

The Individual Day Trader*

Juhani Linnainmaa†

The Anderson School at UCLA

November 2005

*I especially thank Mark Grinblatt and Matti Keloharju for their numerous comments and suggestions. I have also benefited from the comments by Michael Brennan, Monica Piazzesi, and Mark Seasholes (AFA discussant). I also thank the participants at AFA 2004 meetings for helpful suggestions. An earlier version of this paper was circulated under the title “The Anatomy of Day Traders.” I gratefully acknowledge the financial support provided by the Foundation of Emil Aaltonen. Henri Bergström from the Finnish Central Securities Depository and Pekka Peiponen, Timo Kaski, and Jani Ahola from the Helsinki Exchanges provided data used in this study. Matti Keloharju provided invaluable help with the FCSD registry. All errors are mine.

†Correspondence Information: Juhani Linnainmaa, The Anderson School at UCLA, Suite C4.01, 110 Westwood Plaza, CA 90095-1481, tel: (310) 909-4666, fax: (310) 206-5455, <http://personal.anderson.ucla.edu/juhani.linnainmaa>, email: julinnai@anderson.ucla.edu.

The Individual Day Trader

Abstract

This paper shows that individual day traders are reluctant to close losing day trades. They even sell other stocks from their portfolios to finance the unintended purchases. This disposition to ride losers has significant long-term welfare consequences. Day traders hurt their portfolios' performance up to -6% in three months after a holdings change. The changes in individuals' exposure to market-wide shocks cause this underperformance: individuals systematically migrate towards small technology stocks with low B/M ratios. We find a negative relation between day trading profits and long-term performance: active day traders have the highest day trading profits but they hurt their long-term performance the most. Our results suggest that behavioral biases can push investors towards portfolios they might feel uncomfortable holding under other circumstances.

“It seems incredible that knowing the game as well as I did... I did precisely the wrong thing. The cotton showed me a loss and I kept it. The wheat showed me a profit and I sold it out. It was an utterly foolish play... Of all speculative blunders there are few greater than trying to average a losing game.”

—Edwin Lefévre (1923). *Reminiscences of a Stock Operator.*

1 Introduction

Day traders have attracted a significant amount of attention in the popular press. A search of the LexisNexis Academic service reports over 2,500 articles on day traders. A popular Internet bookstore lists over 100 titles for aspiring day traders, including at least one murder mystery, Stephen W. Frey’s *The Day Trader*, which the *New York Times* (February 10, 2002) calls “*fast-paced*.” Academic research has not addressed this phenomenon, not because of a lack of interest, but rather a lack of data. Barber and Odean (2001) point out that “*little is known about [day traders’] trading strategies, because firms that cater to day traders have been generally reluctant to provide access to the trading records of their clients.*”

We address this gap by examining the complete trading records of all *individual* (non-professional) day traders in Finland. We study individual day traders’ reluctance to close day trades at a loss because such a behavior could have important welfare consequences: if individuals systematically end up holding stocks they initially chose for short-term speculation, their portfolios may become something they would not hold under different circumstances. This paper examines this story: do individuals hold on to losing day trades and, if so, does this behavior affect their welfare? This paper’s focus on the welfare consequences of behavioral biases has implications beyond day traders: it would be invaluable to show that behavioral biases can push individuals towards “unsuitable” (i.e., too risky) portfolios.¹

We find that individual day traders not only hold on to losers but also sell their other holdings to extend their intended one-day trading horizons. This behavior has a dramatic impact on day traders’ long-term performance. The individual day traders’ *unintended* port-

¹Note that this possibility is different from the usual studies of the disposition effect where the differences between an investor’s holdings are not considered beyond their unrealized gains and losses. (Section 1.1 discusses the literature.) The possibility we consider is that the stocks that investors end up holding are different (for example, riskier) from the ones they let go.

folio changes are consistently poor: we find raw underperformance from -2% (less active day traders) to -6% (more active day traders) in three months between the new and old portfolios. Individual day traders' systematic migration away from the general stock market and towards small technology stocks with low B/M ratios accounts for most of the long-term underperformance. Day traders' exposure to the market-wide shocks changes over time because of the disposition effect.

An important methodological finding is that the short-term day trading profits and long-term performance are *negatively* related. The short-term profits are increasing in (earlier) day trading activity² but the opposite is true for the day traders' long-term performance: the active day traders hurt their long-term performance *more* than the casual day traders. This shows that an analysis of day trading profits alone overestimates individual day traders' performance because of the joint behavior of the disposition effect, day trading profits, and long-term performance. Although we study individual day traders, our results pertain to a broad class of individuals: those who approach the stock market as an opportunity for a quick profit. This paper finds that the pursuit for a quick profit combined with the unwillingness to accept losses has significant detrimental consequences.

