Do Investors Trade Too Much?

Terrance Odean

The American Economic Review, Vol. 89, No. 5. (Dec., 1999), pp. 1279-1298.

Do Investors Trade Too Much?

By TERRANCE ODEAN*

Trading volume on the world's markets seems high, perhaps higher than can be explained by models of rational markets. For example, the average annual turnover rate on the New York Stock Exchange (NYSE) is currently greater than 75 percent[1] and the daily trading volume of foreign-exchange transactions in all currencies (including forwards, swaps, and spot transactions) is roughly one-quarter of the total annual world trade and investment flow (James Dow and Gary Gorton, 1997). While this level of trade may seem disproportionate to investors' rebalancing and hedging needs, we lack economic models that predict what trading volume in these market should be. In theoretical models trading volume ranges from zero (e.g., in rational expectation models without noise) to infinite (e.g., when traders dynamically hedge in the absence of trading costs). But without a model which predicts what trading volume should be in real markets, it is difficult to test whether observed volume is too high.

If trading is excessive for a market as a whole, then it must be excessive for some groups of participants in that market. This paper demonstrates that the trading volume of a particular class of investors, those with discount brokerage accounts, is excessive.

Alexandros V. Benos (1998) and Odean (1998a) propose that, due to their overconfidence, investors will trade too much. This paper tests that hypothesis. The trading of discount brokerage customers is good for testing the overconfidence theory of excessive trading because this trading is not complicated by agency relationships. Excessive trading in retail brokerage accounts could, on the other hand, result from either investors' overconfidence or from brokers churning accounts to generate commissions. Excessive institutional trading, too, might result from overconfidence or from agency relationships. Dow and Gorton (1997) develop a model in which money managers, who would otherwise not trade, do so to signal to their employers that they are earning their fees and are not "simply doing nothing."

While the overconfidence theory is tested here with respect to a particular group of traders, other groups of traders are likely to be overconfident as well. Psychologists show that most people generally are overconfident about their abilities (Jerome D. Frank, 1935) and about the precision of their knowledge (Baruch Fischhoff et al., 1977; Marc Alpert and Howard Raiffa, 1982; Sarah Lichtenstein et al., 1982). Security selection can be a difficult task, and it is precisely in such difficult tasks that people exhibit the greatest overconfidence. Dale Griffin and Amos Tversky (1992) write that when predictability is very low, as in securities markets, experts may even be more prone to overconfidence than novices. It has been suggested that investors who behave nonrationally will not do well in financial markets and will not continue to trade in them. There are reasons, though, why we might expect those who actively trade in

* Graduate School of Management, University of California, Davis, CA 95616. This paper is based on my dissertation at the University of California-Berkeley. I would like to thank Brad Barber, Hayne Leland, David Modest, Richard Roll, Mark Rubinstein, Paul Ruud, Richard Thaler, Brett Trueman, and the participants at the Berkeley Program in Finance, the National Bureau of Economic Research behavioral finance meetings, the Conference on Household Financial Decision Making and Asset Allocation at The Wharton School, the Western Finance Association meetings, and the Russell Sage Institute for Behavioral Economics, and seminar participants at the University of California-Berkeley, the Yale School of Management, the University of California-Davis, the University of Southern California, the University of North Carolina, Duke University, the University of Pennsylvania, Stanford University, the University of Oregon, Harvard University, the Massachusetts Institute of Technology, Dartmouth College, the University of Chicago, the University of British Columbia, Northwestern University, the University of Texas, UCLA, the University of Michigan, and Columbia University for helpful comments. I would also like to thank Jeremy Evnine and especially the discount brokerage house which provided the data necessary for this study. Financial support from the Nasdaq Foundation and the American Association of Individual Investors is gratefully acknowledged.

[1] The NYSE website (http://www.nyse.com/public/market/2c/2cix.htm) reports 1998 turnover at 76 percent.

financial markets to be more overconfident than the general population. People who are more overconfident in their investment abilities may be more likely to seek jobs as traders or to actively trade on their own account. This would result in a selection bias in favor of overconfidence in the population of investors. Survivorship bias may also favor overconfidence. Traders who have been successful in the past may overestimate the degree to which they were responsible for their own successes—as people do in general (Ellen J. Langer and Jane Roth, 1975; Dale T. Miller and Michael Ross, 1975)—and grow increasingly overconfident. These traders will continue to trade and will control more wealth, while others may leave the market (e.g., lose their jobs or their money). Simon Gervais and Odean (1999) develop a model in which traders take too much credit for their own successes and thereby become overconfident.

Benos (1998) and Odean (1998a) develop models in which overconfident investors trade more and have lower expected utilities than they would if they were fully rational.[2] The more overconfident an investor, the more he trades and the lower his expected utility. Rational investors correctly assess their expected profits from trading. When trading is costly rational investors will not make trades if the expected returns from trading are insufficient to offset costs [e.g., Sanford J. Grossman and Joseph E. Stiglitz (1980) model rational traders who buy investment information only when the gains in expected utility due to the information offset its cost]. Overconfident investors, on the other hand, have unrealistic beliefs about their expected trading profits. They may engage in costly trading, even when their expected trading profits are insufficient to offset the costs of trading, simply because they overestimate the magnitude of expected profits. Benos (1998) and Odean (1998a) model overconfidence with the assumption that investors overestimate the precision of their information signals. In this framework, at the worst, overconfident investors believe they have useful information when in fact they have no information. These models do not allow for systematic misinterpretation of information. Thus the worst expected outcome for an overconfident investor is to have zero expected gross profits from trading and expected net losses equal to his trading costs.

This paper tests whether the trading profits of discount brokerage customers are sufficient to cover their trading costs. The surprising finding is that not only do the securities that these investors buy not outperform the securities they sell by enough to cover trading costs, but on average the securities they buy underperform those they sell. This is the case even when trading is not apparently motivated by liquidity demands, tax-loss selling, portfolio rebalancing, or a move to lower-risk securities.

While investors' overconfidence in the precision of their information may contribute to this finding, it is not sufficient to explain it. These investors must be systematically misinterpreting information available to them. They do not simply misconstrue the precision of their information, but its very meaning.

The next section of the paper describes the data set. Section II describes the tests of excessive trading and presents results. Section III examines performance patterns of securities prior to purchase or sale. Section IV discusses these patterns and speculates about their causes. Section V concludes.

I. The Data

The data for this study were provided by a nationwide discount brokerage house. Ten thousand customer accounts were randomly selected from all accounts which were active (i.e., had at least one transaction) in 1987. The data are in three files: a trades file, a security number to Committee on Uniform Securities Identification Procedures (CUSIP) number file, and a positions file. The trades file includes the records of all trades made in the 10,000 accounts from January 1987 through December 1993. This file has 162,948 records. Each record is made up of an

[2] Other models of overconfident investors include J. Bradford De Long et al. (1991), Albert S. Kyle and F. Albert Wang (1997), Jordi Caballé and József Sákovics (1998), Kent Daniel et al. (1998), and Gervais and Odean (1999), Kyle and Wang (1997) argue that when traders compete for duopoly profits, overconfident traders may reap greater profits. However, this prediction is based on several assumptions that do not apply to individuals trading common stocks.

account identifier, the trade date, the brokerage house's internal number for the security traded, a buy-sell indicator, the quantity traded, the commission paid, and the principal amount. Multiple buys or sells of the same security, in the same account, on the same day, and at the same price are aggregated. The security number to CUSIP table translates the brokerage house's internal numbers into CUSIP numbers. The positions file contains monthly position information for the 10,000 accounts from January 1988 through December 1993. Each of its 1,258,135 records is made up of the account identifier, the year and month, the internal security number, equity, and quantity. Accounts that were closed between January 1987 and December 1993 are not replaced; thus in the later years of the sample the data set may have some survivorship bias in favor of more successful investors.

There are three data sets similar to this one described in the literature. Gary G. Schlarbaum et al. (1978) and others analyze trading records for 2,500 accounts at a large retail brokerage house for the period January 1964 to December 1970; S. G. Badrinath and Wilbur G. Lewellen (1991) and others analyze a second data set provided by the same retail broker for 3,000 accounts over the period January 1971 to September 1979. The data set studied here differs from these primarily in that it is more recent and comes from a discount broker. By examining discount brokerage records I can rule out the retail broker as an influence on observed trading patterns. Brad M. Barber and Odean (1999a) calculate the returns on common securities in 158,000 accounts. (These accounts are different from those analyzed in this paper, but come from the same discount brokerage.) After subtracting transactions costs and adjusting for risk, these accounts underperform the market. Accounts that trade most actively earn the lowest average net returns. Using the same data, Barber and Odean (1999b) find that men trade more actively than women and thereby reduce their returns more so than do women. For both men and women, they also confirm the principal finding of this paper that, on average, the stocks individual investors buy subsequently underperform those they sell.

This study looks at trades of NYSE, American Stock Exchange (ASE), and National Association of Securities Dealers Automated Quotation (NASDAQ) securities for which daily return information is available from the 1994 Center for Research in Security Prices (CRSP) NYSE, ASE, and NASDAQ daily returns file. There are 97,483 such trades: 49,948 purchases and 47,535 sales. 62,516,332 shares are traded: 31,495,296 shares, with a market value of \$530,719,264, are purchased and 31,021,036 shares, with a market value of \$579,871,104, are sold. Weighting each trade equally the average commission for a purchase is 2.23 percent and for a sale is 2.76 percent.[3] Average monthly turnover is 6.5 percent.[4] The average size decile of a purchase is 8.65 and of a sale is 8.68, 10 being the decile of the companies with the largest capitalization.

