Do Behavioral Biases Affect Prices?

Joshua D. Coval and Tyler Shumway*

May 4, 2001

Abstract

This paper documents strong evidence of behavioral biases among Chicago Board of Trade proprietary traders and investigates the effect these biases have on prices. Our traders appear highly loss-averse. Traders who experience morning losses are about 16 percent more likely to assume above-average afternoon risk than traders with morning gains. This behavior has important short-term consequences for afternoon prices, as losing traders are prepared to purchase contracts at higher prices and sell contracts at lower prices than those that prevailed previously. However, during the ten minutes that follow these trades, prices revert strongly to their earlier levels. Consistent with these findings, short-term afternoon price volatility is positively related to the prevalence of morning losses among locals, but overall afternoon price volatility is not.

*Both authors are from the University of Michigan Business School. Joshua Coval can be reached at 734-764-2313 or *coval@umich.edu* and Tyler Shumway can be reached at 734-763-4129 or *shumway@umich.edu*.
This research is supported by the NTT fellowship of the Mitsui Life Center. We are grateful to Sugato Battacharyya, George Benston, Jonathan Berk, Hank Bessembinder, Jonathan Karpoff, Roni Michaely, Lubos Pastor, Matt Spiegel, Mark Taranto, Dick Thaler, Tuomo Vuolteenaho, Ivo Welch, George Wu, and seminar participants at Berkeley, Chicago, Cincinatti, Cornell, Duke, Emory, HBS, Michigan, North Carolina, Stanford, and Yale for helpful comments and suggestions. We also thank Steven J. Cho of the CFTC for helping us obtain the data used for this study.

A number of recent papers in the finance literature have proposed behavioral theories to account for asset pricing anomalies.[1] To provide support for their models' assumptions about investor behavior, these papers draw heavily from the experimental psychology literature, where evidence of cognitive biases is abundant. On the one hand, behavioralists contend that this evidence has been important in prompting researchers to consider heterodox explanations of market anomalies. On the other hand, skeptics argue that there exists so much of such evidence that behavioralists can "psycho-mine" the experimental psychology literature to find support for the particular set of assumptions that allow their models to match otherwise anomalous data. Contributing to the skeptics' argument, many of the behavioral theories rely on biases that are quite different from each other and often produce opposite conclusions about investor behavior. Not surprisingly, strong demand has emerged for empirical work that identifies which of the biases, if any, influence investor decisions. Even stronger is the demand to determine whether behavioral biases have any effect on prices. This paper supplies evidence on both of these issues.

Empirical tests of behavioral models face a number of challenges. First, the models cannot be easily tested with aggregate data. As noted by Campbell (2000), "[Behavioral models] cannot be tested using aggregate consumption or the market portfolio because rational utility-maximizing investors neither consumer aggregate consumption (some is accounted for by nonstandard investors) nor hold the market portfolio (instead they shift in and out of the stock market)." As a result, testing behavioral models is quite difficult without detailed information on the trading behavior of market participants. Unfortunately, given the issues of confidentiality associated with such data, its availability is generally quite low. An additional difficulty is that an investor's horizon, while highly ambiguous in most empirical settings, represents a key dimension to behavioral models. For instance, when fund managers are averse to losses, it is not clear whether their aversion relates to returns at the monthly, quarterly, or annual horizons, or even whether they view losses on positions taken recently as equivalent to losses on positions entered into years ago. Finally, to demonstrate that behavioral biases are important, empirical tests must be positioned to identify a link between biases in individual trader behavior and overall prices.[2]

[1] Examples include theories of overconfidence (Barberis et al.(1998), Daniel et al. (1998), and Odean (1998)) and loss aversion (Benartzi and Thaler (1995), Shumway (1998) and Barberis and Huang (2000)).

[2] Odean (1998), who studies transaction data of clients of a large discount brokerage and uncovers strong evidence of overconfidence, offers an important start in this direction. The empirical challenge is to identify

In this paper, we conduct a series of tests to determine the importance of behavioral biases in the price-setting process. Our tests focus on the trading behavior of market makers in the Treasury Bond futures contract at the Chicago Board of Trade (CBOT). This environment offers a number of advantages in assessing behavioral biases and any consequences they might have for prices. First, because we begin with every transaction made by the market makers in the T-Bond pit over a one-year period (over five million transactions), we have significant power to detect biases in trading behavior. Second, our traders are full-time proprietary traders, who are not trading to satisfy hedging needs, and whose livelihood depends entirely on their ability to trade effectively. To the extent that behavioral biases impair this ability, we can expect their trading behavior to provide a lower bound on the importance of behavioral biases for market participants in general. Also, because our traders are market makers, and do not trade through brokers or other intermediaries, they are far more proximate to the price-setting process. As a result, relative to other market participants, any impact their trading biases have on prices is likely to be more pronounced and therefore easier to detect.

The final, and perhaps most important, benefit of our focus on CBOT market makers is that the relevant horizon is quite clear. While in most settings the horizon over which investor performance is evaluated is ambiguous, CBOT market makers have clear incentives and mechanisms encouraging them to evaluate their performance on a daily basis. The traders receive and review statements at the end of each trading day detailing their performance during the day. Additionally, because most trades are unwound by the end of the day, and traders seldom retain significant positions overnight, all profits or losses during a day can be attributed to trades executed that particular day. Moreover, since the market makers' focus is on reading the order flow, which conveys highly short-lived signals regarding future trading activity, they carry little informational advantage from one day to the next. As a result, the statement CBOT market makers receive at the close of the trading day can be viewed as the perfect report card on their day at work.

Our study tests the null hypothesis of standard, rational investor behavior against a number of popular, though potentially competing, alternative behavioral hypotheses, including overconfidence, the house money effect, and loss-aversion. Specifically, we argue

evidence that their overconfident behavior impacts prices.

that if traders overconfidently interpret realized profits as signals of their ability (a conditional version of overconfidence termed "self-attribution bias"), or if traders are more willing to assume risk when gambling with the "house's money," they will take greater risks as their profits grow.[3] Conversely, if traders are averse to losses incurred at the daily horizon, this will lead to the opposite result: traders will take fewer risks as they become profitable. Thus, our setting allows us to study self-attribution bias and the house money effect on the one hand, and loss-aversion on the other, in an environment in which they yield opposite predictions regarding the relationship between realized profits and subsequent risk-taking.

To examine this relationship, we simply split the trading day into two periods and test whether traders with profitable mornings increase or reduce their afternoon risk-taking. We find strong evidence that CBOT traders are highly loss-averse: they are far more likely to take on additional afternoon risk following morning losses than morning gains. In our sample, a trader with morning losses has a 31.3 percent chance of taking above-average risk in the afternoon, compared to a trader who earns a profit in the morning who has only a 26.9 percent chance. Thus, a losing trader is 16 percent more likely to take above-average afternoon risk than a winning trader. This result shows up robustly across most of our tests, including pooled OLS regressions, panel regressions, and Fama-MacBeth style averages of trader-by-trader (time series) or day-by-day (cross-sectional) regression coefficients. The result is also robust to employing alternate measures of risk: losing traders are 28.6 percent more likely to place an above-average number of afternoon trades and are 9.5 percent more likely to trade at above-average sizes.

Next, to see whether the traders' loss aversion has an impact on prices, we examine whether our traders are more likely to move afternoon prices following morning losses. Specifically, we identify traders as "marginal" or "price-setting" if they purchase at a higher price or sell at a lower price than prevailed previously. For instance, if the previous trade took place at 25, we identify a given market maker as the marginal trader if he purchases at 26 or if he sells at 24. Our results clearly demonstrate that traders are more likely to place such price-moving trades following morning losses. A trader who loses money in the morning is around 15 percent more likely to execute such a trade than a trader who makes

[3]Several papers, including Langer and Roth (1975), document that people take credit for past success and attribute past failure to bad luck. Thaler and Johnson (1990) demonstrate that individual willingness to take risk increases following recent successful outcomes.

money in the morning. Overall, while traders lose money 32.9% of the time, losing traders account for 38% of all afternoon price-setting trades placed by market makers.

To gauge the quality of prices set by traders with morning losses and to assess how permanently they move prices, we monitor the average price change that follows a price-setting trade. If the marginal prices set by losing traders persist for a significant period of time, their loss-averse behavior may have permanent consequences for prices. In addition, if the prices do persist, it suggests that such trading is not so costly to the loss-averse traders – i.e. that they are able to "create their own space."[4] If, on the other hand, prices reverse strongly to previous levels, this raises doubts about the potential importance of loss aversion in influencing prices over the longer term. Moreover, the magnitude of the reversal in prices set by losing traders offers a measure of the costs associated with their loss-averse behavior. Our results indicate significant reversals in prices set by loss-averse traders. During the ten minutes following a price-setting trade, prices reverse 16.4 percent more if the trader experienced morning losses than if he experienced gains. This suggests that the price-setting trades of locals with morning losses have far less permanent an influence than the average price-setting local trade. This finding is highly supportive of arguments made by Friedman (1953) and Fama (1965) against the importance of noise traders in the price formation process.