The rest of the paper is organized as follows. After discussing related research, we describe our data and sample in Section 2. Section 3 contains the day trader disposition effect tests. The performance of day traders is analyzed in Section 4. Section 5 concludes.

1.1 Related Literature

The Disposition Effect. Shefrin and Statman (1985), Odean (1998), Grinblatt and Keloharju (2001) and many others have shown that investors are reluctant to realize losses. This description of individuals' behavior has its roots in the prospect theory of Kahneman and Tversky (1979). This study is different from these other studies because we study *intraday* loss aversion: *when and under what circumstances* do individual day traders abort their day trades. Only other study that partly shares this focus is Locke and Mann (2005): they show that professional traders also exhibit disposition effect. However, in contrast to our findings,

²The average day trader loses money even before brokerage fees but more active day traders are profitable before these fees. The typical day trader loses €100 (approx. \$120) per day after paying the brokerage firm fees. This loss may be a fair compensation for the activity's entertainment value. For example, *The Las Vegas Convention and Visitors Authority* reports that the average gambling budget per trip to Vegas was \$545 in 2004.

the authors find no costs associated with this behavior. Moreover, Locke and Mann's professional traders always close their positions by the end of the day whereas our key result is that individual day traders often go through considerable trouble to keep their losing shares more permanently.

Short-Term Traders. Harris and Schultz (1998) analyze SOES (Small Order Execution System of NASDAQ) bandits who try to profit from the fragmented order flow. These traders earn a small average profit per trade. Garvey and Murphy (2001) analyze a sample of 15 professional day traders who day trade on behalf of a firm for a profit share with the firm. These individuals earn day trading profits by placing limit orders inside the Nasdaq dealers' bid-ask spreads. Seasholes and Wu (2004) study a small sample of very active day traders on the Shanghai Stock Exchange and find that these day traders make money by buying shares that hit their upper price limits and selling them the next day. Jordan and Diltz (2002) examine a sample of *professional* U.S. day traders and find that most of them (two-thirds) lose money—only 1/5 of their sample's day traders are marginally profitable. The authors also find that day trading profits are related to Nasdaq movements. This parallels our finding that individual day traders tilt their portfolio holdings towards small technology stocks.

Barber, Lee, Liu, and Odean (2004) study Taiwanese day traders. They find that the most active day traders perform the best. About 20% of their day traders break even during a six-month period after transaction costs.³ A companion paper to this study by Linnainmaa (2005) uses the same data set as we do and shows that many individuals may day trade because they are uncertain about the profitability of day trading: they day trade and learn from experience. (The commonly suggested alternative to this hypothesis is that day traders are permanently stubborn—i.e., overconfident—about their own abilities: they day trade, lose and day trade again.) Individual day traders' behavior is consistent with the learning hypothesis. For example, day trade outcomes affect trade size and exit decisions even after controlling for the outcomes' wealth effects. Our study differs from these studies in its focus: we identify a strong disposition effect specific to day traders and show that this behavior has significant welfare consequences.

³A possible caveat in the Barber et al. (2004) study is their focus on day trading profits alone. Also Taiwanese day traders appear to ride losers: the authors report that 36% of all day trades in their sample cause changes in portfolio holdings (partial day trades). One of the results we highlight in this paper is the *negative* relation between short- and long-term performance, caused by day traders' tendency to abort losing day trades.

2 Data and Sample

This section describes the institutional setting of the Finnish market and our data sets. We also construct the sample used in our tests and define *day traders*.

2.1 Helsinki Exchanges

Trading on the Helsinki Exchanges (HEX) is divided into sessions. Each trading day starts at 10:10 am with an opening call. Orders that are not executed at the opening remain on the book and form the basis for the continuous trading session. This trading session takes place between 10:30 am and 5:30 pm in a fully automated limit order book, the automated trading and information system (HETI). After-hours trading (5:30 – 5:45 pm) takes place after the continuous trading session and again the next morning (9:30 – 10:00 am) before the next opening call. (Two changes to the trading schedule were made during the sample period. On August 31, 2000, the regular trading session was extended to 6:00 pm and the after-hours session was moved to match this change. On April 10, 2001, an evening session that extended trading hours to 9:00 pm was introduced.)

The HEX trading system displays the five best price levels of the book to the market participants. The public can view this book in a market-by-price form while financial institutions receive market-by-order feed.⁴ Simple rules govern trading on the limit order book. There are no designated market makers or specialists; the market is completely order-driven. An investor trades by submitting limit orders. The minimum tick size is €0.01. An investor who wants immediate execution must place the order at the best price level on the opposite side of the book—consistent with the market convention, we call these orders *market orders*. An investor who wants to buy or sell more shares than what is currently outstanding at the best price level must “walk up or down the book” by submitting separate orders for each price level. If a limit order executes against a smaller order, the unfilled portion stays on the book as a new order. Time and price priority between limit orders is enforced. For example, if an investor submits a buy order at a price level that already has other buy orders outstanding, all the old orders must execute before the new order.