II. Empirical Study

A. *Methodology*

In a market with transaction costs we would expect rational informed traders who trade for the purpose of increasing returns to increase returns, on average, by at least enough to cover transaction costs. That is, over the appropriate horizon, the securities these traders buy will outperform the ones they sell by at least enough to pay the costs of trading. If speculative traders are informed, but overestimate the precision of their information, the securities they buy will, on average, outperform those they sell, but possibly not by enough to cover trading costs. If these traders believe they have information, but actually have none, the securities they buy will, on average, perform about the same as those they sell before factoring in trading costs. Overconfidence in only the precision of unbiased information will not, in and of itself, cause expected trading losses beyond the loss of transactions costs.

If instead of (or in addition to) being overconfident in the precision of their information, investors are overconfident about their ability to interpret information, they may incur average

[3] Weighting each trade by its equity value, the average commission for a purchase is 0.9 and for a sale is 0.8.

[4] I estimate turnover as one-half the average monthly equity value of all trades (purchases and sales) divided by the average equity value of all monthly position statements.

trading losses beyond transactions costs. Suppose investors receive useful information but are systematically biased in their interpretation of that information; that is, the investors hold mistaken beliefs about the mean, instead of (or in addition to) the precision of the distribution of their information. If they believe they are correctly interpreting information that they misinterpret, they may choose to buy or sell securities which they would not have otherwise bought or sold. They may even buy securities that, on average and before transaction costs, underperform the ones they sell.

To test for overconfidence in the precision of information, I determine whether the securities investors in this data set buy outperform those they sell by enough to cover the costs of trading. To test for biased interpretation of information, I determine whether the securities they buy underperform those they sell when trading costs are ignored. I look at return horizons of four months (84 trading days), one year (252 trading days), and two years (504 trading days) following a transaction.[5] Returns are calculated from the CRSP daily return files.

To calculate the average return to securities bought (sold) in these accounts over the T ($T =$ 84, 252, or 504) trading days subsequent to the purchase (sale), I index each purchase (sale) transaction with a subscript i, $i = 1$ to N. Each transaction consists of a security, j_i, and a date, t_i. If the same security is bought (sold) in different accounts on the same day, each purchase (sale) is treated as a separate transaction. The average return to the securities bought over the T trading days subsequent to the purchase is:

$$(1) \qquad R_{P,T} = \frac{\sum_{i=1}^{N} \prod_{\tau=1}^{T} (1 + R_{j_i, t_i+\tau})}{N} - 1,$$

where $R_{j,t}$ is the CRSP daily return for security j on date t. Note that return calculations begin the day after a purchase or a sale so as to avoid incorporating the bid-ask spread into returns.

In this data set, the average commission paid when a security is purchased is 2.23 percent of the purchase price. The average commission on a sale is 2.76 percent of the sale price. Thus if one security is sold and the sale proceeds are used to buy another security the total commissions for the sale and purchase average about 5 percent. To get a rough idea of the effective bid-ask spread I calculate at the average difference between the price at which a security is purchased and its closing price on the day of the purchase and calculate the average difference between the closing price on the day of the sale and the selling price. These are 0.09 percent and 0.85 percent, respectively. I add these together to obtain 0.094 percent as an estimate of the average effective spread for these investors.[6] Thus the average total cost of a round-trip trade is about 5.9 percent. An investor who sells securities and buys others because he expects the securities he is buying to outperform the ones he is selling will have to realize, on average and weighting trades equally, a return nearly 6 percent higher on the security he buys just to cover trading costs.

The first hypothesis tested here is that, over horizons of four months, one year, and two years, the average returns to securities bought minus the average returns to securities sold are less than the average round-trip trading costs of 5.9 percent. This is what we expect if investors are sufficiently overconfident about the precision of their information. The null hypothesis (N1) is that this difference in returns is greater than or equal to 5.9 percent. The null is consistent with rationality. The second hypothesis is that over these same horizons the average returns to securities bought are less than those to securities sold, ignoring trading costs. This hypothesis implies that investors must actually misinterpret useful information. The null hypothesis (N2) is that average returns to securities bought are greater than or equal to those sold.

[5] Investment horizons will vary among investors and investments. Shlomo Benartzi and Richard H. Thaler (1995) have estimated the average investor's investment horizon to be one year and, during this period, NYSE securities turned over about once every two years. At the time of this analysis, CRSP data was available through 1994. For this reason two-year subsequent returns are not calculated for transactions dates in 1993.

[6] Barber and Odean (1999a) estimate the bid-ask spread of 1.00 percent for individual investors from 1991 to 1996. Mark M. Carhart (1997) estimates trading costs of 0.21 percent for purchases and 0.63 percent for sales made by open-end mutual funds from 1966 to 1993.

B. *Significance Testing*

The study compares the average return to purchased securities subsequent to their purchase and the average return to sold securities subsequent to their sale. These returns are averaged over the trading histories of individual investors and across investors. Many individual securities are bought or sold on more than one date and may even be bought or sold by different investors on the same date. Suppose, for example, that one investor purchases a particular stock and that a month later another investor purchases the same stock. The returns earned by this stock over four-month periods subsequent to each of these purchases are not independent because the periods overlap for three months. Because returns to individual stocks during overlapping periods are not independent, statistical tests which require independence cannot be employed here. Instead statistical significance is estimated by bootstrapping an empirical distribution for differences in returns to purchased and sold securities. This empirical distribution is generated under the assumption that subsequent returns to securities bought and securities sold are drawn from the same underlying distribution. The methodology is similar to that of William Brock et al. (1992) and David L. Ikenberry et al. (1995). Barber et al. (1999) test the acceptance and rejection rates for this methodology and find that it performs well in random samples. For each security in the sample for which CRSP return data are available a replacement security is drawn, with replacement, from the set of all CRSP securities of the same size decile and same book-to-market quintile as the original security. Using the replacement securities together with the original observation dates, average returns are calculated for the 84, 252, and 502 trading days following dates on which sales or purchases were observed. For example, suppose that in the original data set security A is sold on October 14, 1987, and August 8, 1989, and is bought on April 12, 1992. If security B is drawn as security A's replacement, then in calculating the average return to replacement securities sold, returns to security B following October 14, 1987, and August 8, 1989, will be computed; and in calculating the average return to replacement securities bought, returns to security B following April 12, 1992, will be computed. Replacements are drawn for each security and then average returns subsequent to dates on which the original securities were purchased and were sold are calculated for the replacement securities. These averages and their differences constitute one observation from the empirical distribution. One thousand such observations are made. The null hypothesis (N2) that the securities investors buy outperform (or equally perform) those they sell is rejected at the α percent level if the average subsequent return of purchases minus that of sales in the actual data is less than the α percentile average return of purchases minus that of sales in the empirical distribution. The null hypothesis (N1) that the securities investors buy outperform (or equally perform) those they sell by at least 5.9 percent (the cost of trading) is rejected at the α percent level if the average subsequent return of purchases minus that of sales minus 5.9 percent in the data set is less than the α percentile average return of purchases minus that of sales in the empirical distribution.

This test tries to answer the following question: Suppose that instead of buying and selling the securities they did buy and sell, these investors had randomly chosen securities of similar size and book-to-market ratios to buy and sell; if each security actually traded were replaced, for all of its transactions, by the randomly selected security, how likely is it that, for the randomly selected replacement securities, the returns subsequent to purchases would underperform returns subsequent to sales by as much as is observed in the data?

C. *Results*

Table 1 presents the principal results in this paper. Panel A reports results for all purchases and all sales of securities in the database. Panels B–F give results for various partitions of the data.[7] The most striking result in Table 1 is that

[7] The empirical distributions used for significance testing for various partitions of the data were derived simultaneously.