Our final set of results examine whether prices exhibit greater volatility on afternoons that follow mornings when trader losses are widespread. This inquiry is closely related to the work of Shiller (1981, 1989), who attributes excess volatility in asset prices to patterns in human behavior. Consistent with the above results, our evidence suggests that loss aversion helps account for the short-term volatility of afternoon prices but cannot account for volatility measured over longer horizons. Following mornings during which overall losses are one-standard deviation larger than usual, expected afternoon volatility measured at the one-second frequency increases by 12 percent. As volatility is measured over longer periods, however, the effect of morning losses disappears. For instance, at the ten-minute horizon, the increase in expected volatility drops to 7 percent. However, our volatility results are not entirely conclusive, since we do not have a sufficiently long time series to explore the volatility hypothesis with much power.

[4] This relates to studies investigating the long-run survival of noise traders, such as De Long et al. (1990), who demonstrate that noise traders can create risk which is priced and prosper in assuming this risk.

The paper proceeds as follows. In Section I, we discuss a variety of behavioral biases and their implications in the daily horizon trade setting. In Section II, we outline our data and tests. Section III presents our results, and Section IV concludes.

I. Self-Attribution, House Money, and Loss-Aversion

Our study examines behavioral biases by focusing on the relationship between profits and subsequent risk-taking activity across the trading day. Before proceeding, it is important to understand how this relationship should look under various sets of assumptions. All of our tests begin with the null hypothesis that profits are not related to future risk-taking activity. This null will be valid in a setting where traders have standard Von Neumann-Morgenstern expected utility, profit opportunities are uncorrelated across the trading day, wealth effects are negligible, margin effects are unimportant, and traders are fully rational.

To understand the potential impact of overconfidence in our setting, it is worth reviewing the signals traders receive. As they trade, CBOT market makers interpret a variety of private signals related to the pit order flow. At any given point in time, they closely monitor brokers currently placing orders, including their identities, the size and direction of their orders, the eagerness with which they attempt to execute their orders, etc. How profitably market makers trade in response to these signals, by adjusting their quotes and managing their positions, depends on both their interpretation of the signals and on luck. To the extent that the signal quality is unknown, and varies across trading days, traders will update their assessment of the signal precision as market movements confirm or disconfirm their reading of the signals. If, however, traders have biased self-attribution, a trader that places profitable trades will become overconfident in the order flow signals. The trader will overly attribute the profits of his trades to his interpretation of the order flow signals and insufficiently attribute the profits to luck. If signal quality varies significantly from day-to-day, traders with self-attribution bias that trade profitably in the morning will become overconfident and will assume above average afternoon risk.

In this context, self-attribution bias is closely related to the "hot hand effect" documented by Gilovich, Valone, and Tversky (1985). Gilovich et al. document that basketball players believe they are far more likely to score following previous successful attempts than following previous misses. They further demonstrate that these beliefs are not justified by their subsequent rate of success – they find little autocorrelation in basketball player field goal percentages. If CBOT market makers become similarly overconfident in their trading ability following successful mornings, they will assume additional afternoon risk. Moreover, if such a relationship exists, the hot hand effect predicts that there will be little relation between morning and afternoon trading performance.

The "house money effect" documented by Thaler and Johnson (1990), also predicts a positive relation between morning profits and afternoon risk-taking. Investigating in an experimental setting the relation between prior outcomes and risk choice, they demonstrate that individuals are more risk-seeking following prior gains than following recent declines in wealth. In our setting, this suggests that traders who have earned profits in the morning will become less risk-averse because they feel they are "gambling with the house money." Thus, consistent with self-attribution bias and the hot hand effect, the house money prediction is that morning trading profits will be positively related to afternoon risk-taking.

A second, opposing alternative hypothesis is that of loss-aversion of the form proposed by Kahneman and Tversky (1979). They propose loss-averse utility functions that are convex over losses and concave over gains. If traders are averse to losses at the daily horizon, then traders that lose money in the morning will respond by assuming greater risk in the afternoon. Kahneman and Tversky (1979) characterize such behavior in the following terms, "[A] person who has not made peace with his losses is likely to accept gambles that would be unacceptible to him otherwise" (p. 287). Conversely, traders that earn profits in the morning will reduce their exposure to risk in the afternoon. These predictions closely resemble the findings of Camerer et al. (1997), who demonstrate that New York City cab drivers set daily targets and work until they reach their goal. This is a particular form of a framing effect, suggesting that the cab drivers frame their problem in terms of reaching a daily earnings goal rather than maximizing their profit. Thus loss aversion and framing yield exactly the opposite prediction from self-attribution bias, the hot hand effect, and the house money effect about the relation between morning profits and afternoon risk-taking.

It is important to note that a relation between morning profits and afternoon risk-taking may be influenced by wealth effects, correlation in profit opportunities, and margin constraints. If wealth effects are important, and traders have declining absolute risk aversion, traders that have become notably wealthier from morning trading will be become less risk-averse in the afternoon. Similarly, if profit opportunities are positively correlated across the trading day, traders will assume high risk on days when they expect to earn high average profits. Finally, if margin constraints are important, morning profits will influence the amount of risk traders are willing and/or allowed to take in the afternoon. Traders who find themselves near their margin constraint following a losing morning may be inclined to reduce their risk-taking to avoid a margin call. Conversely, traders who are below their margin constraint may be inclined to increase their afternoon risk-taking to get above the margin constraint before they are forced to liquidate.

To summarize, our null hypothesis is that afternoon risk will be unrelated to morning profits. If markets are efficient, traders are rational, traders have Von Neumann-Morgenstern utility functions, and wealth effects are negligible, margin constraints are unimportant, and profit opportunities are uncorrelated across the trading day, then we should expect no relationship between morning returns and afternoon risk-taking. Self-attribution bias, hot-hands effect, and the house-money effect all generate an alternative hypothesis on one side of the null: that morning returns will be *positively* related to afternoon risk-taking. Loss-aversion and framing predict the null will be rejected in the other direction: that morning returns will be *negatively* related to afternoon risk-taking.

II. Data and Method

Our primary data consist of the entire history of transactions (audit trail data) from the CBOT Treasury Bond futures pit during all of 1998.[5] The data include identifiers for the buying trader and the selling trader, the price, and the time for each transaction. They also include a code indicating whether each trade is performed on behalf of a customer, on behalf of the trader's clearing firm, on behalf of another trader, or for a trader's own account. Our data include records of over five million futures transactions, 97.4 percent of which involve front-month contracts, which are the focus of our tests. In 96.6 percent of the front-month futures transactions in our data, at least one of the two traders is trading on his own account. There are 1082 different traders in the data. Looking at how frequently

[5] The data was obtained from the CFTC via a Freedom of Information Act filing.

each trader trades for his own account, we identify 426 local traders. Each of our locals executes at least 1,500 trades for his own account, and trades bond futures on at least 100 days over the course of the year. We track each local's trades placed on their own account and the associated inventories and profits throughout each trading session.

Since our hypotheses relate the risk that a trader will take to his profitability, it is important for us to measure both profits and risks correctly. To measure each trader's profits and inventory, we assume that each trader closes out his positions at the end of each day, and thus begins each day with no position. This assumption is supported by the evidence of Kuserk and Locke (1993) and Manaster and Mann (1996) and has been used previously in Manaster and Mann (1996) and Coval and Shumway (2000). Of course, some noise will be present in our calculations to the extent that traders close out positions during evening trading sessions, hedge their positions using options contracts, or place their trades through other locals. Assuming no beginning inventory makes profit calculations simple. We multiply the difference between purchase and sales prices by quantities to arrive at a profit figure for each local at each point in time. Since our hypotheses require us to have a profit figure available at a particular time each day, we add the market value of any inventory, calculated as the current price times the contracts outstanding, and add this to each local's running profit figure to generate a total profit variable for any time of the day.

Measuring the risk each trader takes is less straightforward. Certainly the number of trades a trader places and the average size of these trades will be related to risk he assumes. However, since the level of risk in the T-Bond contract is non-constant across the trading day, estimating each trader's risk requires an estimate of the risk a given position exposes the trader to at different points during a particular day. Therefore, we use historical price change data to model the level of risk throughout the trading day. Using second-by-second price data (time and sales data) from the Futures Industry Institute Data Center, we calculate the front-month futures contract price at the beginning of each minute of each day from 1989 to 1998. These prices are used to calculate the absolute price change from one minute to the next.

To measure the risk a given position faces across the trading day, we employ an ordered logit regression (as in Coval and Shumway (2000)). A logit function of the probability of various potential absolute price changes over the next minute is regressed on the magnitude of price changes in the preceding five minutes and time-of-day dummy variables for each five-minute period during the trading day. The fitted values from this regression are then used to construct an expected absolute price change for each minute of each full trading day in 1998. Since our risk measure is an expected absolute price change, it roughly corresponds to a one-standard deviation measure of price change risk associated with each one-minute interval.