⁴A market-by-price book displays the five levels on both sides of the market but only indicates the total number of shares outstanding at each price level. A market-by-order book shows each order separately and also shows which broker/dealer submitted each order.

The total market value of the 158 companies in the Helsinki Exchanges was €383.14 billion in the middle of the sample period (May 2000). The average realized log-spread during the sample period was 0.43% for the 30 largest stocks and as a low as 0.13% for Nokia, the most actively traded company. Brokers differ in the fees they charge. The most popular online broker charged more active trades a fixed fee of €8.25 and 0.2% of the trade value during the sample period. There was also a flat 28% capital gains tax during the sample period.

2.2 Data

Our data are (1) the complete trading records and holdings information of all Finnish investors, (2) the limit order data for all the stocks listed on the HEX, and (3) transaction data for the period not covered by the limit order data.

1. The Finnish Central Securities Depository registry (FCSD) provided us the investor holdings and trading records for the period from January 1995 to November 2002. Each record includes a date-stamp, price, volume, stock code, a code that identifies the investor type—a domestic institution, a domestic household, or a foreigner—and other demographic information. We classify investors as individuals or institutions for this study. Grinblatt and Keloharju (2000) give the full details of this data set.
2. The limit order data are the supervisory files from the HEX from September 18, 1998 to October 23, 2001. Each entry in this data set is a single order entered into the trading system. Each entry contains a unique order ID, an entry date and time-stamp, a session code, a trade type indicator (i.e., upstairs/downstairs/odd-lot), limit price, volume, the brokerage firm identity, and a set of codes to track the life of the order—an order can expire, be partly or completely filled, be withdrawn, or have its price amended. We use these data to reconstruct the limit order book for each second of every trading day for all the stocks. Data before July 10, 2001 is missing the time-stamp that identifies when an unfilled order is withdrawn.
3. The transactions data the derived from older supervisory files from the HEX for a period from January 1995 to September 17, 1998. Each entry in this data set is a single trade. Each entry contains a unique trade identifier, the price, volume, and brokerage firm identities for the both sides of the trade.

2.2.1 Matching the Data Sets

We match the investor trading records against the limit order data using executed trades to obtain information on, e.g., what type of orders different investors use and when trades take place. Each trade record in the limit order data contains all the same information as the investor trading records *except* the investor identity. We use common elements to link the data sets.

There is no ambiguity in matching two types of trades: trades with unique price-volume combinations and non-unique trades that must originate from the same investor.⁵ We call these trades *uniquely matched trades*. There is no one-to-one link between the data sets for the remaining trades. However, it is often possible to determine almost exactly *when* a trade took place. For example, even if there are two trades with a price/volume combination of 40.1 euros/200 shares, the trades may have occurred almost simultaneously. We follow this idea and compute the lower and upper bound for trade time (i.e., upper and lower time-stamps) for each non-unique entry in the investor trading records. Each investor trading record now contains either the exact time (with the limit order book data for the latter period) of the trade or at least bounds for the time of the trade.

2.2.2 Data Requirements and Subsamples

We use different samples throughout this paper depending on the data requirements:⁶

- Tests that require only daily data use the whole sample period. Exceptions are tests where we first compute daily averages and then take the mean across these first-stage averages. We drop observations before January 1998 in these tests because the relatively small number of day trades before this date. (Day trading became popular in Finland during 1998; see Figure 1 in Linnainmaa 2005).
- Tests that require market and limit order information use the limit order data sample

⁵We say that a trade has a *unique* price-volume combination if, for example, there is only one trade (in one stock-day) with a price of €82 and a volume of 1,200 shares. A trade is non-unique if, say, three trades have the same price-volume combination. In this example, “all must originate from the same investor” would mean that a single investor in the investor data set is the buyer or the seller in all the three trades.

⁶We emphasize that all our results—both the results on the disposition effect as well as the performance results—are both economically and statistically very significant. They are not caused by, e.g., a careful choice of specific subsamples for each test.

period.

- Tests that need to determine *when* individual trades took place use all observations for which the difference between the upper and lower time-stamps is less than ten minutes.
- Tests that use information on *the sequence* of a day trader’s trades during the day—i.e., whether the day trader sold the shares in the morning and then repurchased them later, or whether it was the other way around—use all unambiguous observations. There is no ambiguity when the upper and lower time-stamps show with certainty the direction of the day trade.

2.3 Identifying Individual Day Traders

A *day trade* is a purchase and a sale of the same stock (in any order) on the same day. If the amounts purchased and sold are the same the day trade is *complete*. If the amounts differ, the day trade is *partial*. A *day trader* is an individual investor who day trades at least once.