TABLE 1—AVERAGE RETURNS FOLLOWING PURCHASES AND SALES

	n	84 trading days later	252 trading days later	504 trading days later
Panel A: All Transactions				
Purchases	49,948	1.83	5.69	−24.00
Sales	47,535	3.19	9.00	27.32
Difference		−1.36	−3.31	−3.32
N1		(0.001)	(0.001)	(0.001)
N2		(0.001)	(0.001)	(0.002)
Panel B: Purchases Within Three Weeks of Sales—Sales for Profit and of Total Position—Size Decile of Purchase Less Than or Equal to Size Decile of Sale				
Purchases	7,503	0.11	5.45	22.31
Sales	5,331	2.62	11.27	31.22
Difference		−2.51	−5.82	−8.91
N1		(0.001)	(0.001)	(0.001)
N2		(0.002)	(0.003)	(0.019)
Panel C: The 10 Percent of Investors Who Trade the Most				
Purchases	29,078	2.13	7.07	25.28
Sales	26,732	3.04	9.76	28.78
Difference		−0.91	−2.69	−3.50
N1		(0.001)	(0.001)	(0.001)
N2		(0.001)	(0.001)	(0.010)
Panel D: The 90 Percent of Investors Who Trade the Least				
Purchases	20,870	1.43	3.73	22.18
Sales	20,803	3.39	8.01	25.44
Difference		−1.96	−4.28	−3.26
N1		(0.001)	(0.001)	(0.001)
N2		(0.001)	(0.001)	(0.001)
Panel E: 1987–1989				
Purchases	25,256	0.05	1.47	20.44
Sales	26,732	1.70	4.88	22.95
Difference		−1.65	−3.41	−2.51
N1		(0.001)	(0.001)	(0.001)
N2		(0.001)	(0.001)	(0.006)
Panel F: 1990–1993				
Purchases	29,078	4.67	12.29	32.04
Sales	26,732	5.93	16.44	38.89
Difference		−1.26	−4.15	−6.85
N1		(0.001)	(0.001)	(0.001)
N2		(0.004)	(0.001)	(0.005)

Notes: Average percent returns are calculated for the 84, 252, and 504 trading days following purchases and following sales in the data set trades file. Using a bootstrapped empirical distribution for the difference in returns following buys and following sells, the null hypotheses N1 and N2 can be rejected with p-values given in parentheses. N1 is the null hypothesis that the average returns to securities subsequent to their purchase is at least 5.9 percent greater than the average returns to securities subsequent to their sale. N2 is the null hypothesis that the average returns to securities subsequent to their purchase is greater than or equal to the average returns to securities subsequent to their sale.

for all three follow-up periods and for all partitions of the data the average subsequent return to securities bought is less than that to securities sold. Not only do the investors pay transactions costs to switch securities, but the securities they buy underperform the ones they sell. For example, for the entire sample over a one-year horizon the average return to a purchased security is 3.3 percent lower than the average return to a security sold.

The rows labeled N1 give significance levels for rejecting the null hypothesis that the expected returns to securities purchased are 5.9 percent (the average cost of a round-trip trade) or more greater than the expected returns to securities sold. Statistical significance is determined from the empirical distributions described above; p-values are given in parentheses. For the unpartitioned data (Panel A) N1 can be rejected at all three horizons with $p < 0.001$. The rows labeled N2 report significance levels for rejecting the second null hypothesis (N2) that the expected returns to securities purchased are greater than or equal to those of securities sold (ignoring transactions costs). For the unpartitioned data (Panel A) N2 can be rejected at horizons of 84 and 252 trading days with $p < 0.001$ and at 504 trading days with $p < 0.002$.

These investors are not making profitable trades. Of course investors trade for reasons other than to increase profit. They trade to meet liquidity demands. They trade to move to more, or to less, risky investments. They trade to realize tax losses. And they trade to rebalance. For example, if one security in his portfolio appreciates considerably, an investor may sell part of his holding in that security and buy others to rebalance his portfolio. Panel B examines trades for which these alternative motivations to trade have been largely eliminated. This panel examines only sales and purchases where a purchase is made within three weeks of a sale; such transactions are unlikely to be liquidity motivated since investors who need cash for three weeks or less can borrow more cheaply (e.g., using credit cards) than the cost of selling and later buying securities. All of the sales in this panel are for a profit; so these securities are not sold in order to realize tax losses (and they are not short sales). These sales are of an investor's complete holding in the security sold; so most of these sales are not motivated by a desire to rebalance the holdings of an appreciated security.[8] Also this panel examines only sales and purchases where the purchased security is from the same size decile as the security sold or from a smaller size decile (CRSP size deciles for the year of the transaction); since size has been shown to be highly correlated with risk, this restriction is intended to eliminate most instances where an investor intentionally buys a security of lower expected return than the one he sells because he is hoping to reduce his risk.

We see in Panel B that when all of these alternative motivations for trading are (at least partially) eliminated, investors actually perform worse over all three evaluation periods; over a one-year horizon the securities these investors sell underperform those they buy by more than 5 percent. Sample size is, however, greatly reduced and statistical significance slightly lower. Both null hypotheses can still be rejected.

In Panels C–F the data set is partitioned to test the robustness of these results. Panel C examines the trades made by the 10 percent of the investors in the sample who make the greatest number of trades. Panel D is for trades made by the 90 percent of investors who trade least. The securities frequent traders buy underperform those they sell by a bit less than is the case for the investors who trade least. It may be that the frequent traders are better at security picking. Or it may be that because they hold securities for shorter periods, the average returns in periods following purchases and sales are more alike. If, for example, an investor buys a security and sells it ten trading days later, the 84-trading-day period following the purchase will overlap the 84-trading-day period following the sale on 74 trading days. Thus the returns for the two 84-day periods are likely to be more alike than they would be if there were no overlap. Panel E examines trades made during 1987–1989 and Panel F those made during 1990–1993. For panels C, D, E, and F, we can reject both of the null hypotheses at all three horizons.

D. *Calendar-Time Portfolios*

To establish the robustness of the statistical results presented above, I calculate three measures of performance that analyze the returns on calendar-time portfolios of securities purchased and sold in this data set. The calendar-time portfolio method eliminates the problem of cross-sectional dependence among sample firms, since the returns on sample firms are aggregated into two portfolio returns.[9] These intercept tests test whether the difference in the average subsequent returns to securities purchased and to securities sold in the data set is significantly different than zero. Transactions costs are ignored. Thus the null hypothesis tested here is N2, whether average returns to securities bought are greater than or equal to those sold even before subtracting transactions costs.

I calculate calendar-time returns for securities purchased as follows. For each calendar month t, I calculate the return on a portfolio with one position in a security for each occurrence of a purchase of that security by any investor in the data set during the "portfolio formation period" (of 4, 12, or 24 months) preceding the calendar month t. A security may have been purchased on several occasions during the portfolio formation period. If so, each purchase generates a separate position in the portfolio. Each position is weighed equally. Similarly I form and calculate returns for a portfolio based on sales.

The first performance measure I calculate is simply the average monthly calendar-time return on the "Buy" portfolio minus that on the "Sell" portfolio. Results for portfolio formation periods of 4, 12, and 24 months are reported in Table 2, Panel A. For all three periods the monthly returns on this "long-short" portfolio are reliably negative.

Second, I employ the theoretical framework of the Capital Asset Pricing Model and estimate

[8] The profitability of a sale and whether that sale is of a complete position are determined by reconstructing an investor's portfolio from past trades. Exactly how this is done is described in Odean (1998b). It is possible that there are some cases where it appears that an investor's entire position has been sold, but the investor continues to hold shares of that security acquired before 1987. It is also possible that the investor continues to hold this security in a different account.

[9] This discussion of calendar-time portfolio methods draws heavily on Barber et al's. (1999) discussion and analysis of these methods.

TABLE 2—MONTHLY ABNORMAL RETURNS FOR CALENDAR-TIME PORTFOLIOS

Formation period	4 months	12 months	24 months
Panel A: Raw Returns			
Return	−0.293***	−0.225***	−0.137**
	(0.081)	(0.071)	(0.067)
Panel B: CAPM Intercept			
Excess return	−0.311***	−0.234***	−0.152**
	(0.080)	(0.073)	(0.068)
Beta	0.036**	−0.012	0.020
	(0.018)	(0.020)	(0.018)
Panel C: Fama-French Three-Factor Intercept			
Excess return	−0.249***	−0.207***	−0.136**
	(0.075)	(0.070)	(0.065)
Beta	−0.001	−0.007	0.008
	(0.019)	(0.020)	(0.019)
Size coefficient	0.031	0.075***	0.068***
	(0.028)	(0.026)	(0.024)
HML coefficient	−0.138***	−0.051	−0.025
	(0.035)	(0.032)	(0.029)

Notes: Raw returns (Panel A) are $R_{Bt} - R_{St}$, where R_{Bt} is the percent return in month t on a equally weighted portfolio with one position in a security for each occurrence of a purchase of that security by any investor in the data set in the 4, 12, or 24 months (the formation period) preceding month t and R_{St} is the percent return in month t on a equally weighted portfolio with one position in a security for each occurrence of a sale of that security by any investor in the data set in the 4, 12, or 24 months preceding month t. The CAPM intercept is estimated from a time-series regression of $R_{Bt} - R_{St}$ on the market excess return $R_{mt} - R_{rf}$. The Fama-French three-factor intercept is estimated from a time-series regressions of $R_{Bt} - R_{St}$ on the market excess return, a zero-investment size portfolio (SMB_t), and a zero-investment book-to-market portfolio (HML_t). Standard errors are in parentheses.

***,**Significant at the 1- and 5-percent level, respectively. The null hypothesis for beta (the coefficient estimate on the market excess return) is H_0: $\beta = 1$.

Jensen's alpha (Michael C. Jensen, 1969) by regressing the monthly return of the buy-minus-sell portfolio on the market excess return. That is, I estimate:

$$(2) \quad R_{Bpt} - R_{Spt} = \alpha_p + \beta_p(R_{mt} - R_{ft}) + \epsilon_{pt}$$

where:

R_{Bpt} = the monthly return on the calendar-time portfolio based on purchases;
R_{Spt} = the monthly return on the calendar-time portfolio based on sales;
R_{Mt} = the monthly return on a value-weighted market index;
R_{ft} = the monthly return on T-bills;[10]
β_p = the market beta; and
ϵ_{pt} = the regression error term.