Finally, a trader's risk is calculated by multiplying each minute's risk measure by the trader's position at the beginning of the minute, and adjusting the trader's risk for the minute by any changes in inventory, and therefore risk, that occur during the minute. Again, our measure is roughly the standard deviation of wealth the trader assumed during a given minute. We then can calculate the cumulative risk a trader has assumed up to a given point each day by summing the measure of risk across all of the previous minutes. We term this risk measure the "total dollar risk." Although we view this to be the proper way to measure trader-specific risk, we verify that our results are robust to employing alternative risk measures, such as number of trades and average trade size.

III. The Evidence

This section details the evidence we obtain on each of our hypotheses. The first hypothesis that we examine in detail is that locals at the CBOT exhibit behavioral biases. In particular, we examine the relationship between each trader's profits in the morning and the risk that he takes in the afternoon. If profit opportunities are uncorrelated across the trading day and wealth effects are negligible, any relationship between morning profits and afternoon risk taking indicates that traders exhibit behavioral biases. The second hypothesis that we examine relates each trader's performance in the morning to his probability of setting prices in the afternoon. In particular, we examine whether traders with morning losses tend to be buying or selling when the price moves up or down. Our third hypothesis concerns the permanence of the prices set by traders with losses. We estimate the expected price ten minutes after a price change, conditional on whether the trader that moved the prices had losses or gains in the morning. Our final hypothesis relates aggregate morning losses to afternoon volatility. We test whether mornings with widespread losses lead to more volatile afternoons.

A. Summary Statistics

To examine whether CBOT locals exhibit loss aversion, we look at the relation between morning trading performance and afternoon risk taking. Since the trading day at the CBOT begins at 7:20 a.m. and ends at 2:00 p.m., we split the trading day into a morning period before 11:00 a.m., and an afternoon period after 11:00 a.m. We round the midway point of the trading day (10:40 a.m.) to 11 a.m. because it is somewhat closer to the midway point of the average trader's lunch break. Then, for each trader, we calculate morning and afternoon profits, and we calculate morning and afternoon values for each of the three risk measures: total dollar risk, number of trades, and average trade size. With morning and afternoon profits and risks defined, we are almost ready to test our hypotheses.

We can examine our first hypothesis simply by relating the risk a trader takes in the afternoon to the trader's profit or loss in the morning. However, because traders face margin constraints, results from simple regressions may be misleading. If traders who experience large morning losses face binding margin constraints, they may be forced to liquidate their holdings and assume very little risk in the afternoon. Alternatively, if traders lose enough to trigger constraints but their trading is not immediately restricted, they may take an inordinate amount of risk in hopes that they can win back enough to avoid margin calls.

To control for trader heterogeneity with respect to margin constraints and risk tolerance in general, we normalize morning profits and afternoon risk. We first calculate the standard deviation of each trader's morning profits and then divide each trader's morning profit observations by his profit standard deviation to calculate the trader's normalized profit. We denote our measure of the normalized morning profits of trader i on date t as $\pi^M_{i,t}$. We perform a similar calculation to normalize each trader's afternoon risk. We calculate trader-specific means and standard deviations of each of our afternoon risk measures, then demean each trader's daily afternoon risk and divide by the trader-specific standard deviation. We denote trader i's normalized measure of afternoon risk on date t as $\text{Risk}^A_{i,t}$, where $\text{Risk}^A_{i,t}$ may be used to reflect trader i's total dollar risk, total number of trades, or average trade size on date t. In this way, our profit measure and our three risk measures have standard deviations of one for each trader.[6]

In Table I we report the mean and standard deviation of morning and afternoon measures of profits, number of trades, average trade size, and total dollar risk. Several points emerge from Table I that are worth noting. First, our panel of local trading days contains 82,595 observations. In 67 percent of the observations, the given local traded profitably during the morning. This is supported by the overall average of locals' standardized morning profit, which at 0.137 is significantly positive. Since this number is scaled by the trader-specific standard deviation, it means that each trader earns on average 14 percent of one standard deviation in his morning returns per morning. Afternoon profits are also positive on average, but at 0.077 are slightly lower.[7]

Next, we turn to our three measures of risk. Since we have demeaned our risk measures by their trader-specific averages, the overall averages are equal to zero by construction. However, when we take averages of observations associated with profitable and losing mornings, notice that the afternoon risk measures are far higher following losing mornings than following profitable mornings. Traders with losing mornings place more trades (0.124 vs. -0.066), place trades with larger average size (0.086 vs. -0.046), and assume greater total dollar risk (0.204 vs. -0.098) than those with profitable mornings.

However, we can also see from Table I that traders with losing mornings are not otherwise equivalent to traders with profitable mornings. They enter the afternoon having assumed far greater morning risk, on average, and they enter the afternoon with significantly larger outstanding inventories. For instance, traders with losing mornings have placed 6.6% of one standard deviation more trades during the morning than average. Also, traders with losing mornings begin the afternoon with 21.7% of one standard deviation larger outstanding inventories than average. Thus, it is quite important that we control for these factors in a regression setting.[8]

Looking at afternoon returns, we see that they are only slightly lower following losing

[6]While these normalizations allow for a more sensible economic interpretation, they are not necessary for our main set of results.

[7]The standardized morning and afternoon profit measures correspond to average dollar profits of $1599 and $871, respectively.

[8]Of course, it is possible that the morning measures of inventory and risk are large for traders with morning losses because they have already become risk-seeking. Consistent with this, our results are far stronger when when the morning risk and inventory controls are omitted.

mornings than following profitable mornings (0.078 vs. 0.076). This suggests that the additional afternoon risk traders assume following losing mornings is not terribly costly from an expected return standpoint. Moreover, to the extent that traders seek to increase the spread in their afternoon returns following morning losses (e.g. due to loss aversion), we see that they are successful in achieving this objective, as the standard deviation of their overall afternoon return is significantly larger following losing mornings than following winning mornings (1.22 vs. 0.88). Now we turn to the regression setting to see whether these results are significant and robust to controlling for other factors.

B. Morning Losses Lead to Afternoon Risk-taking

In Table II, we present results of regressions of afternoon risk-taking on morning profits. Included in the regression is the absolute value of each trader's outstanding morning (11:00 a.m.) inventory, demeaned, and normalized by each trader's standard deviation of outstanding morning inventory. We include normalized inventory for three reasons. First, as noted above, traders with losing mornings tend to have larger outstanding positions heading into the afternoon, and we would like to control for the additional afternoon risk introduced by this position. Second, if traders do not begin each day with zero inventory, including each trader's absolute inventory may attenuate the bias that results in our measurement of profit and risk. Finally, in order to account for the possibility that traders unwind losing and winning positions in different ways,[9] we include a term interacting morning profits and morning inventory. Specifically, our regression takes the following form:

$$\mathrm{RISK}^{A}_{i,t} = \alpha + \beta_{\pi}\pi^{M}_{i,t} + \beta_{I}|\mathrm{INV}^{M}_{i,t}| + \beta_{\pi I}\pi^{M}_{i,t}\cdot|\mathrm{INV}^{M}_{i,t}| + \beta_{R}\mathrm{RISK}^{M}_{i,t} + \varepsilon_{i,t}, \tag{1}$$

where $\mathrm{RISK}^{A}_{i,t}$ is one of the three normalized afternoon measures of risk for trader i on date t, $\pi^{M}_{i,t}$ is trader i's date t morning profit, $|\mathrm{INV}^{M}_{i,t}|$ is the absolute value of trader i's outstanding position (measured in thousands of contracts) at the end of the morning on date t, $\mathrm{RISK}^{M}_{i,t}$ is trader i's morning risk measure on date t, and $\varepsilon_{i,t}$ is the error term.

We estimate our regression in a variety of ways. First, we estimate a simple pooled-OLS regression. We also conduct Fama-MacBeth style regressions in which we conduct trader-

[9]This possibility is suggested by the findings of Shefrin and Statman (1985), Odean (1998), and Locke and Mann (1999).

by-trader regressions and average the coefficients across traders, and we conduct day-by-day regressions and average the coefficients across days. The Fama-MacBeth regressions serve two purposes. First, they check whether our results are driven by cross-sectional or time-series correlation in residuals. Also, they test whether our results are driven more by particular traders or by particular days. We also conduct the panel regression with fixed effects for both traders and days and panel-corrected standard errors.[10] In Panel A, we report the results of these regressions with the number of afternoon trades as the measure of afternoon risk-taking. Panel B reports results of regressions using average trade size as the dependent variable. Finally, in Panel C, we document the results using total dollar risk.