It is important to emphasize several issues about this definition. First, we focus on individual investors who *act* as day traders. We *do not* limit our analysis to individuals who day trade more professionally. This definition is somewhat different from the industry definition. For example, the NASD and SEC define a “day trader” as a professional with the primary goal of earning a living through the day trading profits.⁷ The purpose of this study is to analyze the behavior and performance of individual day traders, no matter how occasional their day trading may be. Second, although our criterion for what constitutes a day trader—a single day trade is enough—appears relaxed, we make sure it is inconsequential in our tests. For example, if some of our “day trades” are, e.g., trading errors and not genuine short-term speculation, our tests are more conservative. More importantly, our analysis of day trader *performance* takes a cautious approach and uses non-overlapping periods: an identification period and a subsequent performance evaluation period. This guarantees that our empirical results do not suffer from an endogeneity problem.

⁷In testifying in front of the Permanent Subcommittee on Investigations of the Senate Committee on Governmental Affairs in 1999, President of NASD Regulation Mary Schapiro said “...the term “day trading,” as commonly used within the industry, generally refers to the trading activities of the “professional day trader,” that is an individual who conducts intra-day trading in a focused and consistent manner, with the primary goal of earning a living through the profits derived from this trading strategy.” The SEC report “Report of Examinations of Day-Trading Broker-Dealer” (February 25, 2000) uses a similar definition.

There are 1,055,505 individual investors in the FCSD registry with at least some asset holdings during the sample period. A total of 22,529 investors day trade at least once; 21.8% of them day trade at least 10 times and 6.0% day trade at least 50 times. The total number of day trades is 300,894; 44.1% of these day trades are *partial*. While the average number of day trades per investor is 13.4, the median number is only two: this shows that the sample consists of a large number of relatively inactive day traders but also contains some very active day traders. For example, the most active day trader in the sample day trades 1,715 times. Linnainmaa (2005) discusses the demographics of the same sample of individual day traders. For example, these individuals are predominantly male, approximately ten years younger than the rest of the investor population, and often have little investment experience before the start of their day trading careers. The investors we classify as day traders are also otherwise very active. The total number of trades from them is 4,419,077. This represents 49.0% of the total order-flow from all Finnish individuals and 63.8% of the total turnover during the sample period.

3 Day Traders' Disposition to Ride Losers

This section examines whether day traders abort day trades when faced with a loss. We test the following hypotheses:

Hypothesis 1. *Individuals' speculation with short-term price movements is the main reason for observing day trades, not individuals' use of target prices.*

Hypothesis 2. *Day trades that an investor has to close (e.g., because of the impossibility of keeping a short position open overnight) have lower realized returns than day trades where the investor has discretion to keep the position open.*

Hypothesis 3. *Individuals abort day trades when the stock price moves against them.*

Hypothesis 4. *Individuals take losses from failed day trades close to the end of the trading day.*

The first hypothesis, which we call the “no target prices” hypothesis, is very important for our story. For example, it is possible that investors’ use of target prices generates *all* the day trades in our sample: an investor purchases shares and sets a price target, selling the

shares when this target is crossed. If this is a valid description of how investors behave, a day trade would arise from a same-day crossing of the target price. The second hypothesis tests whether day trade returns are different depending on how constrained the investor is. If day traders prefer to leave their losing positions open whenever they can, day trades that are more difficult to keep open have lower *realized* returns.⁸ The third and fourth hypotheses are predictions of the disposition effect: whether day traders abort day trades and whether they delay taking losses.

3.1 Individuals' Use of Target Prices

This section distinguishes between two possible stories for why we observe same-day round-trips in the data. First, these may be genuine short-term speculation. Second, these may result from individual investors' use of target prices (Schlarbaum, Lewellen, and Lease 1978). These alternatives make different predictions about same-day round-trip trades:

1. The “target price” scenario:
 - Investors *never* realize losses from same-day round-trip trades
 - The probability that a same-day round-trip trade is completed is monotonically increasing in the holding-period return

The second prediction says that if investors use target prices, a higher return is more likely to breach the (unobserved) target price. For example, suppose that individuals' target prices are uniformly distributed between 0 and $k \gg 0$. Then, the probability of observing a round-trip trade is identically zero for negative returns and linearly ($1/k$) increasing in the domain of positive returns.

2. The “genuine day trades” scenario:
 - Investors realize losses from the same-day round-trip trades
 - The probability that the investor sells the shares on the same day jumps up at the loss/gain point (i.e., at zero net return).

⁸An important assumption underlying this hypothesis is that most day traders do not possess private information: i.e., individuals do not engage in riskier (constrained) strategies only when they have private information.