The subscript p denotes the parameter estimates and error terms for the regression of returns for calendar-time portfolios with a p month formation period. Results from these regressions are reported in Table 2, Panel B. Excess return estimates (α) are reliably negative for all three portfolio formation periods (4, 12, and 24 months).

Third, I employ an intercept test using the three-factor model developed by Eugene F. Fama and Kenneth R. French (1993). I estimate the following monthly time-series regression:

$$(3) \quad R_{Bpt} - R_{Spt} = \alpha_p + \beta_p(R_{mt} - R_{ft}) + z_p SMB_t + h_p HML_t + \epsilon_{pt}$$

where SMB_t is the return on a value-weighted portfolio of small stocks minus the return on a value-weighted portfolio of big stocks and HML_t is the return on a value-weighted portfolio of high book-to-market stocks minus the return on a value-weighted portfolio of low book-to-market stocks.[11]

Fama and French (1993) argue that the risk of common stock investments can be parsimoniously summarized as risk related to the market, firm size, and a firm's book-to-market ratio. I measure these three risk exposures using the coefficient estimates on the market excess return $R_{mt} - R_{ft}$, the size zero-investment portfolio (SMB_t), and the book-to-market zero-investment portfolio (HML_t) from the three-factor regressions. Portfolios with above-average market risk have betas greater than one, $\beta_P > 1$. Portfolios with a tilt toward large (growth) stocks relative to a value-weighted market index have size (book-to-market) coefficients less than zero, $z_P < 0$ ($h_P < 0$).

The regression yields parameter estimates of α, β, z, and h. The error term in the regression is denoted by ϵ_t. The estimate of

[10] The return on T-bills is from *Stocks, Bonds, Bills, and Inflation: 1997 Yearbook* (Ibbotson Associates, 1997).

[11] The construction of these factors is described in Fama and French (1993). I thank Kenneth French for providing these data.

the intercept term (α) provides a test of the null hypothesis that the difference in the mean monthly excess returns of the "buy" and "sell" calendar-time portfolios is zero.[12] As reported in Table 2, Panel C, excess returns for this model are reliably negative for all three portfolio formation periods (4, 12, and 24 months). There is some evidence that, compared to the stocks they sell, these investors tend to buy smaller, growth stocks. After adjusting for size and book-to-market effects, there is no evidence of systematic differences in the market risk (β) of the stocks they buy and sell.

E. *Security Selection vs. Market Timing*

The posttransaction returns of the stocks these investors purchase are lower than those they sell. This underperformance could be due to poor choices of which stocks to buy and sell or poor choices of when, in general, to buy stocks and when to sell them. That is, the underperformance may be caused by inferior security selection or inferior market timing (or both).

To test whether the underperformance is due to poor security selection, I repeat the analysis of Section II, subsection B, using market-adjusted returns rather than raw returns. From each return subsequent to a purchase or a sale, I subtract the return on the CRSP value-weighted index for the same period. This adjustment removes the effect that market timing might have on performance. Results for all investors during the entire sample period are reported in Table 3. The differences in the market-adjusted returns subsequent to purchases and sales are reliably negative at all three horizons (4, 12, and 24 months) and are similar to the difference in raw returns subsequent to purchases and sales reported in Table 1, Panel A. For example, over the following 12 months, market-adjusted returns to purchases are 3.2 percent less than market-adjusted returns to sales, while raw returns to purchases are 3.3 percent less than raw returns to sales. This supports the hypothesis that these investors make poor choices of which stocks to buy and which to sell.

TABLE 3—AVERAGE MARKET-ADJUSTED RETURNS FOLLOWING PURCHASES AND SALES

	n	84 trading days later	252 trading days later	504 trading days later
Purchases	49,948	−1.33	−2.68	−0.68
Sales	47,535	0.12	0.54	2.89
Difference		−1.45	−3.22	−3.57
N1		(0.001)	(0.001)	(0.001)
N2		(0.001)	(0.001)	(0.002)

Notes: Average percent returns in excess of the CRSP value-weighted index are calculated for the 84, 252, and 504 trading days following purchases and following sales in the data set trades file. Using a bootstarpped empirical distribution for the difference in market-adjusted returns following buys and following sells, the null hypotheses N1 and N2 can be rejected with p-values given in parentheses. N1 is the null hypothesis that the average market-adjusted returns to securities subsequent to their purchase is at least 5.9 percent greater than the average market-adjusted returns to securities subsequent to their sale. N2 is the null hypothesis that the average market-adjusted returns to securities subsequent to their purchase is greater than or equal to the average market-adjusted returns to securities subsequent to their sale. This table reports results for all investors over the entire sample period.

To test whether these investors exhibit an ability to time their entry and exit from the market, I examine whether their entry or exit from the market in one month predicts the next month's market return. I first calculate monthly order imbalance as the dollar value of all purchases in a month divided by the dollar value of all purchases and all sales in that month. I then regress the current month's return of the CRSP value-weighted index on the previous month's order imbalance:

$$(4)\quad R_{mt} = a + b\left(\frac{Buys_{t-1}}{Buys_{t-1} + Sells_{t-1}}\right) + e_t.$$

The coefficient estimate (b) for order imbalance is statistically insignificant ($t = -0.4$, $R^2 = 0.0$). This suggests that poor market timing does not make an important contribution to the subsequent underperformance of the stocks these investors buy relative to those they sell.

[12] The error term in this regression may be heteroskedastic, since the number of securities in the calendar-time portfolio varies from month to month. Barber et al. (1999) find that this heteroskedasticity does not significantly affect the specification of the intercept test in random samples.

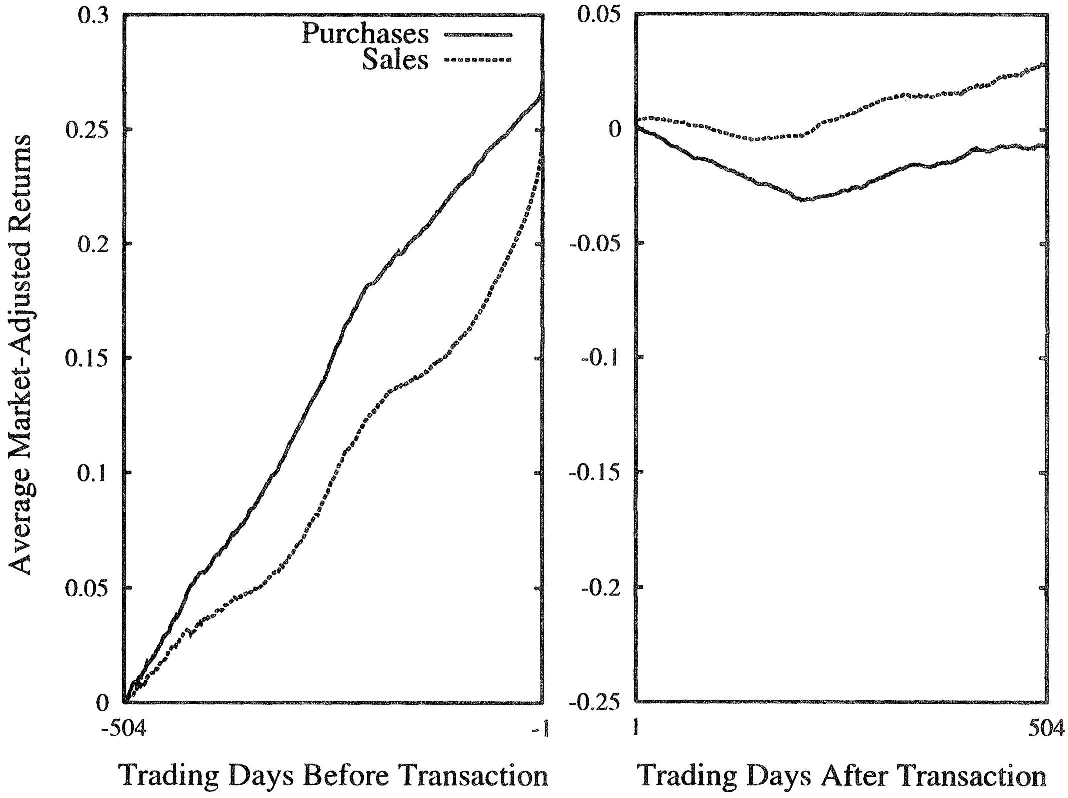

FIGURE 1. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX FOR ALL SECURITIES BOUGHT AND SOLD

Note: 46,830 bought; 44,265 sold.

III. Returns Patterns Before and After Transactions

The securities the investors in this data set buy underperform those they sell. When the investors are most likely to be trading solely to improve performance (Table 1, Panel B), performance gets worse. It appears that these investors have access to information with some predictive content, but they are misinterpreting this information. It is possible that they are misinterpreting a wide variety of information, such as accounting data, technical indicators, and personal knowledge about an company or industry. A simpler explanation is that many of them are misinterpreting the same information. One information set readily available to most investors is recent historical returns.

This section describes return patterns to securities before and after they are purchased and sold by individual investors.