As we can see, consistent with the results presented in Table I, our regressions indicate that traders are significantly loss-averse. The results are highly significant across the most of the different specifications. The regressions indicate that a one-standard deviation decrease in morning profits leads the average trader to place 11% to 23% more afternoon trades than normal (Panel A), place afternoon trades which are 5% to 11% larger than normal (Panel B), and assume total dollar risk which is up to 1.5% larger than normal (Panel C). The significance of morning profits in explaining afternoon total dollar risk (Panel C) is only significant in one of the four specifications (the Fama-MacBeth average of daily cross-sectional regressions) and is of lower economic significance than that using other risk measures. The regressions using average afternoon trade size (Panel B) include somewhat fewer observations than the others because they only include traders who place at least one afternoon trade. The fact that traders with profitable mornings place afternoon trades which are of smaller size than average suggests the results are not entirely driven by a "framing effect" similar to the taxi cab findings of Camerer et al. (1997). Traders with profitable mornings are not only more likely to stop trading in the afternoon – those that remain tend to trade less aggressively (in lower sizes) than normal.

As expected, the inventory terms are highly significant, indicating that traders with large midday positions assume additional afternoon risk as they unwind them (or that losing traders are already expanding their positions in order to assume greater afternoon risk). The morning risk variables come in highly significant as well, indicating that traders who assume significant morning risk tend to continue to do so in the afternoon. To make

[10] Our panel-corrected standard error estimates adjust for contemporaneous correlation and heteroskedasticity across traders.

sure our results are not driven by outliers or by the behavior of traders facing margin constraints, we rerun our regressions on the subset of traders whose morning profits did not deviate from zero by more than two standard deviations. Under this specification, which we do not report due to space considerations, the results are considerably stronger in economic and statistical terms. For instance, the regressions using total dollar risk all become highly significant when the outliers are removed.

As an additional robustness check of our results, in Table III we conduct logit regressions to see whether a trader's likelihood of assuming greater-than-average afternoon risk depends on whether the trader incurred morning losses. Specifically, we estimate the logit model defined by

$$\text{Prob}(\text{RISK}^{A}_{i,t} > 0) = \frac{\exp X'\beta}{1 + \exp X'\beta}, \tag{2}$$

where

$$X'\beta = \alpha - \beta_{\pi} I(\pi^{M}_{i,t} < 0) + \beta_{I}|\text{INV}^{M}_{i,t}| + \beta_{\pi I} I(\pi^{M}_{i,t} < 0) \cdot |\text{INV}^{M}_{i,t}| + \beta_{R}\text{RISK}^{M}_{i,t},$$

and where $I(\pi^{M}_{i,t} < 0)$ is an indicator variable which is equal to one if trader i's morning profit on date t is negative. Note that the term $-\beta_{\pi} I(\pi^{M}_{i,t} < 0)$ enters the regression equation with a negative sign in order to make the expected sign of the coefficients negative under loss aversion.

Again, the morning losses enter significantly in almost all regressions. The fact that the Fama-MacBeth regression by date yields consistently stronger estimates than the regression by trader (as was the case in Table II), implies that our results are being driven more by the profits/risk-taking relation across traders on particular days than by the relation across days for particular traders. Our binary results also offer an alternate estimate of the economic significance of our results. Using the pooled logit regressions of Table III, and evaluating them with dependent variables at their means, traders that lose money in the morning increase their probability of assuming above-average afternoon risk from 26.9 to 31.3 percent. This represents an increase in likelihood of slightly more than 16 percent. Overall, Tables II and III make a strong case that market makers at the CBOT behave in a loss-averse manner.

C. Semi-Parametric Regressions

In their calibrations of loss-averse utility functions, Kahneman and Tversky (1979) estimate that individual utility functions have a kink at zero and convexity over losses that is roughly equal to their concavity over gains. To investigate whether the afternoon responses of traders with morning losses aresymmetric to those with morning gains, we conduct a series of semi-parametric regressions that permit a non-linear relationship between morning returns and afternoon risk-taking. Specifically, we rank traders each day according to their normalized morning profit and assign them to one of twenty profitability groups.[11] We then conduct daily cross-sectional regressions of the following form:

$$\mathrm{RISK}^{A}_{i,t} = \alpha + \sum_{j=1}^{20} \beta_{\pi,j} D_{i,j,t} + \beta_I |\mathrm{INV}^{M}_{i,t}| + \sum_{j=1}^{20} \beta_{\pi I,j} D_{i,j,t} \cdot |\mathrm{INV}^{M}_{i,t}| + \beta_R \mathrm{RISK}^{M}_{i,t} + \varepsilon_{i,t}, \quad (3)$$

where $D_{i,j,t}$ is a dummy variable which is equal to one if trader i's morning profit ranks in group j on date t. We then average the cross-sectional regression coefficients across time and calculate the corresponding standard errors. Figure I plots the average of the morning profit regression coefficients for each of the twenty profit percentile groups when total afternoon dollar risk is used as the dependent variable. A kernel-smoothed line is plotted across these coefficients and two-standard error bands are included to reflect significance.

The figure highlights a significant asymmetry between the responses of traders with profitable mornings and those with losing mornings. Consistent with the above results, traders with losing mornings increase their afternoon risk-taking significantly. Moreover, the smaller the morning losses, the smaller the increase in afternoon risk-taking. The relationship holds up to the until around the 30th percentile of morning profitability. Given that this is the point in the ranking where the traders earn zero profits, it is highly consistent with the findings of Kahneman and Tversky (1979), who estimate that utility functions shift at zero from convex to concave with a kink.

As we move into the positive range of trader morning profitability, the picture changes. Traders with profitable mornings all take on relatively similar, below-average, levels of afternoon risk. Only traders who experience extremely high profits differ at all in their

[11] To ensure the zero-profit level does not move around across the year, we create the profitability bins using the entire year of morning profit observations.

afternoon risk-taking. Traders in the final few percentiles exhibit a slight increase in their afternoon risk-taking, though the increase is not statistically significant, and is economically much smaller than the increased risk-taking of traders with losing mornings. Thus, it appears that the relationship between morning profitability and afternoon risk-taking is not a symmetric one, with losing trader behavior more sensitive to the level of their losses than winning traders to the level of their gains.

D. Profits and Risk-Taking across Days

As we have argued, there are compelling institutional and behavioral factors which justify a one-day horizon for our traders. However, it is likely that all traders do not exclusively evaluate their profits at the daily horizon and that other horizons are important. One of the important aspects of testing at a daily horizon in our setting is that, because traders seldom enter the trading day with outstanding positions, our traders can attribute their performance during the morning entirely to trades executed that morning. Although this possibility disappears if we test at an hourly horizon, it does not if we move to the multi-day setting – profits earned during a particular day can be attributed entirely to decisions made that day. Thus, to see whether our findings are exclusive to the one-day horizon, or whether they are detectable at lower frequencies, we examine the relationship between profits and risk-taking across trading days.

To examine profits and risk-taking across days, we compare overlapping pairs of trader-days. Specifically, we ask whether profits on one day explain a trader's level of risk-taking the next. An attractive feature of the multi-day setting is that, unlike the morning-afternoon tests, we do not need to worry about a trader having an outstanding position following a losing day which influences our measurement of his risk-taking activity on the following day. Since we assume that traders hold no position overnight, inventory is always zero at the beginning of each day and traders must enter trade to incur risk. We estimate the regressions employed above without the inventory controls, simply regressing a trader's level of daily risk on his previous day's profit and previous day's risk. As in the earlier tests, we again normalize the profit and risk measures by traders, though we now use daily averages and standard deviations.

Daily regressions of risk-taking on profits corresponding to those presented in Tables II and III were conducted. When we use continuous measures of risk and profits (as in Table II), no detectable relationship exists between profit and risk across trading days, whether we run pooled OLS, panel, or Fama-MacBeth style regressions, and regardless of whether we remove outliers. To conserve space, we do not report the results of these regressions. In Table IV, we report the results using logit regressions (as in Table III) and risk measured by total dollar risk. As we can see, the results are consistently insignificant (similarly insignificant results obtain using other risk measures). This suggests that horizon effects can be quite important in identifying loss-averse behavior and that, for our set of traders, loss-aversion is only pronounced at the daily horizon.

E. Morning Profits and Afternoon Price Leadership

Having documented strong evidence of loss aversion among our traders, we now turn to the question of whether this loss aversion matters for prices. We begin by identifying trades placed by locals that move the price in a direction consistent with their trade. In a futures pit, traders do not post bid and ask prices as market makers do on an exchange floor. Rather, a group of traders stands ready to buy at a particular price, and a group stands ready to sell at a different, higher price. When large buy orders arrive from customers outside the pit, the orders generally are filled at the higher price. Large sales orders similarly go through at the lower price. The posted futures price therefore oscillates between the effective bid and ask price throughout the day. We identify trades that cause the posted price to change from bid to ask (or from ask to bid) because of the purchase or sale of a local trader for his own account. Specifically, we compare the price of a given local trade to the price of the previous trade. If the local purchases at a price that was higher than the previous price, we identify the trade as responsible for having raised the price. Likewise, if the local sells at a price that is lower than the previous price, we identify the trade as responsible for lowering the price. Although the actual bid and ask prices at a given point in time are not recorded by the CBOT, under conditions when they are well-defined for market participants, price-setting trades will resemble market orders. Locals do not execute price-setting trades very frequently. In order to be certain that an identified price move is caused by local, we drop all trades that occur during the same second. Across all locals and days, the average number of price-moving trades per afternoon is 0.218 per trader, for a total of around 93 price-moving trades executed by locals on a given afternoon.

