and vice versa for the lowest priced contracts. However, contracts with intermediate prices appear relatively efficiently priced. Further, even the effect at the extremes is transitory. In fact, price-frequency relationships move exactly as predicted by Daniel, Hirshleifer and Subrahmanyam (1998). Initial trading prices in 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.

Our baseline logit results mirror the frequency analysis. In addition, the logit models show that past IEM prices do not improve predictive accuracy. However, at intermediate horizons, adding outside information can improve accuracy in a manner that is also consistent with Daniel, Hirshleifer and

Subrahmanyam (1998). Traders appear to have a transitory overreaction to outside information at intermediate horizons.

We show the economic significance of the transitory overconfidence. A trader can achieve a significant positive Sharpe ratio for a static strategy of buying and holding the lowest priced contracts at intermediate horizons. Dynamic trading strategies may also lead to significant profits. However, traders engaging in either strategy would drive prices toward more efficient levels and reduce the profits from the strategies.

Our evidence informs research on prediction markets by providing systematic tests of efficiency for binary prediction markets. Our method contributes to measuring accuracy in discrete outcome forecasting. Any model that (1) predicts a probability of an outcome and (2) can be used under similar repeated conditions can be tested in the manner we develop in this paper. Our evidence informs research on behavioral finance by showing clear evidence (1) against a longshot bias affecting prices, (2) against overweighting of low probability events affecting prices and (3) for a transitory, information based overconfidence bias. Finally, evidence on the IEM suggest the markets behave very similarly to naturally occurring markets, bolstering the case for the external validity of our results.

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 University of Iowa Tippie College of Business. Traders invest their own money to trade, bearing the real-money risks and reaping the real-money rewards of their

activity.⁴ The markets discussed here have thousands of registered traders and, often, hundreds of active traders. They run from four to five weeks each and were liquidated and reinitialized monthly for more than five years.

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.

We study the Microsoft Price Level and Computer Industry Returns Markets. Prospectuses for these markets appear in the on-line appendices. 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).⁵ Each market contains a complete “slate” of binary contracts, one of which will liquidate at \$1 depending on the state of the world on the liquidation date. Other contracts expire worthless. 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, initial slates consist 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.⁶ Thus, these contracts are simple binary options.

⁴ Because IEM contracts are real futures contracts, the IEM is under the regulatory purview of the Commodity Futures Trading Commission (CFTC). The CFTC issued “no-action” letters 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.

⁵These dates were chosen because the contract values are linked to underlying values of stocks on option expiration dates (the third Fridays of each month).

⁶Cutoffs 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 when stock prices deviate significantly from the cutoff and contract prices reach extreme levels (close to 0 or 1). Only one split occurred in the data set used here. This split created a “middle” range contract in addition to the “L” and “H” contracts. For consistency, only the original “L” and “H” contracts are used in the data analysis here. Making a different choice or omitting data from this month does not change results.

In the Computer Industry Returns Market, a slate consists of four contracts. Each contract corresponds to an underlying security. 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 if its underlying security has the highest return from third Friday to third Friday.⁷ Thus, these contracts are also simple state contingent claims, though with a somewhat more complex interrelationship than the contracts in the Microsoft Price Level Market.

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 slate 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.⁸ Let p_t^j be the price of contract j on date t and $q(V_T^j = \mathbf{1} | I_t) := q_t^j$ 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:

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

⁷We 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.

⁸See 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 equals 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, Rietz and Ross, 1999).

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.⁹

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 models¹⁰ 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.¹¹ 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.¹² Since the contract price should represent the probability that a contract will liquidate at \$1 conditional on all available

⁹ Some 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.

¹⁰ 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.

¹¹The selection of the particular base contract is irrelevant. It has no effect on the estimates or predictions.

¹²We 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

information, price ratios should be sufficient statistics for forecasting probabilities of payoffs. In particular, consider using as independent variables:

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

If prices do indeed equal probabilities, then it should be the case that $\alpha_t^j = \mathbf{0}$ and $\beta_t^{j,j} = \mathbf{1}$ for all j and $\beta_t^{i,j} = \mathbf{0}$ 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 implied by efficient markets), 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 serves as the means 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 prices are unbiased and efficient, we should observe estimates consistent with $\alpha_t^j = \mathbf{0}$ and $\beta_t^{j,j} = \mathbf{1}$ for all j and $\beta_t^{i,j} = \mathbf{0}$ 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 $\beta_t^{j,j} > \mathbf{1}$. 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.^{13,14} Paying too much for a low-payoff-

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.