Figures 1 and 2 graph average market-adjusted returns in excess of the CRSP value-weighted index for sales and purchases of securities in the database from two years (504 trading days) before the transaction until two years after it.[13] If such graphs were made for all purchases and sales in the entire market, the

[13] The average market-adjusted return for a set of N transactions for a period of T trading days following each transaction is calculated as:

$$R_{P,T} = \frac{\sum_{i=1}^{N} \left(\prod_{\tau=1}^{T} (1 + R_{j_i,t_i+\tau}) - \prod_{\tau=1}^{T} (1 + R_{M,t_i+\tau}) \right)}{N}$$

where j_i, t_i, and $R_{j,t}$ are defined as in equation (1) and $R_{M,t}$ is the day t return on the CRSP value-weighted market index excluding distributions. If the calculation is done for the CRSP value-weighted market index inclusive of distributions, daily market-adjusted returns are, on average, one basis point lower. This change in indices has virtually no

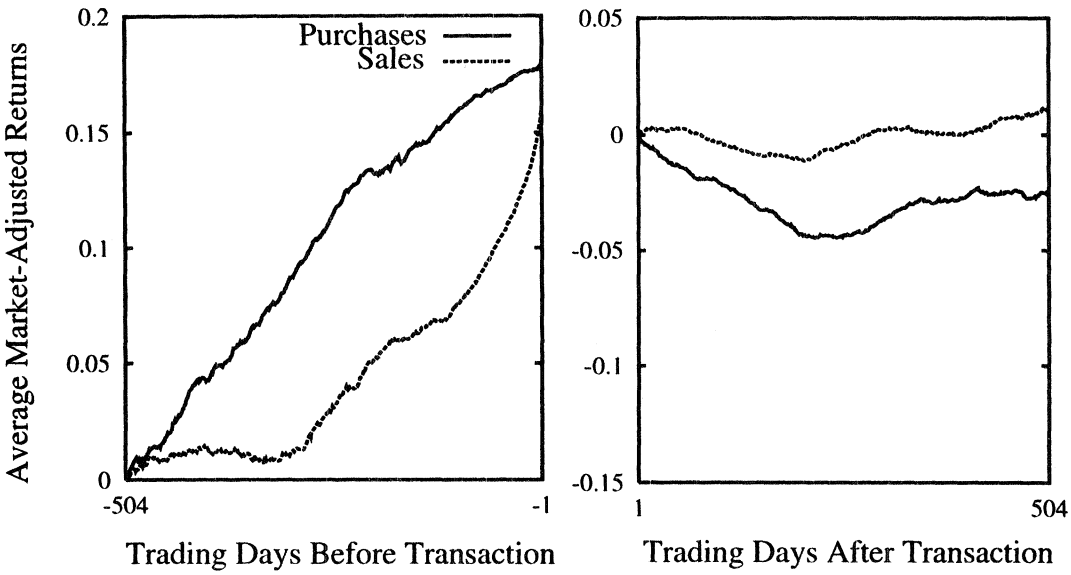

FIGURE 2. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX FOR SECURITIES BOUGHT AND SOLD BY THE 90 PERCENT OF INVESTORS WHO TRADED LEAST

Note: 20,870 bought; 20,803 sold.

paths for returns to sales and to purchases would coincide, since for every purchase there is a sale. The differences in these paths here reflect differences in returns to the securities that these traders in aggregate sold to and bought from the rest of the market.

Figure 1 graphs average market-adjusted returns for all purchases and all sales of securities in the data set for which daily returns are available from CRSP. On average these investors both buy and sell securities which have outperformed the market over the previous two years. This is consistent with the findings of Josef Lakonishok and Seymour Smidt (1986) and others that trading volume is positively correlated with price changes. The securities the investors buy have appreciated somewhat more than those they sell over the entire previous two years, while the securities they sell have appreciated more rapidly in the months preceding sales. Securities purchased underperform the market over the next year, while securities sold perform about as well as the market over the next year. If there were no predictive information in the purchase or sale of a security, and if investors traded in a mix of securities representative of the market, we would expect securities to perform about as well as the market after being purchased or sold. If trading were concentrated in a particular segment of the market, such as small capitalization companies, we would expect that if there were no predictive value to a transaction these securities would perform about as well, relative to the market, as the segment of the market from which they were drawn.[14] In the overall sample the securities that were bought and sold are from about the same average size deciles (8.65 and 8.68). Nevertheless, securities purchased subsequently underperform those sold. The difference in average market-adjusted returns to purchases and sales is statistically significant at the three time horizons for which it is tested: 84, 252, and 504 trading days (Table 3).

As discussed at the end of the Section II,

effect on the market-adjusted returns of purchases and sales relative to each other.

[14] The data analyzed in these graphs extends from 1985 through 1994. Barber and Lyon (1997) find that big firms outperformed small firms from 1984 to 1988 and that small firms outperformed big firms from 1989 to 1994.

subsection C, when securities are held only a short time between purchase and sale, the average returns to purchases and sales over longer horizons will tend to converge. Investors who trade most frequently tend to hold their positions for shorter periods than those who trade less. Active traders may also have shorter trading horizons and so looking at returns one to two years after a transaction may not be relevant for the most active traders. Concentrating on trades of the 90 percent of investors who trade the least accentuates, and facilitates identifying, differences in the returns patterns of securities purchased and sold. Figure 2 graphs average market-adjusted returns for the purchases and sales made by these investors. The differences in returns to purchases and sales is greater in Figure 2 than in Figure 1. Prior to the transaction, purchases have been rising steadily for two years; sales, on the other hand, only started rising a little over a year before the sale but have risen more rapidly in recent months. After a purchase the market-adjusted returns to securities fall over the next eight months or so, nearly as rapidly as they rose over the eight months prior to the purchase. The difference in market-adjusted returns to securities bought and to securities sold following the transactions are statistically significant for all three time horizons at which I have tested, 84, 252, and 504 trading days ($p < 0.001$).

While investors buy and sell securities that have, on average, appreciated prior to purchase or sale, some of the securities they buy and sell have depreciated. The decision to buy or sell a previous winner may be motivated differently than the decision to buy or sell a previous loser. In Figures 3 and 4 the purchases and sales of the 90 percent of investors who trade least are partitioned into previous winners and losers. A security that had a positive raw return over the 126 trading days (six months) preceding a purchase or sale is classified as a previous winner. A security that had negative raw return over this period is a previous loser. Because of the selection criteria, market-adjusted returns are steep and nearly straight for both winners and losers during the evaluation period (-126 to -1 trading days).

In Figure 3 previous winners that are bought by the infrequent traders outperform the market by 60 percent over the entire two years preceding a purchase. They then underperform the market by about 5 percent over the next two years. Previous winners that are sold outperform the market by almost 40 percent over the 15 months before the sale; over the 24th to 16th month before the sale their return is similar to the market's. After the sale they outperform the market by 3 percent over the next two years. Using the tests described in Section II, subsection B, the differences in market-adjusted returns subsequent to transactions for previous winners sold and previous winners bought are statistically significant for time horizons at which I have tested, namely, 84 trading days ($p = 0.002$), 252 trading days ($p = 0.001$), and 504 trading days ($p = 0.001$).

Figure 4 graphs average market-adjusted returns for previous losers that are bought and sold by the infrequent traders. Those that are bought rise, relative to the market, nearly 4 percent over the 24th to 18th month prior to a purchase; then they fall 28.5 percent. Securities sold rise about 1 percent (relative to the market) over the 24th to 19th month prior to the sale and then fall 24.5 percent. After being purchased previous losers continue to underperform the market by about 5.5 percent over the next year. They regain most of this loss in the next year. Previous losers which are sold outperform the market by 1 percent over the next three months. They then lose 5 percent more than the market over the next nine months and finally regain some of this loss. The difference in market-adjusted returns to previous losers bought and previous losers sold following the transactions are statistically significant for the first two time horizons at which I have tested, namely, 84 trading days ($p = 0.001$), and 252 trading days ($p = 0.003$). The difference is not statistically significant for 504 trading days.

In Figures 3 and 4 we see that both securities that previously outperformed the market and those that previously underperformed it, underperform it subsequent to being purchased. There is another class of securities, recent initial public offerings, that have neither previously outperformed or underperformed the market. Figure 5 graphs the average market-adjusted returns for a proxy for newly issued securities over the two years following a purchase. Purchases are included in this graph if the beginning date for the security's listing in the CRSP

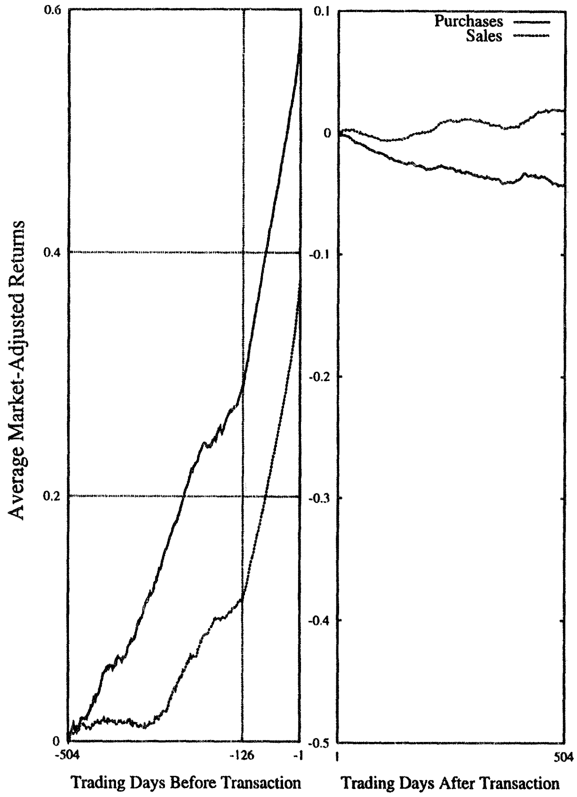

FIGURE 3. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX FOR SECURITIES BOUGHT AND SOLD BY THE 90 PERCENT OF INVESTORS WHO TRADED LEAST, FOR SECURITIES THAT HAD POSITIVE RAW RETURNS OVER THE 126 TRADING DAYS PRECEDING A PURCHASE OR A SALE