The first prediction is that day traders sometimes have to take losses from their day trades. The second prediction follows from the prospect theory of Kahneman and Tversky (1979) where agents treat gains and losses differently.

3.1.1 Methodology

We examine how often investors complete same-day round-trip trades conditional on the realized or unrealized same-day return on the purchase. First, we construct a sample where each observation is a single investor-stock-day with a purchase of shares. Each observation contains the number of shares purchased and sold, together with the average purchase, sale, and same-day closing prices. We limit the sample to observations where the investor does not previously own any shares in the company. Second, we compute two types of returns: we compute the realized sell-buy spread for round-trip observations and the unrealized same-day return for “purchases only” observations. We define these returns as:

$$\text{realized sell-buy spread} = \ln P_s - \ln P_b \quad (1)$$

$$\text{unrealized same-day return} = \begin{cases} \ln c - \ln P_b & \text{if } V_b > V_s \\ \ln P_s - \ln c & \text{if } V_s > V_b \end{cases} \quad (2)$$

where P_b and P_s are the average purchase and sale prices, V_b and V_s are the number of shares purchased and sold, and c is the same-day closing price.⁹ Eq. 2 assumes that an investor could close his position at a price close to the same-day closing price.¹⁰ Third, we assign all investor-stock-day observations into 22 bins using this return variable and compute what proportion of purchases in each bin are round-trip trades.

3.1.2 Results

Figure 1 shows that we investors’ use of target prices is not the source of day trades. Two strong patterns support this conclusion. First, individuals often take losses from day trades. *None* of these loss-generating day trades—a total of 51,598 day trades, or 1/3 of all day trades—can be explained by the target price proposition. Second, instead of a monotonic

⁹Note that only the $V_b > V_s$ case is used here to compute unrealized same-day returns. We state the definition for the other case (“more sold than purchased”) for future reference.

¹⁰We ignore brokerage firm fees throughout Section 3. The spread is implicitly included because we use actual trade prices.

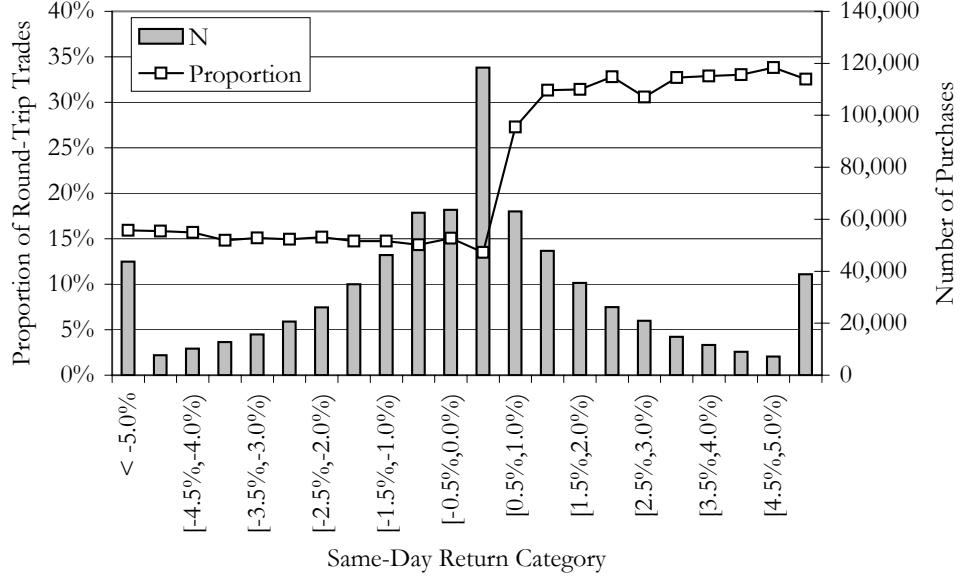


Figure 1: The Proportion of Round-Trip Trades Conditional on the Same-Day Return. This figure shows what fraction of purchases is reversed during the same day conditional on the same-day return. The sample consists of all investor-stock-day observations from investors with at least one round-trip trade between January 1995 and November 2002. We exclude investor-stock-day observations where the investor already owns shares in the company. The return on the x -axis is defined for round-trip trades as the log-difference between the average sale and purchase prices. The *unrealized* return for stock-days with only purchases is defined as the log-difference between the same-day closing price and the purchase price. The line reports what fraction of purchases become round-trip trades conditional on the unrealized or realized return. The lowest return bin contains purchases with $r < -5\%$ and the highest bin contains purchases with $r \geq 5\%$. The solid bars report the number of purchases in each return interval.