Having identified trades that have moved the price, we then ask whether traders place more price-moving trades following losing mornings than following profitable mornings. Specifically, we regress the number of price-setting trades placed by a trader on a given afternoon (relative to his average) on the trader's morning returns and his morning inventory. Our regressions take the following form:

$$\Delta_{i,t}^{A} - \bar{\Delta}_{i}^{A} = \alpha + \beta_{\pi}\pi_{i,t}^{M} + \beta_{I}|\mathrm{INV}_{i,t}^{M}| + \beta_{\pi I}\pi_{i,t}^{M} \cdot |\mathrm{INV}_{i,t}^{M}| + \varepsilon_{t}^{i}, \tag{4}$$

where $\Delta_{i,t}^{A}$ is the number of price-setting trades placed in the afternoon of day t by trader i, $\bar{\Delta}_{i}^{A}$ is trader i's average number of price-setting trades per afternoon, and other variables are as defined earlier.

The results of pooled OLS, fixed effects, and Fama-MacBeth style cross-sectional and time-series regressions are reported in Table V. In all specifications the results are highly significant. A trader who experiences a one standard deviation loss in the morning is between two and five percent more likely to place a price-setting afternoon trade than he is on an average afternoon. Since traders only place price-setting trades approximately 22 percent of the time, even a two percent increase in likelihood is economically significant. Experiencing morning losses makes a trader between 10 and 20 percent more likely to be a price setter than an equivalent trader with morning gains. Again, the results are highly significant across all the regression specifications, and appear strong both in the time-series and the cross-section. The results also obtain if binary specifications of morning profit and afternoon risk-taking replace the continuous variables used in Table V, or if an ordered logit model with trader-specific fixed effects replaces the demeaned regression specification reported. The results are markedly stronger when traders with large morning losses or gains are removed from the sample.

Consistent with the earlier results on average afternoon trade size, these results suggest that traders who incur morning losses are not passively assuming more afternoon risk. To obtain the additional risk, they appear to be frequently hitting existing limit orders at prices that are less favorable than those of previous trades. This implies that traders with morning losses cannot easily assume the additional afternoon risk they desire and must give up an "edge" to obtain the additional exposure. Although one might suspect that losing traders are only giving up an edge to unwind existing positions and not to open new positions, the results remain highly significant if we focus only prices set by inventory expanding trades and ignore trades that are closing out existing positions.

While a subset of the afternoon trades of traders with losing mornings appear to move prices, if they are motivated by loss aversion, the price impact of such trades should be less permanent than that of trades which are information-based. Moreover, given that these traders are likely to unwind their positions as the day progresses, we should expect the price impact of position-initiating trades to disappear as the day moves forward. We investigate these issues in the following section.

F. Loss Aversion and Price Permanence

The price-setting trades we identify in the previous section appear to be motivated by loss-averse traders eager to assume additional afternoon risk to improve their odds of recovering morning losses. If this is indeed the case, and the trades are not based on information about the fundamental value of the futures contract, we should expect them to have a more transitory impact on prices than trades based on information. Moreover, an examination of the quality of price-setting trades placed by loss-averse traders will give us a further idea of their afternoon trade performance in relation to their morning profitability. To pursue this, we compare the price permanence of price-setting trades placed by traders with morning losses to those placed by traders with morning profits. If afternoon trades placed by traders with profitable mornings are more likely to be informed trades than those placed by traders with losing mornings, we should expect the prices set by profitable traders to be far more permanent than those set by traders trying to recover their losses.

To examine price permanence, we regress the price change which takes place during the ten minutes following the price-setting trade on the price change caused by the trade, the previous two price changes, and the interaction of all three price changes with a dummy variable indicating whether the trader experienced a morning loss. As a robustness check, we add the interaction of price changes with the time between each price change to the specification. The regression takes the following form:

$$P_{j+n} - P_j = \alpha + \sum_{l=1}^{3} \Big\{ [P_{j-l+1} - P_{j-l}] \cdot \Big[\beta_{P,l} + \beta_{M,l} I(\pi_{i,t}^{M} < 0) + \beta_{T,l}(\tau_{j-l+1} - \tau_{j-l}) \Big] \Big\} + \varepsilon_j, \tag{5}$$

where P_{j+n} is the first price that is different from the previous transaction price at least ten minutes after the transaction at price P_j, and where $\tau_j - \tau_{j-1}$ is the number of minutes that transpire between the first transaction at price P_j and the first transaction at P_{j-1}. The results appear in Table V.

Table V actually reports estimated coefficients from three different regression models. Model 1 simply relates the price change over ten minutes to the current and two previous price changes by constraining $\beta_{M,j}$ and $\beta_{T,j}$ to equal zero for $j = 1, 2, 3$. The regression coefficient for the current price change is remarkably close to negative one (-0.9540) and the coefficient for the previous price change is very close to negative one half (-0.5137). Given the nature of trading in the CBOT futures pit, these coefficients are very economically sensible. They capture the short-term mean reversion of second-by-second futures prices.

As mentioned previously, most of the price movement in the futures pit can be considered bid-ask bounce. The effective bid-ask spread of the CBOT T-Bond futures pit is generally less than one price tick. This implies that when the price moves, it almost always does so by just one tick. Large purchases move the price up one tick while large sales move it down a tick. Therefore, short-run prices are extremely mean reverting. It can be helpful to illustrate with an example. One common price series might be 24, 25, 24, 25, 24, 25, where 24 and 25 represent thirty-seconds of a percentage of the face value of a $100,000 bond. Clearly the true price is somewhere between 24 and 25 in the preceding example. At the end of this price series, the expected price change over the next ten minutes is $-0.9540(1) + -0.5137(-1) + -0.1888(1) = -0.63$, implying an expected price of 24.37. The point of this calculation is simply to illustrate that the negative coefficient estimates for $\beta_{P,j},\ j = 1, 2, 3$, represent predictable, extremely short-term mean reversion.

Examining the coefficient estimates of Models 2 and 3, we see that the price-setting trades of traders with morning losses are reversed much more forcefully than the price-setting trades of traders without losses. While the estimates of $\beta_{M,1}$ are relatively small and statistically only marginally significant in the two models, the estimates of $\beta_{M,2}$ are large and significant in both models. This has intuitive appeal. Because the price change over the next ten minutes is driven almost exclusively by bid-ask bounce, conditioning on information beyond the current price change does not improve our forecast. However, conditioning on the previous price change adds a great deal of explanatory power to our forecast because the previous price change indicates whether the current price change is part of a short-term trend or simply bid-ask bounce.

To see the importance of the penultimate price change, compare the price series 24, 25, 24, 25, 24, 25, to 23, 24, 23, 24, 25, 26. Note that the only difference in price changes is that the penultimate price change is negative in the first case (24 - 25) and positive in the second case (25 - 24). Thus, the price change that immediately precedes the current price change differentiates between true movements in the price and a change from bid to ask or from ask to bid price. Given the coefficients in Model 2 of Table VI, if a trader that made money in the morning is responsible for moving the price up in our second example price series, the expected price in ten minutes reverts back to 24.45. However, if a trader that lost money in the morning is responsible for moving the price, the expected price in ten minutes is just 24.20. Similar numbers can be calculated for a price decline.

The results of Table VI make it clear that the prices set by traders with losses in the morning are reversed much more dramatically than those set by traders with gains in the morning. This implies two important inferences. First, because the trades of losing traders have only temporary price impact, other traders in the pit appear to regard them as "noise" trades, and appear comfortable trading aggressively against them. Any impact of the traders' behavioral biases on prices appears to be eliminated rapidly by other market participants. Second, because the prices set by losing traders are reversed so dramatically, trading to make up morning losses appears to be costly. Thus, Table VI provides strong evidence confirming that loss aversion is driving the behavior documented in Tables II and III.

G. Aggregate Morning Losses and Afternoon Price Volatility

In our final set of tests, we ask whether the price-setting trades executed by locals with morning losses cause afternoon prices to be more volatile than they would be if locals had no behavioral biases. Our measure of price volatility is the standard deviation of price changes measured at one-second, one-minute, five-minute, ten-minute, and half-day frequencies. Similar to our other regressions, we demean each measure of price volatility and normalize it by its standard deviations. To investigate our volatility hypothesis, we regress normalized afternoon volatility on the volatility in the corresponding morning and several measures of the prevalence of morning losses among local traders. Specifically, our regressions are as follows:

$$\sigma^A_{h,t} = \alpha + \beta_\sigma \sigma^M_{h,t} + \beta_\Pi \Pi^M_t + \varepsilon_t, \tag{6}$$

where $\sigma^A_{h,t}$ measures the abnormal volatility of afternoon price changes on date t measured at frequency h, $\sigma^A_{h,t}$ measures the abnormal volatility of morning price changes on date t measured at frequency h, and Π^M_t measures aggregate morning losses on day t.