Note: 9,688 bought; 12,250 sold.

daily returns file is no more than five trading days prior the date of the purchase. This is not a perfect proxy for new issues, but it does give us some indication of how new issues perform after being purchased. When the trades of all investors are considered, 398 purchases meet this "new issue" criteria. Only 25 sales meet the criteria; because of this small sample size sales are not graphed. (If sales are graphed their return pattern is very similar to that of the purchases.) The "new issues" that the investors buy underperform the market by an average of about 25 percent over the 14 months following the purchase. They recover about half of this loss in the next ten months. The underperformance of the market by new issues noted here is consistent with, though more extreme than, Jay R. Ritter's (1991) and Tim Loughran and Ritter's (1995) findings that after the first day's close initial public offerings tend to underperform the market. When compared to the empirical benchmark distribution the underperformance of these new issues is statistically significant ($p < 0.05$) over the 84-trading-day horizon. The underperformance is not statistically significant for the 252- and 504-trading-day horizons.

Figure 6 graphs average market-adjusted returns over the 20 trading days preceding a transaction for securities bought and sold by the 90 percent of investors who traded least. In this graph securities are classified as previous winners or losers on the basis of their raw returns over the period of 146 to 21 trading days (the seventh through the second month) preceding a purchase or a sale. The securities which investors sell rise sharply in the 20 days preceding a

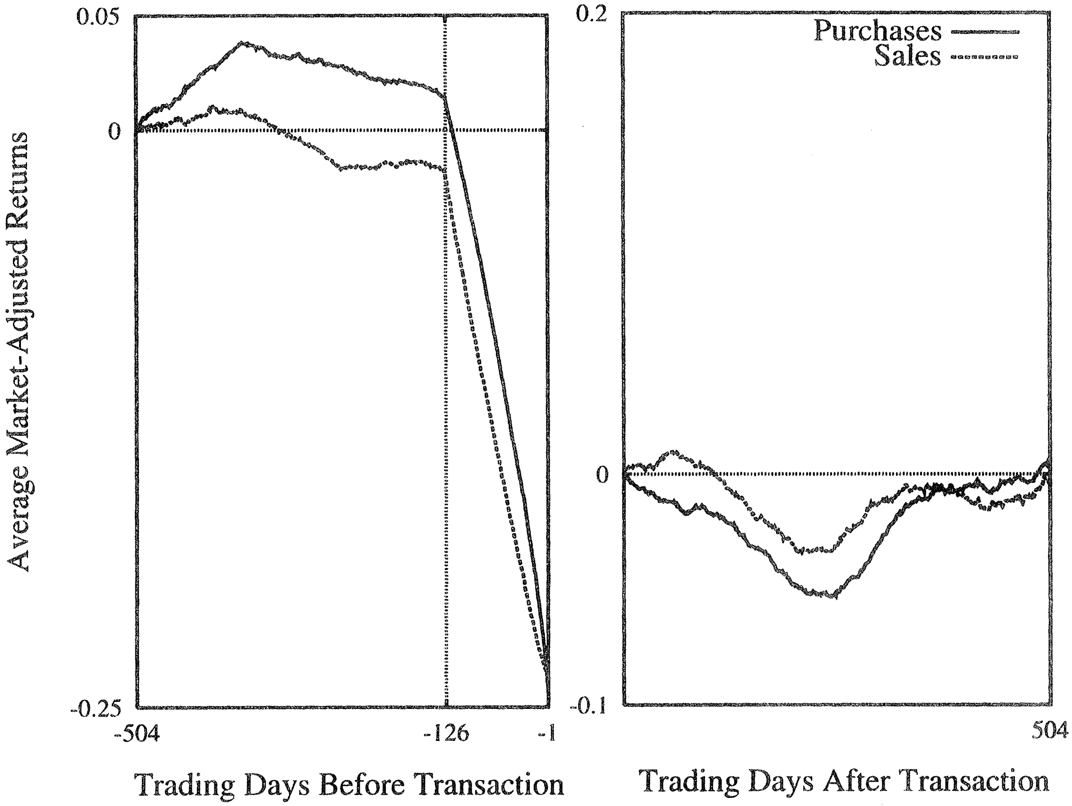

FIGURE 4. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX FOR SECURITIES BOUGHT AND SOLD BY THE 90 PERCENT OF INVESTORS WHO TRADED LEAST, FOR SECURITIES THAT HAD NEGATIVE RAW RETURNS OVER THE 126 TRADING DAYS PRECEDING A PURCHASE OR A SALE

Note: 8,971 bought; 6,602 sold.

sale; previously winning securities rise 4.1 percent and previous losers rise 2.8 percent. Previous winners that they buy also rise while losers they buy fall. When compared to the empirical benchmark distributions, the 20-trading-day market-adjusted returns for previous winners bought, previous winners sold, previous losers bought, and previous losers sold are all significantly different than 0 ($p < 0.001$ in all four cases).

IV. Discussion

The previous section identifies a number of regularities in the return patterns of securities before they are bought or sold by individual investors. These investors buy securities that have experienced greater absolute price changes over the previous two years than the ones they sell (Figures 3 and 4). They buy similar numbers of winners and losers, but they sell far more winners than losers (Figures 3 and 4). Investors sell securities which have risen sharply in the weeks prior to sale. This is true for securities that were previous winners and for previous losers (Figure 6).

I propose that, at least in part, these patterns can be explained quite simply. The buying patterns are caused by the large number of securities from which investors can choose to buy and by the tendency of investors to let their attention be directed towards securities that have experienced abnormally good or bad performance. The selling patterns result from investors' reluctance to sell short and from the disposition effect (i.e., investors' reluctance to realize losses).

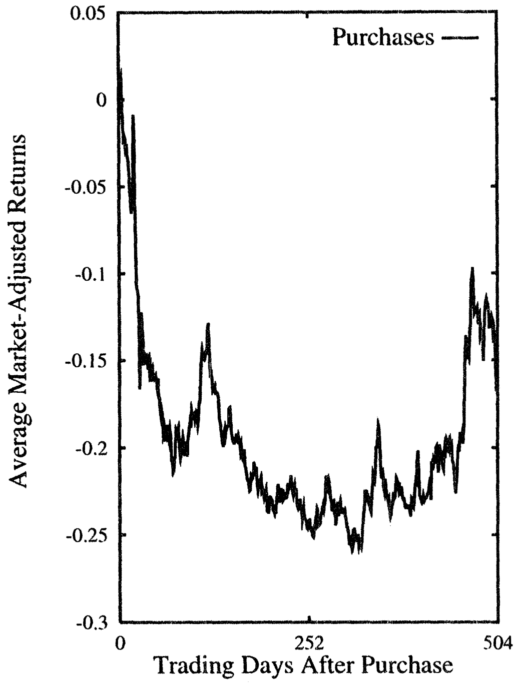

FIGURE 5. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX FOR SECURITIES BOUGHT THAT WERE ISSUED (LISTED ON CRSP) WITHIN FIVE DAYS PRIOR TO PURCHASE

Note: 398 bought.

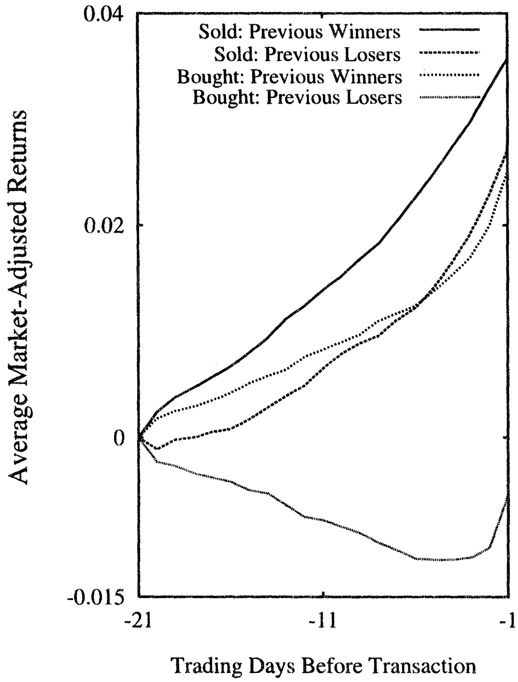

FIGURE 6. AVERAGE RETURNS IN EXCESS OF THE CRSP VALUE-WEIGHTED INDEX OVER THE 20 DAYS PRECEDING A TRANSACTION FOR SECURITIES BOUGHT AND SOLD BY THE 90 PERCENT OF INVESTORS WHO TRADED LEAST

Notes: Previously profitable securities had positive raw returns over the period from 146 to 21 trading days preceding a purchase or a sale. Previously not profitable securities had negative raw returns over the same period. 26,434 previous winners sold. 17,078 previous losers sold. 26,133 previous winners bought. 18,964 previous losers bought.