relation between the holding-period return and the probability of a round-trip, there is a jump up at the point when the investor can close his trade at profit. (Note that this turn to profitability happens approximately at 0.4%—not at 0%—because of commissions.) The prospect theory predicts precisely this type of behavior. Moreover, this is inconsistent with the target price hypothesis because the proportion of round-trip trades should be equal to zero up to this point and then slope upwards. These strong results against the target price hypothesis are very important because the rest the study rests upon the assumption that investors’ day trades reflect short-term speculation. Although *it is* possible that some day trades arise from the use of target prices, true speculation swamps such behavior in the data.¹¹

¹¹ As additional evidence against the target price hypothesis, note that day trading became popular in Finland during 1998 (Linnainmaa 2003). The number of day trades per year from 1997 to 2002 are: 3,180 (1997), 8,149 (1998), 23,180 (1999), 100,866 (2000), 93,037 (2001), and 68,191 (up to November 2002). This time-variation is difficult to explain by investors’ use of target prices; for example, the stock prices fell in 2001 whereas they rose in 1998 but there were 11.4 times more day trades in 2001.

3.2 Constraints to Close and Day Trading Returns

A day trader faces two types of situations where he may be forced to terminate the day trade against his will:

- *Short Positions.* If a day trader has taken a short position, he has to terminate the trade irrespective of its performance.
- *Liquidity Problems.* If a day trader has purchased shares (with intent of selling them the same day), the day trader can only abort the day trade if (1) he has enough liquid funds to finance the new purchase or (2) is willing to sell something else from his portfolio to finance the purchase.

This section studies whether the returns from completed day trades are consistent with day traders being affected by this “constraint mechanism”—i.e., whether investors keep losing day trades open if they can, biasing realized returns upwards. We predict that day trades that constrain the investor more have low *realized* returns. The intuition is that the returns from these constrained day trades are closer to the “true profitability of day trading” whereas the disposition effect contaminates other observations.

3.2.1 Methodology

We classify all day trades into nine categories:

<i>Complete Day Trades, $V_b = V_s$</i>	
1	OBS Own shares; Buy first, sell later
2	OSB Own shares; Sell first, buy later
3	DOBS Do not own shares; Buy first, sell later
4	DOSB Do not own shares; Sell first, buy later
<i>Partial Day Trades, $V_b \neq V_s$</i>	
5	OBS Own shares; Buy first, sell later
6	OSB Own shares; Sell first, buy later
7	DOBS Do not own shares; Buy first, sell later
8	DOSB Do not own shares; Sell first, buy later
9	ERROR Overnight short position

The last category, ERROR, contains day trades where the investor leaves open an overnight short position. This means that the investor is likely to incur a monetary penalty from the brokerage firm. We drop all day trade observations that cannot be classified (i.e., when the intraday sequence of purchases and sales is unknown because of the time-stamps).

We make the following predictions about realized returns. First, the DOSB strategy has the lowest return because it contains cases where an investor must cover a short position. Second, the OSB strategy has the highest return: an investor selling a stock from his portfolio can freely decide not to buy it back if the stock price increases after the sale. Third, the OBS and DOBS strategies fall in the middle: an investor's ability to keep the position open depends on the availability of liquid assets. The ERROR observations could have high or low returns:

- The returns are *high* if investors leave illegal short positions open when they are very confident that the monetary gain from an overnight fall in the stock price offsets the brokerage firm's monetary penalty.
- The returns are *low* if investors leave illegal short positions open when they do not have enough funds to cover the short position (i.e., the stock price has increased sharply), or are reluctant to close the position because of losses.

We analyze the relative performance of day trades by computing the average realized sell-buy spread (Eq. 1) for each category. We also test a hypothesis specific to partial day trades: these day trades may be instances where the investor has been unable to close the entire position at profit. We test this by computing the unrealized return (Eq. 2) for each partial day trade's residual position (i.e., $V_b - V_s$).¹²

3.2.2 Results

The results in Table 1 support the constraint hypothesis: the realized returns are lower when the day trader is more constrained. The relation is exactly as predicted for *complete* day trades: the OSB strategy has the highest return (2.39%) while the NOSB strategy has the lowest return (0.45%). The two strategies where the investor first buys a stock fall in the middle. The ordering is nearly correct also for *partial* day trades: NOBS has the lowest

¹²For example, if a day trader bought 300 shares but only sold 100, the residual position is +200 shares.

Table 1: Realized Day Trading Returns Conditional on Strategy

All day trades are classified into categories based on whether the investor already owns shares (**O**...) or not (**DO**...) and whether the investor initiates the day trade with a buy (...**BS**) or with a sell (...**SB**). We create separate categories for complete ($V_b = V_s$) and partial ($V_b \neq V_s$) day trades. The **ERROR** category contains day trades where a short position is left open overnight. Day trades that cannot be classified are dropped (see text). *Realized Sell-Buy Spread* is the log-spread between the average sale and purchases prices. *Unrealized Same-Day Return* is the log-spread between the same day close and the average purchase price or between the average sale price and the same day close, depending on the sign of the residual position ($V_b - V_s$). *% Positive* is the proportion of day trades where the spread or the same-day unrealized return is positive.