¹³ 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).

probability contract here corresponds to betting too much on the longshot in a horse race. This also accords with Kahneman and Tversky's (1979) proposition from Prospect Theory that people over-weight 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.¹⁵

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

¹⁴ 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:

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

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 reflects 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.

¹⁵This 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 bias.

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. We test these propositions 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 betting 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 the 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.¹⁶ 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.¹⁷

¹⁶ This 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.

¹⁷ While 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.

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). The table also shows average payoffs for contracts in each range, differences between average payoffs and prices, numbers of observations, and t-tests for differences between payoffs and prices. The data is consistent with overconfidence. Prices fall significantly below payoffs for contracts in the \$0.0 to \$0.2 range at most horizons longer than one day. Prices are significantly above payoffs for contracts in the \$0.8 to \$1.0 range at 4- and 14-day horizons. Table 2 aggregates across price quintiles (i.e., the lowest 20% of prices, next lowest 20%, etc). Aggregating this way shows clearly that the effect at the low end is driven primarily by very low priced contracts. In the lowest quintile, the average price and probability of payoff were \$0.0016 and \$0.0135, respectively. Prices corresponding to the Trump election and Brexit vote discussed in the introduction typically would fall in the 2nd or 3rd quintile, where there are no significant effects. Intermediate contracts appear to be priced efficiently.

There are several limitations inherent in this simple frequency analysis. 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. This creates a negative correlation across outcomes that can be addressed by logit and multinomial models. In addition, regression models allow us to analyze possible sources of inefficiency and search for data that may provide better predictions. Nevertheless, the

frequency analysis serves as a quick summary of the data. In spite of weaknesses, its results mirror the logit models that follow .

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.

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.¹⁸ 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.

¹⁸Using simple logit models to predict the probability of winning for all horses without considering the multinomial nature of each individual race, Asch and Quandt (1986) suggest that late changes in odds help predict winning frequencies at racetracks. In financial markets, short-horizon “momentum” effects are observed in stock market data (e.g., Jegadeesh and Titman, 1993). This might carry over the IEM if traders over- or under-react to news, creating serial correlations in prices.

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. In Model III, we add the ratio of the current Microsoft stock price to the cutoff. 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. This implies, a trader at these horizons could do better at predicting the chances of payoff by knowing the current Microsoft price relative to the cutoff in addition to knowing only the "H" contract price on the IEM.¹⁹ Thus, Model III provides limited evidence against market efficiency. 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.²⁰ 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

¹⁹This is information that the traders should have had before closing prices are determined because the closing prices in the IEM are measured at midnight.

²⁰The selection of the particular base contract is irrelevant. It has no effect on the estimates or predictions.

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 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 removing the cross-price restrictions imposed in Model I 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.²¹ 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 bias at longer horizons. Further, typically at intermediate or long horizons, additional information can help improve explanatory power in logit models for predicting payoffs. Next, we discuss the risks and return to investment strategies designed to exploit these inefficiencies.

²¹Again, 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.

III. Investment Strategies

Given the overconfidence bias evident in the IEM, one might conjecture that there are profitable investment opportunities based on price alone. To evaluate risk and returns for different contracts, we assume that the logit models give the best prediction of the actual probability of payoff.²² We go beyond simply observing a bias by extending results to the risk-return tradeoff of 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.²³ The under-pricing of low-priced contracts and over-pricing of high-priced contracts is apparent. 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,

²²This is analogous to observing 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). Assuming that the logit models give the best predictions is similar to the usual assumption in the racetrack literature that ex-post winning frequencies are the best estimate of the ex-ante actual probabilities.

²³This is similar to Figure 3 in Borghesi (2009), but shows the evolution of price deviations from frequencies through time.

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 maximum possible 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. However, implementable strategies would likely have achieved lower profits for two reasons: (1) high trading

volumes would have resulted in adverse price movements and (2) the strategies effectively required foresight that would have been unattainable at the time.