Measuring aggregate morning losses is not a simple task because it is not clear how losses should aggregate. Therefore, we pick three different ways to aggregate losses. First, we simply calculate the fraction of locals with losses at 11:00 a.m. Second, we calculate the average of $\pi^M_{i,t}$ across traders each day. Finally, to measure the prevalence of losses among traders who are particularly loss-averse, we sum the product of an indicator variable that is one when morning profits are positive, $I(\pi^M_{i,t} > 0)$, and trader i's β_π coefficient in the OLS regression,

$$I(\mathrm{RISK}^A_{i,t} > 0) = \alpha - \beta_\pi I(\pi^M_{i,t} < 0) + \beta_I |\mathrm{INV}^M_{i,t}| + \beta_{\pi I} \pi^M_{i,t} \cdot |\mathrm{INV}^M_{i,t}| + \beta_R \mathrm{RISK}^M_{i,t} + \varepsilon_{i,t}, \tag{7}$$

where the measure of risk employed, $\mathrm{RISK}_{i,t}$, is the total dollar risk. Since our regression relates aggregate risk to aggregate losses, the sample size is limited to the 236 trading days in our data. Also, because serial correlation may be a problem in our sample, all the estimates we report are adjusted to account for first-order autocorrelation. The results are reported in Table VII for each of the loss measures and at a number of frequencies.

Table VII contains evidence suggesting that afternoon price volatility is related to morning market maker profitability. Consider first the coefficient on the fraction of losses (times negative one) and on normalized morning profits ($\pi^M_{i,t}$). We see that while they are both of the correct sign – indicating that mornings with a high fraction of locals experiencing losses and mornings with high average trader losses precede afternoons that are more volatile than usual – neither of the coefficients is statistically significant. The final measure of morning losses, the fraction of traders with morning profits weighted by the trader-specific measure of loss aversion, yields stronger results. At the one-second frequency, the coefficients are negative and significant in statistical and economic terms. A one-standard deviation decrease in the weighted average of morning profitability (0.0048) leads to a 12.2 percent increase in expected afternoon second-by-second volatility. This is consistent with the results of Table V, as traders with morning losses place additional price-setting afternoon trades to assume additional risk. As we move to the one and five-minute frequencies, the statistical and economic significance of the results declines slightly. When volatility is measured over ten minutes or over the entire afternoon, the results lose much of their economic significance and, as a result, all statistical significance. A one-standard deviation decrease in the weighted average of morning profitability now leads to only a 5.2 percent increase in overall afternoon volatility. Thus, although loss-averse traders appear to have a short-term influence on prices, consistent with the results of Table VI, their influence largely disappears during the ten minutes following their trades. However, it is important to note that while there appears to be a relationship between morning losses and afternoon volatility that is consistent with our earlier findings, the results are far from conclusive. Since our tests employ only a single observation per day, the power of our regressions is somewhat limited.

IV. Conclusion

Although behavioral finance has recently become an extremely popular area of asset pricing research, relatively little empirical evidence exists in direct support of behavioral theories. This is due, in part, to the fact that behavioral models cannot be tested as easily as traditional asset pricing models. Because aggregate consumption data or market returns data reflects the decisions of both rational and behaviorally biased traders, the standard tests of restrictions imposed by the Euler equations of rational, utility-maximizing agents are inapplicable. Proper assessment of behavioral theories require detailed information on the trading strategies of various market participants, and, until recently, such information has been difficult to come by.

This paper offers a detailed look at the trading behavior of a set of professional market makers and directly tests both for biases in their behavior and the consequences such biases may have for prices. Our traders are highly proximate to the price-setting process and they generally close out their positions by the end of each trading day, making the horizon over which they evaluate their performance quite clear. These factors provide us with significant power to identify conditions under which behavioral biases are likely to be important in influencing prices.

We find strong evidence that our traders are highly loss-averse. They assume significantly more afternoon risk following morning losses than following morning gains. In their eagerness to assume greater afternoon risk, they place price-setting trades more frequently, purchasing contracts at higher prices, and selling contracts at lower prices. However, afternoon prices set by traders with morning losses reverse substantially more than those set by traders with morning gains. This suggests that any price impact resulting from the traders' behavioral biases dissipates extremely quickly. Consistent with this, we find that mornings with widespread losses lead to increases in short-run afternoon volatility but no increase in volatility measured over longer intervals.

Unfortunately, because of the nature of the data, market, and trader horizons, most of our power to detect effects on prices is concentrated at the microstructure frequency. However, considering that this is one of the most competitive markets in the world, it is reasonable to expect our results to apply at lower frequencies in markets where trader horizons are longer and prices are set less efficiently. Future work should investigate the extent to which this is the case.

References

Barberis, Nicholas C. and Ming Huang, 2000, Mental accounting, loss aversion, and individual stock returns, Working paper, University of Chicago.

Barberis, Nicholas C., Andrei Shleifer, and Robert W. Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 309-343.

Benartzi, Shlomo, and Richard H. Thaler, 1995, Myopic Loss Aversion and the Equity Premium Puzzle, *Quarterly Journal of Economics*, Vol. 110, 73-92.

Camerer C, Babcock L, Loewenstein G, Thaler R, 1997, Labor supply of New York City cabdrivers: One day at a time, *Quarterly Journal of Economics*, 112, 407-441.

Coval, Joshua D., and Tyler Shumway, 2001, Is sound just noise?, *Journal of Finance*, forthcoming.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and overreactions, *Journal of Finance* 53, 1839-1885.

De Long, J.B., A. Shleifer, L.H. Summers, and R.J. Waldmann, 1990, Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98, 703-738.

Fama, Eugene F., 1965, The Behavior of Stock Market Prices, *Journal of Business* 38, 34-105.

Friedman, Milton, 1953, The Case for Flexible Exchange Rates, in *Essays in Positive Economics*, Chicago: University of Chicago Press.

Gilovich, T., R. Vallone, and A. Tversky, 1985, The Hot Hand in Basketball: On the Misperception of Random Sequences, *Cognitive Psychology*, 17, 295-314.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 236-291.

Kuserk, G.J., and P.R. Locke, 1993, Scalper Behavior in Futures Markets - and Empirical-Examination, *Journal of Futures Markets*, 13, 409-431.

Langer, Ellen J., and Jane Roth, 1975, Heads I win tails it's chance: The illusion of control as a function of the sequence of outcomes in a purely chance task, *Journal of Personality and Social Psychology*, 32, 951-955.

Locke, Peter R., and Steven C. Mann, 1999, Do professional traders exhibit loss realization aversion?, Working paper, Texas Christian University.

Manaster, Steven, and Steven C. Mann, 1996, Life in the pits: Competitive market making and inventory control, *Review of Financial Studies* 9, 953-975.

Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *Journal of Finance*

53, 1775-1798.

Odean, Terrance, 1998, Volume, Volatility, Price, and Profit When All Traders Are Above Average, *Journal of Finance* 53, 1887-1934.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421-436.

Shiller, Robert J., 1989, *Market volatility* (MIT Press, Cambridge, Massachusetts).

Shumway, Tyler, 1998, Explaining returns with loss aversion, Working paper, University of Michigan Business School.

Thaler, Richard H., and Eric J. Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, *Management Science* 36, 643-660.

Table I

Summary Statistics

Table I reports a number of summary statistics for the sample. The sample consists of the trading experience of 426 local traders at the Chicago Board of Trade's Treasury Bond Futures pit over the 236 full trading days during 1998. The variable LA Coefficient $\cdot I(\pi^M_{i,t} > 0)$ represents the daily fraction of traders with morning gains, weighted by a trader-specific loss aversion coefficient as described in Table VII.

Panel A: Statistics by Trader-Day

	Morning		Afternoon	
Variable	Mean	St. Dev.	Mean	St. Dev.
All Trader-Days (N = 82595)				
Profits	0.137	1.009	0.077	1.003
Number of Trades	0.000	0.997	0.000	0.997
Average Trade Size	0.000	0.997	0.000	0.997
Total Dollar Risk	0.000	0.997	0.000	0.997
Absolute Inventory	0.000	0.997	0.000	0.997
Traders with Profitable Mornings (N = 55877)				
Profits	0.500	0.760	0.078	0.881
Number of Trades	-0.035	0.986	-0.066	0.980
Average Trade Size	-0.063	0.967	-0.046	0.989
Total Dollar Risk	-0.114	0.879	-0.098	0.887
Absolute Inventory	-0.104	0.867	-0.065	0.924
Traders with Losing Mornings (N = 26718)				
Profits	-0.620	1.046	0.076	1.219
Number of Trades	0.066	1.013	0.124	1.016
Average Trade Size	0.119	1.040	0.086	1.006
Total Dollar Risk	0.238	1.173	0.204	1.170
Absolute Inventory	0.217	1.197	0.123	1.110

Panel B: Statistics by Day

Variable	Mean	St. Dev.	Minimum	Maximum
Afternoon Price Changes	621.8703	215.383	195.00	1582.00
Fraction with $\pi^M_{i,t} < 0$	0.3238	0.049	0.20	0.50
Average $\pi^M_{i,t}$	0.1350	0.115	-0.52	0.56
LA Coefficient $\cdot I(\pi^M_{i,t} > 0)$	0.0224	0.0048	0.0094	0.0361

Table II

Morning Profits and Afternoon Risk-Taking

Table II reports the results of a number of different regressions relating morning profits to afternoon risk-taking by locals at the CBOT. All regressions have the basic form,

$$\text{RISK}^{A}_{i,t} = \alpha + \beta_{\pi}\pi^{M}_{i,t} + \beta_{I}|\text{INV}^{M}_{i,t}| + \beta_{\pi I}\pi^{M}_{i,t} \cdot |\text{INV}^{M}_{i,t}| + \beta_{R}\text{RISK}^{M}_{i,t} + \varepsilon_{i,t}.$$

T-statistics are in parentheses. Risk is measured in three different ways, as the number of afternoon trades, the average size of afternoon trades, or the cumulative risk-weighted inventory of each trader. The sample contains 82,595 local-days. The standard errors of the fixed-effects PCSE results are allowed to be heteroskedastic and concurrently correlated across locals.