In Section II, formal hypotheses are subjected to rigorous tests. In this section, conjectures are proposed to explain the return patterns described in Section III. These conjectures are not, however, tested.

Investors face a formidable challenge when looking for a security to buy. There are well over 10,000 securities to be considered. These investors do not have a retail broker available to suggest purchase prospects. While the search for potential purchases can be simplified by confining it to a subset of all securities (e.g., the S&P 500), even then the task of evaluating and comparing each security is beyond what most nonprofessionals are equipped to do. Unable to evaluate each security, investors are likely to consider purchasing securities to which their attention has been drawn. Investors may think about buying securities they have recently read about in the paper or heard about on the news. Securities that have performed unusually well or poorly are more likely to be discussed in the media, more likely to be considered by individual investors and, ultimately, more likely to be purchased.

Once their attention has been directed to potential purchases, investors vary in their propensity to buy previous winners or previous losers. The null hypothesis that the probability of buying previous winners (or losers) is the same for all investors in this data set can be rejected ($p < 0.001$) using a Monte Carlo test described in the Appendix. The separation between those who buy previous winners and those who buy previous losers is greatest for securities which have experienced large price changes.

It may be that those who buy previous winners believe that securities follow trends while those who buy previous losers believe they revert. The investors who believe in trend may buy previous winners to which their attention has been directed, while those who believe in reversion buy previous losers to which their attention has been directed. If investors were as willing to sell securities short as to buy, we might expect them to actively sell as well as to actively buy securities to which their attention was directed. But mostly these investors do not sell short—less than 1 percent of the sales in this data set are short sales. The cost of shorting is high for small investors who usually receive none of the interest on the proceeds of the short sale. Furthermore short selling is not limited in liability and may be considered too risky by many investors.

While theoretical models of financial markets often treat buying and selling symmetrically, for most investors the decision to buy a security is quite different from the decision to sell. In the first place, the formidable search problem for purchases does not apply to sales. Since most investors do not sell short, those seeking a security to sell need only consider the ones they already own. This is usually a manageable handful—in this data set the average number of securities, including bonds, mutual funds, and options as well as stocks, per account is 3.6. Investors can carefully consider selling each security they own regardless of the attention given it in the media.

Though the search for securities to sell is far simpler, in other respects the decision to sell a security is more complex than the decision to buy. When choosing securities to buy, an investor only needs to form expectations about the future performance of those securities. When choosing securities to sell, the investor will consider past as well as future performance. If the investor is rational he will want to balance the advantages or disadvantages of any tax losses or gains he realizes from a sale against future returns he expects a security to earn. If an investor is psychologically motivated he may wish to avoid realizing losses and prefer to sell his winners. In Figures 3 and 4 investors sell nearly twice as many previous winners as previous losers. Using this same data set, Odean (1998b) shows that these investors strongly prefer to sell their winning investments and to hold on to their losing investment even though the winning investments they sell subsequently outperform the losers they continue to hold. Jeffrey Heisler (1997) and Chip Heath et al. (1999) find that investors display similar behavior when closing futures contracts and exercising employee stock options. This behavior is predicted by Hersh Shefrin and Meir Statman's disposition theory (1985) and, in more general terms, by Daniel Kahneman and Tversky's prospect theory (1979). It appears that for many investors the decision to sell a security is more influenced by what that security has done than by what it is likely to do.

Disposition theory predicts that investors will evaluate investments relative to a reference point or "break even" price. An investment sold for more than its reference point will be perceived as a gain. An investment sold for less will be perceived as a loss. Investors do not like to accept a loss so investments above the reference point are more likely to be sold than those below it. The reference point for an investment is sometimes assumed to be its purchase price. However for investments that have been held over a wide range of prices, purchase price may be only one determinant of the reference point. For example, a homeowner who bought his house for $100,000 just before a real-estate boom, and had the house appraised for $200,000 after the boom, may no longer feel he is "breaking even" if he sells his house for $100,000. Alternatively, suppose an investor buys a security at $20 a share. The share price falls over a few months to about $10 where it stays for the next year. If the share price then starts to rise rapidly, the investor may happily choose to sell for much less than $20, because his reference point has fallen below the original purchase price.

Suppose that reference points are moving averages (with some weighting function) of past prices.[15] When securities appreciate quickly they gain relative to their moving averages. A security that has lost value in recent months will probably be below its reference point. If the security rises rapidly over a few weeks, it might pass its reference point and thus become a candidate for a sale.

[15] Heath et al. (1998) find that the decision to exercise employee security options is a function maximum price of the underlying security over the previous year.

Attention focusing, the disposition effect, and the reluctance to sell short explain some of the security return patterns noted in Figures 1–6. These are patterns that precede sales and purchases. They are indications of the trading practices and preferences of investors. It is useful to understand these patterns, but it is not surprising that they exist. The patterns that are surprising to find are those that follow purchases and sales. These patterns indicate that these investors are informed but misuse their information. In Figures 1 and 2 the securities investors buy underperform those they sell. When these trades are partitioned into purchases and sales of previous winners (Figure 3), the previous winners investors buy underperform those they sell. These winners have been outperforming the market for at least two years prior to being purchased. After purchase they underperform the market.

It is possible that the return pattern for previous winners is caused by investors who buy at the top of a momentum cycle. Narasimhan Jegadeesh and Sheridan Titman (1993) document momentum patterns in security returns. They sort securities into those which have performed well or poorly during six-month formation periods. In the subsequent year the securities that previously did well continue to outperform those that previously did poorly. After one year these trends reverse somewhat. John R. Nofsinger and Richard W. Sias (1999) find that the reversals are mostly confined to securities with high percentages of individual investor ownership. If the rise of momentum securities is, in part, driven by the purchases of "momentum traders," then, when the last momentum trader has taken his position, the rise may stall. If momentum traders have pushed price beyond underlying value then the price is likely to fall when new information becomes available. Individual investors who follow momentum strategies may be among the last momentum traders to buy these securities and among the first to suffer losses when trends reverse. Some of the underperformance of securities these investors buy relative to those they sell may be due to mistiming of momentum cycles.

The same reasoning would not necessarily apply on the down side. Investors who follow momentum strategies might not sell securities that have fallen simply because they do not already own these securities and they do not like to sell short. If they do own securities that have fallen they may choose not to sell them because of disposition effects (i.e., they do not like to realize losses).

These explanations for the return patterns found in these data are speculations. Further research is needed to understand why individual investors choose the securities they choose and why they choose so poorly.

Whenever it is suggested that investors behave suboptimally the question arises: "why don't they learn?" It is possible that they do learn, but slowly. Equity markets are noisy places to learn. Most of the inferences drawn in this paper could not be made with the sample sizes available to most investors. It is likely that many investors never make the sort of evaluative comparisons made here. They do not, for example, routinely look up the performance of a security they sold several months ago and compare it to the performance of a security they bought in its stead. The disposition effect, too, may slow learning. Investors tend to sell winning investments and hold on to losers. If they weigh realized gains more heavily than "paper" losses when evaluating their personal performance, they may feel they are doing better than they are. During the seven years covered by the data, 55 percent of the original accounts drop out of the sample. About half of these drop in the first year, perhaps as a response to the market crash of October 1987. While there are many reasons to close an account, some investors may have closed their accounts because they did learn that they were not as good at picking securities as they had anticipated.

In aggregate the investors in this study make trading choices which lead to below-market returns. This does not mean these investors lose money. 1987 through 1993 were good years to be in the stock market and most of these investors are probably happy that they were.

The discount brokerage customers in this study make some poor trading choices. Other groups of traders make bad choices as well. Jensen (1968), Lakonishok et al. (1992), and Burton G. Malkiel (1995) show that active money managers underperform relevant market indices. While this may indicate poor judgment, agency considerations could also motivate active managers to make choices they would not otherwise make. Investors with discount brokerage accounts are studied in this paper for two reasons. First, a discount brokerage firm was generous enough to make the data available. Second, discount customers trade mostly for themselves and without agency concerns; they are therefore well suited for testing behavioral theories of finance. It would be instructive to repeat this study for other groups of traders.

This is a study of the trading of individual investors with discount brokerage accounts. What effect, if any, the trading of these investors will have on market prices will also depend on the trading of other market participants who may follow very different trading practices.

V. Conclusion

This paper takes a first step towards demonstrating that overall trading volume in equity markets is excessive by showing that it is excessive for a particular group of investors: those with discount brokerage accounts. These investors trade excessively in the sense that their returns are, on average, reduced through trading. Even after eliminating most trades that might be motivated by liquidity demands, tax-loss selling, portfolio rebalancing, or a move to lower-risk securities, trading still lowers returns. I test the hypothesis that investors trade excessively because they are overconfident. Overconfident investors may trade even when their expected gains through trading are not enough to offset trading costs. In fact, even when trading costs are ignored, these investors actually lower their returns through trading. This result is more extreme than is predicted by overconfidence alone.