IV. Summary and Discussion

We systematically test the predictive accuracy of binary prediction markets. The question is more challenging than simply asking whether the most likely forecast outcome actually occurs or not. The question is important because, while binary prediction markets have become the most popular prediction markets, the evidence for efficiency to date rests largely on evidence from linear markets.

Our evidence shows a transitory overconfidence bias at intermediate horizons. At intermediate horizons, the most likely events will occur somewhat less often than predicted. The least likely events will occur slightly more often than predicted. However, especially at the low-probability end, this effect is especially concentrated on very low priced contracts. Even this effect disappears as the event approaches. Recent high-profile “misses” by prediction markets have forecast probabilities in ranges where, on election eve, there is still a considerable amount of residual uncertainty. On election eves, there was still about a 15% chance Britain would vote to leave the EU and an 18% chance that Trump would win the election. For three reasons, it is likely that these realizations arise from the residual uncertainty inherent in any election. First, while low, these prices are in the range where we observe no significant difference between forecast probabilities and observed frequencies according to Table 2. Second, these were election eve forecasts, when the evidence suggests transitory biases have gone away even in Table 1. Third, the IEM ran both binary and linear prediction markets in the 2016 Presidential Election. The linear market indicated a close race and was fairly accurate (predicting the vote shares with a 2.57 percentage point average absolute prediction error).²⁴

The implication for prediction markets is that well designed markets produce relatively efficient forecasts for mid-range probabilities and at short horizons. However, some transitory overconfidence

²⁴ As noted in footnote 1, the IEM differs from most markets in that “winning” is based on who takes the majority of the popular vote. On election eve, the IEM forecast a 0.775 probability that Clinton would win the popular vote. The fact that she did win the popular vote doesn’t make the IEM “right” any more than the fact that Trump won the Electoral College made “wrong” the markets based on the Electoral College (e.g. those discussed in Silver, 2016).

may affect market prices, especially at intermediate horizons and for very low or high priced contracts. We find no evidence for a longshot bias or overweighting low probability events affecting prices. The idea of supplementing frequency analysis with logit models and adding outside information can be applied whenever researchers need to test forecast efficiencies for repeated discrete events. Because the IEM is a real-money market with price dynamics that mirror naturally occurring markets, we believe our results are externally valid.

Three issues should give forecasters using binary prediction markets some caution. First, researchers need to guard against the temptation to conclude that observing the most likely forecast outcome means the forecast was “correct” and observing a less likely outcome means the forecast was “incorrect.” In fact, if the markets are producing efficient probability forecasts, we should observe such outcomes with frequencies that parallel the forecast probabilities. Second, market and contract design can be very important. Markets that impose transaction and/or profit fees can lead to mis-pricing. Markets where there is no clear arbitrage restrictions limiting prices to valid probability measures can lead to mis-pricing. Third, mis-specified contracts (e.g., contracts with poorly defined state-space/payoff correspondences) can lead to misinterpretation of the results.

While prediction markets as forecasting tools now have a thirty year history, there are still vast opportunities for future research. One area where there is little research is identifying distributions of possible outcomes associated with combinations of prediction market forecasts. While some researchers, most notably Nate Silver, aggregate prediction markets with other sources of information, most still only work with point forecasts. Only Berg, Geweke and Rietz (2010) have investigated outcome distributions that would be consistent with prices of contracts that forecast different parts of the distribution. This would be valuable in trying to understand the amount of residual uncertainty in advance of an event. Another promising area of research is to look at individual trader activity to identify which traders drive the overconfidence bias and how they do so. This may yield ways of mitigating and/or adjusting for the bias in future markets. Finally, it would be interesting to determine whether even the transitory bias remains in future prediction markets. If future traders engage in dynamic trading strategies designed to

exploit the bias, these strategies themselves should counter the bias. In fact, traders acting in this manner could be the reason for prices converging to efficient levels near liquidation dates.