Panel A: Dependent Variable: Afternoon Number of Trades

Method	α	β_{π}	β_{I}	$\beta_{\pi I}$	β_{R}
Pooled OLS	0.0187	-0.1281	0.0528	0.0167	0.2509
	(5.55)	(-30.5)	(15.6)	(16.2)	(74.3)
FM by Trader	0.0320	-0.1530	0.0588	0.0361	0.2433
	(13.9)	(-15.8)	(9.72)	(9.51)	(30.2)
FM by Date	0.0187	-0.2322	0.0577	0.0488	0.1880
	(0.73)	(-28.3)	(13.1)	(10.5)	(28.9)
Fixed Effects PCSE	-	-0.1148	0.0309	0.0155	0.1406
	-	(-25.6)	(9.17)	(15.4)	(27.6)

Panel B: Dependent Variable: Afternoon Average Trade Size

Method	α	β_{π}	β_{I}	$\beta_{\pi I}$	β_{R}
Pooled OLS	-0.0015	-0.0481	0.0523	0.0047	0.2162
	(-0.40)	(-10.1)	(13.7)	(4.13)	(54.4)
FM by Trader	-0.0528	-0.1076	-0.0205	0.0175	0.2088
	(-1.06)	(-3.14)	(-0.23)	(4.26)	(25.1)
FM by Date	0.0032	-0.0995	0.0539	0.0187	0.1724
	(0.22)	(-10.5)	(8.91)	(2.98)	(27.9)
Fixed Effects PCSE	-	-0.0540	0.0540	0.0037	0.1514
	-	(-11.1)	(13.6)	(3.26)	(31.0)

Panel C: Dependent Variable: Afternoon Total Dollar Risk

Method	α	β_{π}	β_{I}	$\beta_{\pi I}$	β_{R}
Pooled OLS	0.0001	0.0001	0.6150	0.0013	0.2703
	(0.04)	(0.05)	(208.0)	(2.26)	(91.3)
FM by Trader	-0.0001	-0.0061	0.6281	0.0069	0.2354
	(-0.09)	(-1.26)	(58.8)	(2.38)	(25.7)
FM by Date	-0.0004	-0.0142	0.5934	0.0153	0.0273
	(-0.06)	(-2.99)	(59.0)	(3.11)	(36.6)
Fixed Effects PCSE	-	-0.0011	0.6123	0.0016	0.2693
	-	(-0.45)	(179.9)	(3.07)	(78.0)

Table III

Binary Results for Morning Profits and Afternoon Risk-Taking

Table III reports the results of a number of different logit and regression models relating morning profits to afternoon risk-taking by locals at the CBOT. All models measure both morning profits and afternoon risk in a binary form, and the logit models have the basic form,

$$\text{Prob}(\text{RISK}^{A}_{i,t} > 0) = \frac{\exp X'\beta}{1 + \exp X'\beta},$$

where

$$X'\beta = \alpha - \beta_{\pi} I(\pi^{M}_{i,t} < 0) + \beta_I |\text{INV}^{M}_{i,t}| + \beta_{\pi I} I(\pi^{M}_{i,t} < 0) \cdot |\text{INV}^{M}_{i,t}| + \beta_R \text{RISK}^{M}_{i,t}.$$

T-statistics are in parentheses. The sample contains 82,595 local-days.

Panel A: Prob(Afternoon Number of Trades > Mean Trades)

Method	α	β_{π}	β_I	$\beta_{\pi I}$	β_R
Pooled Logit	-0.4785	-0.4446	0.1161	-0.0887	0.4119
	(-53.7)	(-28.8)	(11.4)	(-6.07)	(54.4)
FM by Trader	-0.7460	-0.5560	0.1232	-0.0913	0.5075
	(-3.58)	(-13.5)	(3.92)	(-1.82)	(6.89)
FM by Date	-0.5764	-0.5084	0.1413	-0.1110	0.3962
	(-10.5)	(-21.1)	(7.92)	(-3.66)	(24.1)

Panel B: Prob(Afternoon Average Trade Size > Mean Size)

Method	α	β_{π}	β_I	$\beta_{\pi I}$	β_R
Pooled Logit	-0.4722	-0.1518	0.0818	-0.0030	0.4281
	(-46.5)	(-8.80)	(7.01)	(-0.18)	(45.7)
FM by Trader	-0.5605	-0.1637	-0.1153	0.1865	0.4916
	(-4.56)	(-1.34)	(-0.82)	(1.70)	(15.2)
FM by Date	-0.5302	-0.1993	0.1036	0.0126	0.3660
	(-15.9)	(-8.77)	(5.42)	(0.44)	(26.3)

Panel C: Prob(Afternoon Total Dollar Risk > Mean Risk)

Method	α	β_{π}	β_I	$\beta_{\pi I}$	β_R
Pooled Logit	-0.9625	-0.2087	2.0650	-0.4847	1.3770
	(-70.6)	(-9.62)	(67.0)	(-11.2)	(59.1)
FM by Trader	-0.7379	-0.0831	4.5221	0.9834	1.6630
	(-14.7)	(-1.30)	(11.6)	(1.19)	(20.2)
FM by Date	-0.0972	-0.1937	2.5306	0.4702	1.4885
	(-31.1)	(-6.71)	(30.1)	(5.21)	(34.8)

Table IV

Binary Results for Day's Profit and Subsequent Day's Risk-Taking

Table IV reports the results of a logit model relating a given day's profits to the subsequent day's risk-taking by locals at the CBOT. The model measures both morning profits and afternoon risk in a binary form, and the model has the basic form,

$$\text{Prob}(\text{RISK}_{i,t+1} > 0) = \frac{\exp X'\beta}{1 + \exp X'\beta},$$

where

$$X'\beta = \alpha - \beta_{\pi} I(\pi_{i,t} < 0) + \beta_R I(\text{RISK}_{i,t} < 0).$$

T-statistics are in parentheses. The sample contains 82,169 local-days.

Dependent Variable: Daily Total Dollar Risk			
Method	α	β_{π}	β_R
Pooled OLS	1.0792	-0.0073	-0.4752
	(10750)	(0.18)	(778.34)
FM by Trader	1.0749	0.0210	-0.3911
	(55.90)	(1.00)	(-13.75)
FM by Date	1.1426	0.0105	-197.24
	(34.27)	(0.51)	(-20.55)

Table V

Morning Profits and Afternoon Price Leadership

Table V reports the results of a number of different regressions relating afternoon risk-taking to afternoon price leadership by locals at the CBOT. All regressions have the basic form,

$$\Delta_{i,t}^{A} - \bar{\Delta}_{i}^{A} = \alpha + \beta_{\pi}\pi_{i,t}^{M} + \beta_{I}|\text{INV}_{i,t}^{M}| + \beta_{\pi I}\pi_{i,t}^{M} \cdot |\text{INV}_{i,t}^{M}| + \varepsilon_{i,t}.$$

where $\Delta_{i,t}^{A}$ is the number of price-setting trades made by trader i on the afternoon of date t. This is compared to its average level for trader i, $\bar{\Delta}_{i}^{A}$, and then regressed on the explanatory variables used above. T-statistics are in parentheses. Regressions on an indicator variable that is equal to 1 when morning profits are positive result in qualitatively identical inferences, as do ordered logit regressions with fixed effects by trader.