I examine return patterns before and after the purchases and sales made by these investors. The investors tend to buy securities that have risen or fallen more over the previous six months than the securities they sell. They sell securities that have, on average, risen rapidly in recent weeks. And they sell far more previous winners than losers. I suggest that these patterns can be explained by the difficulty of evaluating the large number of securities available for investors to buy, by investors' tendency to let their attention be directed by outside sources such as the financial media, by the disposition effect, and by investors' reluctance to sell short. Return patterns after purchases and sales are more difficult to understand. It is possible that some of these investors are among the last buyers to contribute to the rise of overvalued momentum securities and are among the first to suffer losses when these securities decline. What is more certain is that these investors do have useful information which they are somehow misinterpreting.

Appendix

I use a Monte Carlo simulation to test the hypothesis that investors vary in their propensity to buy previous winners and previous losers. Two test statistics are employed: the proportion of accounts buying only previous winners or only previous losers, and the average of $|N_w - N_l|$ where N_w and N_l are the number of previous winners and previous losers purchased in an account. These two statistics are first calculated from the data and then simulated under the null hypothesis that each investor has the same probability of buying a previous winner as every other investor. For the simulation the probability of buying a previous winner is set to be the empirically observed ratio of previous winners bought to previous winners plus previous losers bought. Observations are taken only from accounts with more than one purchase of a previous winner or previous loser. For each account the same number of simulated purchases are generated as are observed in the sample. Each simulated purchase is drawn as either a previous winner or previous loser. When simulated purchases have been drawn for each account the two test statistics are calculated. This process is repeated 1,000 times and for each test statistic the 1,000 observations constitute a simulated distribution. When previous winners (losers) are simply defined to be securities which had a positive (negative) return over the six months prior to purchase (as in Figures 3 and 4), the average number of purchases per account is 8.4. The fraction of accounts buying only previous winners or previous losers is 0.265, while in the 1,000 simulations the largest fraction of accounts buying only previous winners or previous losers is 0.252. In the actual data $|N_w - N_l|$ is 3.6 while in the 1,000 simulations the largest value of $|N_w - N_l|$ is 2.6. Using either statistic we can reject the null hypothesis that each investor has the same propensity for buying winners and losers ($p < 0.001$). If big winners are defined to be securities that returned 60 percent or more in the previous six months (about the average in Figure 3) and big losers are those that returned -40 percent or less (about the average in Figure 4), then of the 1,197 investors who bought more than one big winner or big loser (4.5 such purchases on average), 555 bought only big winners or only big losers. In 1,000 simulations based on the assumption that all investors had the same probability as each other for buying big winners (or big losers), at most 457 bought only winners or only losers. The null hypothesis that each investor has the same propensity for buying big winners and big losers can be rejected ($p < 0.001$).

REFERENCES

Alpert, Marc and Raiffa, Howard. "A Progress Report on the Training of Probability Assessors," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press, 1982, pp. 294–305.

Badrinath, S. G. and Lewellen, Wilbur G. "Evidence on Tax-Motivated Securities Trading Behavior." *Journal of Finance*, March 1991, *46*(1), pp. 369–82.

Barber, Brad M. and Lyon, John D. "Firm Size, Book-to-Market Ratio, and Security Returns: A Holdout Sample of Financial Firms." *Journal of Finance*, June 1997, *52*(2), pp. 875–83.

Barber, Brad M.; Lyon, John D. and Tsai, Chih-Ling. "Holding Size While Improving Power in Tests of Long-Run Abnormal Stock Returns." *Journal of Finance*, February 1999, *54*(1), pp. 165–202.

Barber, Brad M. and Odean, Terrance. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." *Journal of Finance*, 1999a (forthcoming).

______. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." Working paper, University of California, Davis, 1999b.

Benartzi, Schlomo and Thaler, Richard H. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, February 1995, *110*(1), pp. 73–92.

Benos, Alexandros V. "Aggressiveness and Survival of Overconfident Traders." *Journal of Financial Markets*, October 1998, *1*(3–4), pp. 353–83.

Brock, William; Lakonishok, Josef and LeBaron, Blake. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance*, December 1992, *47*(5), pp. 1731–64.

Caballé, Jordi and Sákovics, József. "Overconfident Speculation with Imperfect Competition." Working paper, Universitat Autnoma de Barcelona, Spain, 1998.

Carhart, Mark M. "On Persistence in Mutual Fund Performance." *Journal of Finance*, March 1997, *52*(1), pp. 57–82.

Daniel, Kent; Hirshleifer, David and Subrahmanyam, Avanidhar. "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance*, December 1998, *53*(6), pp. 1839–86.

De Long, J. Bradford; Shleifer, Andrei; Summers, Lawrence H. and Waldmann, Robert J. "The Survival of Noise Traders in Financial Markets." *Journal of Business*, January 1991, *64*(1), pp. 1–19.

Dow, James and Gorton, Gary. "Noise Trading, Delegated Portfolio Management, and Economic Welfare." *Journal of Political Economy*, October 1997, *105*(5), pp. 1024–50.

Fama, Eugene F. and French, Kenneth R. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, February 1993, *33*(1), pp. 3–56.

Fischhoff, Baruch; Slovic, Paul and Lichtenstein, Sarah. "Knowing with Certainty: The Appropriateness of Extreme Confidence." *Journal of Experimental Psychology*, November 1997, *3*(4), pp. 552–64.

Frank, Jerome D. "Some Psychological Determinants of the Level of Aspiration." *American Journal of Psychology*, April 1935, *47*(2), pp. 285–93.

Gervais, Simon and Odean, Terrance. "Learning to Be Overconfident." Working paper, University of Pennsylvania, 1999.

Griffin, Dale and Tversky, Amos. "The Weighing of Evidence and the Determinants of Confidence." *Cognitive Psychology*, July 1992, *24*(3), pp. 411–35.

Grossman, Sanford J. and Stiglitz, Joseph E. "On the Impossibility of Informationally Efficient

Markets." *American Economic Review*, June 1980, *70*(3), pp. 393–408.

Heisler, Jeffrey. "Loss Aversion in a Futures Market: An Empirical Test." *Review of Futures Markets*, June 1994, *13*(3), pp. 793–822.

Heath, Chip; Huddart, Steven and Lang, Mark. "Psychological Factors and Security Option Exercise." *Quarterly Journal of Economics*, May 1999, *114*(2), pp. 601–27.

Ibbotson Associates. *Stocks, bonds, bills, and inflation: 1997 yearbook.* Chicago: Ibbotson Associates, 1997.

Ikenberry, David L.; Lakonishok, Josef and Vermaelen, Theo. "Market Underreaction to Open Market Share Repurchases." *Journal of Financial Economics*, October 1995, *39*(2–3), pp. 181–208.

Jegadeesh, Narasimhan and Titman, Sheridan. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, March 1993, *48*(1), pp. 65–91.

Jensen, Michael C. "The Performance of Mutual Funds in the Period 1945–1964." *Journal of Finance*, May 1968, *23*(2), pp. 389–416.

______. "Risk, the Pricing of Capital Assets, and Evaluation of Investment Portfolios." *Journal of Business*, April 1969, *42*(2), pp. 167–247.

Kahneman, Daniel and Tversky, Amos. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, March 1979, *47*(2), pp. 263–91.

Kyle, Albert S. and Wang, F. Albert. "Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test?" *Journal of Finance*, December 1997, *52*(5), pp. 2073–90.

Lakonishok, Josef; Shleifer, Andrei and Vishny, Robert W. "The Structure and Performance of the Money Management Industry." *Brookings Papers on Economic Activity*, Microeconomics 1992, pp. 339–79.

Lakonishok, Josef and Smidt, Seymour. "Volume for Winners and Losers: Taxation and Other Motives for Stock Trading." *Journal of Finance*, September 1986, *41*(4), pp. 951–74.

Langer, Ellen J. and Roth, Jane. "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*, December 1975, *32*(6), pp. 951–55.

Lichtenstein, Sarah; Fischhoff, Baruch and Phillips, Lawrence. "Calibration of Probabilities: The State of the Art to 1980," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press, 1982, pp. 306–34.

Loughran, Tim and Ritter, Jay R. "The New Issues Puzzle." *Journal of Finance*, March 1995, *50*(1), pp. 23–51.

Malkiel, Burton G. "Returns from Investing in Equity Mutual Funds 1971 to 1991." *Journal of Finance*, June 1995, *50*(2), pp. 549–72.

Miller, Dale T. and Ross, Michael. "Self-Serving Biases in Attribution of Causality: Fact or Fiction?" *Psychological Bulletin*, March 1975, *82*(2), pp. 213–25.

Nofsinger, John R. and Sias, Richard W. "Herding and Feedback Trading by Institutional and Individual Investors." *Journal of Finance*, December 1999, *54*(6).

Odean, Terrance. "Volume, Volatility, Price, and Profit When All Traders Are Above Average." *Journal of Finance*, December 1998a, *53*(6), pp. 1887–934.

______. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, October 1998b, *53*(5), pp. 1775–98.

Ritter, Jay R. "The Long-Run Performance of Initial Public Offerings." *Journal of Finance*, March 1991, *46*(1), pp. 3–27.

Schlarbaum, Gary G.; Lewellen, Wilbur and Lease, Ronald C. "Realized Returns on Common Stock Investments: The Experience of Individual Investors." *Journal of Business*, April 1978, *51*(2), pp. 299–325.

Shefrin, Hersh and Statman, Meir. "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence." *Journal of Finance*, July 1985, *40*(3), pp. 777–90.