Panel A: <i>Realized Sell-Buy Spreads</i>					
	N	Mean	s.e.	Md	% Positive
<i>Complete Day Trades ($V_b = V_s$)</i>					
OBS	23,803	1.17%	0.02%	1.26%	68.5%
OSB	18,121	2.39%	0.03%	1.96%	78.6%
DOBS	56,597	0.99%	0.02%	1.12%	66.4%
DOSB	21,042	0.45%	0.03%	0.60%	58.8%
All	119,563	1.14%	0.01%	1.17%	67.3%
<i>Partial Day Trades ($V_b \neq V_s$)</i>					
OBS	21,065	0.15%	0.03%	0.31%	52.9%
OSB	21,904	1.57%	0.03%	1.27%	66.9%
DOBS	10,618	-0.15%	0.04%	0.03%	50.3%
DOSB	1,454	0.22%	0.12%	0.34%	54.5%
ERROR	1,726	0.65%	0.10%	0.56%	59.4%
All	56,767	0.66%	0.02%	0.67%	58.1%
Panel B: <i>Unrealized Same-Day Returns for Partial Day Trades</i>					
	N	Mean	s.e.	Md	% Positive
OBS	21,029	-0.36%	0.03%	-0.21%	43.2%
OSB	21,864	0.16%	0.03%	0.00%	50.5%
DOBS	10,599	-0.91%	0.05%	-0.56%	40.7%
DOSB	1,453	-0.60%	0.09%	-0.31%	40.6%
ERROR	1,703	-0.89%	0.11%	-0.62%	36.6%
All	56,648	-0.28%	0.02%	-0.17%	45.3%

return (-0.15%) and NOSB the second lowest (0.22%); OSB has the highest return (1.57%).

These results suggest that day traders leave their trades open when they can.

Panel B supports the hypothesis that partial day trades are partial precisely because the investor has been unable to sell the remaining shares at profit. The unrealized return is positive in 45% of all the cases (before brokerage fees) and only the least constrained day trade type (OSB) has a non-negative mean.¹³ Finally, day trades that result in an overnight

¹³In unreported work, we examine when individuals close their residual positions. The majority of these

short position (ERROR) perform very poorly: the residual position’s unrealized return is positive only in one out of three cases.

3.3 Intended Day Trades

3.3.1 Methodology

This section examines *intended day trades*: cases where an investor purchases shares with intent of later selling them the same day but aborts the day trade when the stock price falls. The investor may have to liquidate his *other* positions to finance the new, “unexpected” purchase because the investor did not expect to keep the position open. This type of behavior leaves its marks into the data: intended day trades are purchases with poor unrealized returns, accompanied by sales of other holdings.

Our strategy for testing for intended day trades is straightforward. First, because an investor’s day trades are concentrated,¹⁴ a purchase (without sales) soon after a day trade is *more likely* an intended day trade. Hence, we categorize purchases depending on how long it has been since the investor’s previous day trade and examine whether the purchases are “different” across bins. If there are intended day trades, we should observe (1) an upward sloping trend in the same-day unrealized returns and (2) a downward sloping trend in how often investors sell other shares to finance the new purchase.

We first classify all purchases based on the number of days from the investor’s previous day trade. We drop (1) purchases that more than three months away and (2) purchases before January 1998. We assign the remaining purchases into eight distance categories and compute the average unrealized same-day return (Eq. 2) for each day-bin. We also compute how often investors simultaneously sell their other holdings. This generates $2 * 8 = 16$ daily time-series with approximately 1,200 observations each. Finally, we compute the time-series means and standard errors for each distance bin.

positions (54%) is closed the next day; the average position is kept open for 6 trading days. The investors usually end up realizing a loss: the average holding-period return is -1.15% (a t -value of -13.9). These figures are computed from those partial day trades where an investor previously owned no shares in the company.

¹⁴An earlier version of this paper used the sample from Section 3.1 to examine the duration between day trades: “what is the probability of observing a day trade conditional on the log-number of trading days from the previous day trade?” The coefficient for this duration variable is -0.45 with a t -value of -223.0 , confirming that day trades are clustered.

Table 2: Intended Day Trades

This table examines purchases without same-day sales that take follow soon after a day trade by the same investor (“intended day trades”). We discard purchases more than three months after the previous day trade and those before January 1998. All remaining observations are classified into eight categories based on the number of trading days from the previous day trade. *Unrealized Same-Day Return* in Panel A is the log-difference between the closing price and the average purchase price. *Other Stocks Sold* in Panel B shows how often purchases are accompanied by a simultaneous sale of the investor’s other holdings. The average unrealized same-day return and the average “other stocks sold” proportion is first computed for each day-bin to create $2 * 8$ daily time-series. This table reports the time-series means and standard errors. N is the number of daily observations. The last line reports the pair-wise difference between the outermost and the innermost category.