References

- Asch, P, and RE Quandt, 1986, Racetrack Betting: The Professors' Guide to Strategies, Dover, MA: Auburn House.
- Authors, J, 2016, "Why Prediction Markets Still Have a Future," Financial Times, March 11, 2013, online at: <https://www.ft.com/content/9e282b2a-8a48-11e2-9da4-00144feabdc0>, accessed 2/14/16.
- Barber, BM and T Odean, 2001, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," Quarterly Journal of Economics, 261-292.
- Berg, JE, R Forsythe, FD Nelson and TA Rietz (2008) "Results from a Dozen Years of Election Futures Markets Research," in CR Plott and VL Smith, editors, The Handbook of Experimental Economics Results, Elsevier Science, Amsterdam, pp. 742-751.
- Berg, JE, J Geweke and TA Rietz, 2010 "Memoirs of an indifferent trader: Estimating forecast distributions from prediction markets," Quantitative Economics, 1:163–186.
- Berg, JE, FD Nelson and TA Rietz (2008) "Prediction Market Accuracy in the Long Run," International Journal of Forecasting, 24, 283-298.
- 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.
- Bonderanko, O, and P Bossaerts (2000) "Expectations and Learning in Iowa," Journal of Banking and Finance, 24(9):1535-1555.
- Borghesi, R, 2009, "An Examination of Prediction Market Efficiency: NBA Contracts On TradeSports," The Journal of Prediction Markets, 3(2): 65-77.
- Daniel, K, D Hirshleifer and A Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Overreactions," The Journal of Finance, 58(6): 1839-1885.
- The Economist, 2016, "Who Said Brexit was a Surprise?" June 24, 2016, online at: <https://www.economist.com/blogs/graphicdetail/2016/06/polls-versus-prediction-markets>, accessed 2/14/16.
- 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, TA Rietz and TW Ross, 1999, "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets," Journal of Economic Behavior and Organization, 39, 83-110.
- 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.
- Grossman, SJ, and JE Stiglitz, 1980, "On the Impossibility of Informationally Efficient Markets," The American Economic Review, 70, 393-408.
- 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.
- Majumder, SR, D Diermeier, TA Rietz and LAN Amaral (2009) "Price Dynamics in Political Prediction Markets," Proceedings of the National Academy of Sciences of the United States of America, 106(3): 679-684.
- 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.
- 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.
- Silver, N, 2016, "Why Five ThirtyEight Gave Trump a Better Chance Than Almost Anyone Else," November 11, 2016, online at: <http://fivethirtyeight.com/features/why-fivethirtyeight-gave-trump-a-better-chance-than-almost-anyone-else/> accessed 2/14/16.
- Thaler, RH, and WT Ziemba, 1988, "Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries," The Journal of Economic Perspectives, 2, 161-174.
- Tziralis, G, and I Tatsiopoulos, 2007, "Prediction Markets: An Extended Literature Review," The Journal of Prediction Markets, 1: 75-91.
- Ziemba, WT, and DB Hausch, 1986, Betting at the Racetrack, Vancouver and Los Angeles: Dr. Z Investments, Inc.

On-Line 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 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.

On-Line 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.