Dependent Variable: Afternoon Price Leadership ($\Delta_{i,t}^{A} - \bar{\Delta}_{i}^{A}$)

Method	α	β_{π}	β_{I}	$\beta_{\pi I}$
Pooled OLS	-0.0086	-0.0233	0.0148	0.0027
	(-4.9)	(-10.8)	(8.7)	(5.1)
FM by Trader	-0.0084	-0.0285	0.0155	0.0071
	(-6.6)	(-7.6)	(7.1)	(4.8)
FM by Date	-0.0049	-0.0454	0.0134	0.0095
	(-1.1)	(-11.4)	(5.8)	(3.3)
Fixed Effects PCSE	-	-0.0263	0.0126	0.0032
	-	(-11.8)	(7.3)	(6.5)

Table VI
Price Permanence Regressions

Table VI reports the results of several time-series regressions that relate the change in price over the next ten minutes, conditional on a current change in price, to the profitability of the local trader that executes the price-setting trade. All regressions have the basic form,

$$P_{j+n} - P_j = \alpha + \sum_{l=1}^{3} \Big\{ [P_{j-l+1} - P_{j-l}] \cdot \Big[\beta_{P,l} + \beta_{M,l} I(\pi_{i,t}^{M} < 0) + \beta_{T,l}(\tau_{j-l+1} - \tau_{j-l}) \Big] \Big\} + \varepsilon_j$$

where P_{j+n} is the first price that is different from the previous transaction price at least ten minutes after the transaction at price P_j, and where $\tau_j - \tau_{j-1}$ is the number of minutes that transpire between the first transaction at price P_j and the first transaction at P_{j-1}. All estimates are corrected for fifth order autocorrelation, and T-statistics are reported in parentheses. The sample size is 42,906, which corresponds to the number of price-setting trades executed by locals during the year. The dependent variable has a mean of 0.00184, with a standard deviation of 0.1244.

Dependent Variable: Price Change over Ten Minutes ($P_{j+n} - P_j$)

Estimate	Model 1	Model 2	Model 3
α	0.0022	0.0022	0.0022
	(3.49)	(3.49)	(3.49)
$\beta_{P,1}$	-0.9540	-0.9221	-0.9558
	(-72.21)	(-49.56)	(-50.70)
$\beta_{P,2}$	-0.5137	-0.4550	-0.4984
	(-25.86)	(-17.50)	(-18.80)
$\beta_{P,3}$	-0.1888	-0.1696	-0.1948
	(-11.02)	(-7.98)	(-8.80)
$\beta_{M,1}$	-	-0.0658	-0.0342
	-	(-2.49)	(-1.29)
$\beta_{M,2}$	-	-0.1396	-0.1040
	-	(-3.44)	(-2.54)
$\beta_{M,3}$	-	-0.0484	-0.0304
	-	(-1.34)	(-0.84)
$\beta_{T,1}$	-	-	-0.0030
	-	-	(-10.31)
$\beta_{T,2}$	-	-	-0.0009
	-	-	(-1.51)
$\beta_{T,3}$	-	-	-0.0003
	-	-	(-2.78)

Table VII

Aggregate Morning Losses and Afternoon Price Changes

Table VII reports the results of a few time series regressions relating aggregate morning losses to the volatility of afternoon price changes. All regressions have the basic form,

$$\sigma^{A}_{h,t} = \alpha + \beta_{\sigma}\sigma^{M}_{h,t} + \beta_{\Pi}\Pi^{M}_{t} + \varepsilon_{t}$$

where $\sigma^{A}_{h,t}$ measures the volatility of afternoon price changes on date t measured at frequency h, $\sigma^{A}_{h,t}$ measures the volatility of morning price changes on date t measured at frequency h, and Π^{M}_{t} measures aggregate morning losses on day t in one of three different ways: (a) the fraction of locals with losses at 11:00 a.m.; (b) the average of $\pi^{M}_{i,t}$ across traders; and (c) the sum of the product of the indicator variable that is one when morning profits are positive, $I(\pi^{M}_{i,t} > 0)$, and trader i's β_{π} coefficient in the OLS regression,

$$I(\text{RISK}^{A}_{i,t} > 0) = \alpha - \beta_{\pi} I(\pi^{M}_{i,t} < 0) + \beta_{I}|\text{INV}^{M}_{i,t}| + \beta_{\pi I}\pi^{M}_{i,t} \cdot |\text{INV}^{M}_{i,t}| + \varepsilon_{i,t}.$$

The sample size is 236 trading days, all estimates are corrected for first order autocorrelation, and T-statistics are in parentheses.

Dependent Variable: Abnormal Afternoon Volatility ($\sigma^{A}_{h,t}$)

Loss Measure	Frequency	α	β_{Δ}	β_{Π}
Fraction with Losses · -1	One Second	-0.48	0.41	-1.47
		(-1.24)	(6.82)	(-1.25)
Average Normalized Profit	One Second	0.04	0.43	-0.32
		(0.40)	(6.97)	(-0.61)
LA Coefficient $\cdot I(\pi^{M}_{i,t} > 0)$	One Second	-0.57	0.39	-25.43
		(-2.09)	(6.66)	(-2.13)
LA Coefficient $\cdot I(\pi^{M}_{i,t} > 0)$	One Minute	-0.57	0.19	-25.23
		(-1.96)	(3.08)	(-2.02)
LA Coefficient $\cdot I(\pi^{M}_{i,t} > 0)$	Five Minutes	-0.53	0.30	-24.01
		(-1.90)	(5.00)	(-1.97)
LA Coefficient $\cdot I(\pi^{M}_{i,t} > 0)$	Ten Minutes	-0.34	0.22	-15.13
		(-1.17)	(3.65)	(-1.21)
LA Coefficient $\cdot I(\pi^{M}_{i,t} > 0)$	Half-Day	0.64	0.11	-10.78
		(2.27)	(1.48)	(-0.78)

Figure 1: Morning Profit Percentile and Afternoon Risk-Taking

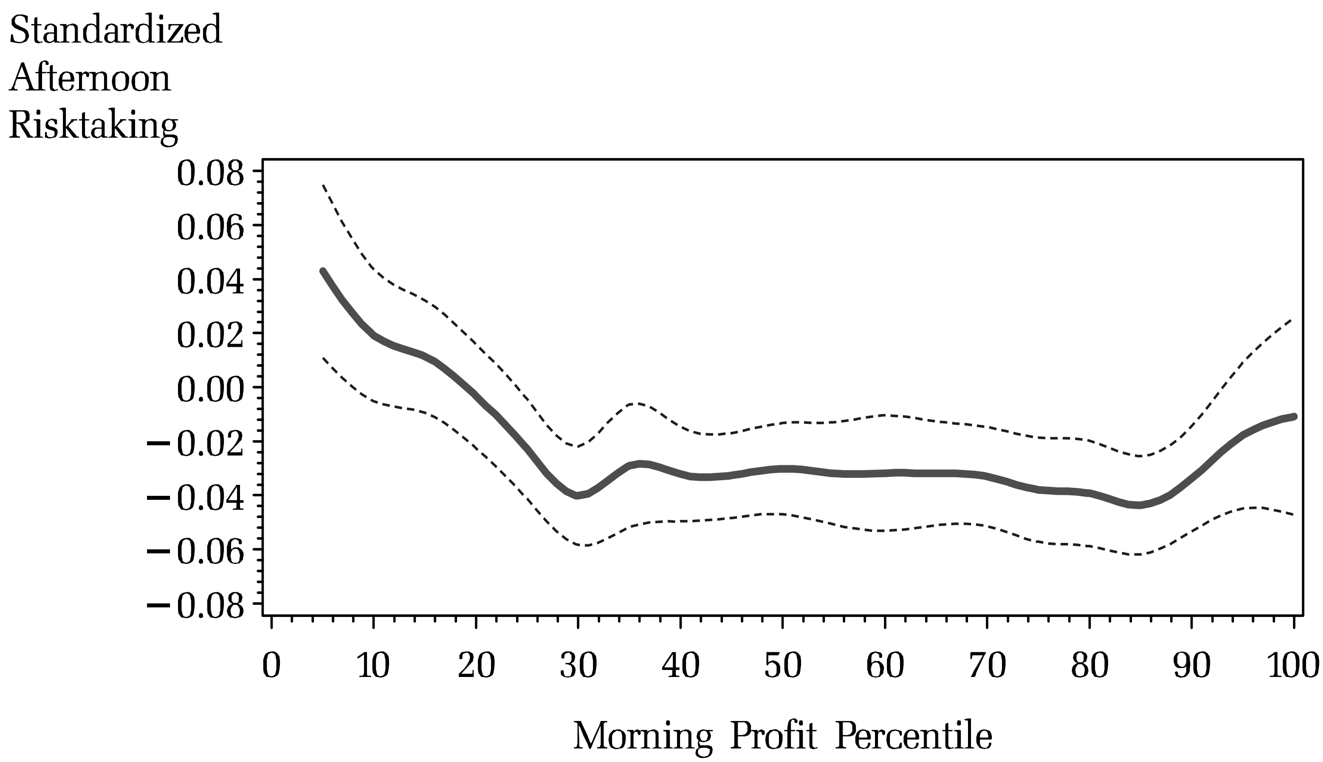

This figure plots the time-series averages of 236 daily cross-sectional semi-parametric regressions of afternoon total dollar risk on morning profit percentile. The regressions are kernel-smoothed and the dashed lines reflect two-standard error bands of the time series-averaged regressions.