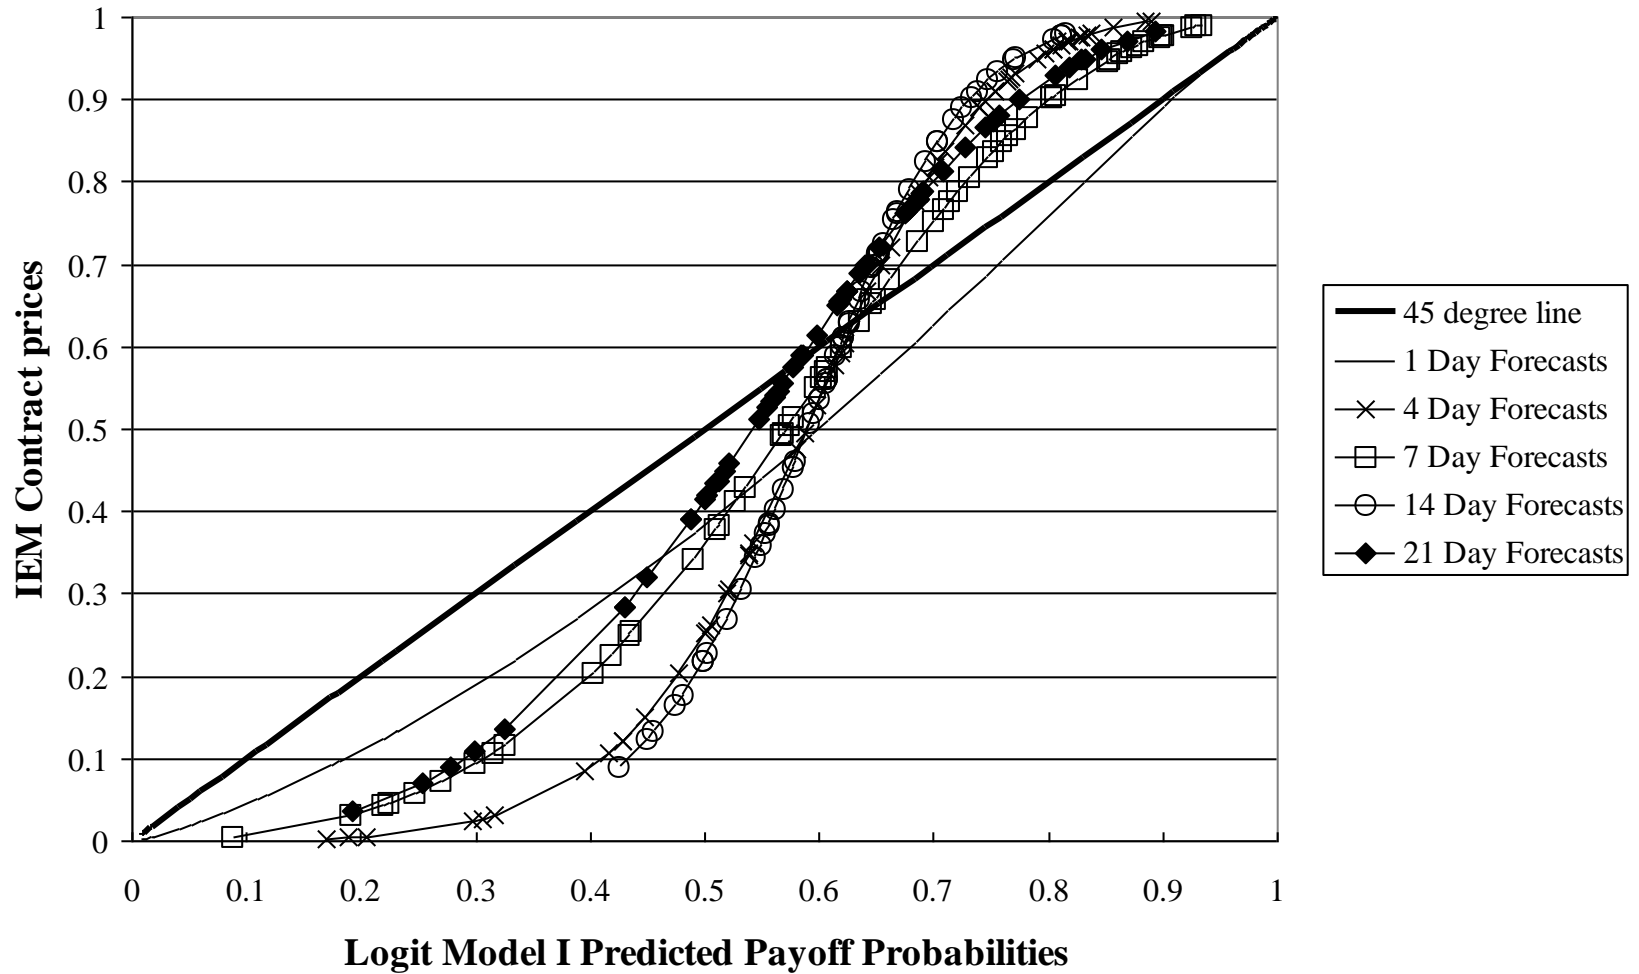


Figure 4: Logit Model I predicted payoff probabilities versus contract prices for the MSH contract in the Microsoft Price Level Market. Logit Model I uses the log price ratio as the only independent variable to explain the payoff probability.




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	\$(0.0712)	\$(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.