Panel A: *Unrealized Same-Day Return*

Days from the Previous Day Trade	Mean	s.e.	N
[1]	-0.387%	0.037%	1,208
[2]	-0.281%	0.045%	1,185
[3]	-0.235%	0.044%	1,169
[4]	-0.203%	0.053%	1,157
[5]	-0.257%	0.073%	1,135
[6, 10]	-0.196%	0.045%	1,217
[11, 21]	-0.121%	0.041%	1,229
[22, 63]	-0.121%	0.039%	1,231
Difference [22, 63] – [1]	0.262%	0.027%	1,208

Panel B: *Other Stocks Sold*

Days from the Previous Day Trade	Mean	s.e.	N
[1]	55.6%	0.503%	1,212
[2]	46.9%	0.618%	1,185
[3]	43.4%	0.651%	1,172
[4]	40.6%	0.698%	1,159
[5]	38.4%	0.714%	1,138
[6, 10]	35.0%	0.455%	1,220
[11, 21]	31.0%	0.379%	1,230
[22, 63]	25.5%	0.286%	1,231
Difference [22, 63] – [1]	-30.1%	0.524%	1,212

3.3.2 Results

The results in Table 2 strongly support the existence of intended day trades. First, purchases soon after a day trade have both economically and statistically significantly lower unrealized same-day returns than later purchases. This is consistent with the proposition that early purchases are “intended but failed” speculation attempts: an investor aborts a day trade

because of a fall in the stock price. For example, a purchase one day after the previous day trade has an unrealized same-day return of -0.39% . In contrast, the average return on a purchase one to three months later is only -0.12% . (The time-series difference has a t -value of -9.6 .)

Second, purchases soon after a day trade are often accompanied by liquidation of an investor's other holdings. For example, sales of other holdings accompany "one day after the previous day trade" purchases 56% of the time. This proportion drops monotonically towards later purchases and is only one-fourths in the "one to three months" bin. This dramatic drop is both economically and statistically significant, and shows that the investors' disposition to ride losers can affect day traders' portfolios in two ways:

1. **The direct effect.** When a day trader aborts a day trade, he keeps shares originally not intended for long-term holding.
2. **The indirect effect.** When a day trader aborts a day trade, he often has to liquidate other positions to finance the new purchase.

The finding that individual day traders' liquidate other holdings to finance their new purchases is important also for another reason. It shows that our individual day traders' trading horizons are, indeed, just one day: the investors struggle to keep their new shares because their original intention was to trade at the margin.

3.4 Delayed Closing of Positions

3.4.1 Methodology

This section examines whether investors wait until the end of the day before taking losses from day trades. Two effects may push day traders to terminate losing day trades towards the end:

1. **Disposition to Ride Losers.** Investors speculating about short-term price movements wait as long as possible before accepting loss. For example, suppose that a day trader does not have any enough funds to keep the position open overnight. Then, the disposition to ride losers means that the day trader realizes losses at the last possible moment.

2. Market Making Strategy. Individuals can profit from day trading in two ways: by correctly predicting price movements or by offering liquidity to the market with limit orders. An individual who tries to act as a market maker may lose when he has to close his inventory at the end of the day with a market order.

We test whether the data supports the first mechanism while controlling for the second effect. Note that the intraday timing itself should not (exogenously) affect the profitability: day traders would otherwise migrate towards this high profitability period. We begin by constructing a sample of all trades (purchases and sales) that *terminate* day trades. For example, if an investor buys shares in the morning and sells them in the afternoon, we only include the sale into the sample. We call these observations *terminating trades*. We drop observations where the difference between the terminating trade's upper and lower time-stamps is more than ten minutes.

We then examine the delayed closing of losing day trades in two ways. In the first approach, we divide the trading day into five-minute intervals and compute what fraction of day trades in each interval result in a loss. We examine whether this proportion changes towards the end of the trading day as a first-pass analysis. The second approach estimates a logistic regression where the dependent variable is set to one if the day trader loses money. We include an intercept and the following variables on the RHS of the model:

- the log-number of seconds remaining until the close
- a dummy variable for a market-order initiated trade
- the log-spread at the time of the trade
- a dummy variable for the direction of the terminating day trade

(We also include several control variables; these variables are described in Table 3.) We also include an interaction term between the market order dummy and the spread. These variables control for the “market making” story. The disposition effect hypothesis predicts that the *Time Remaining* coefficient is negative.

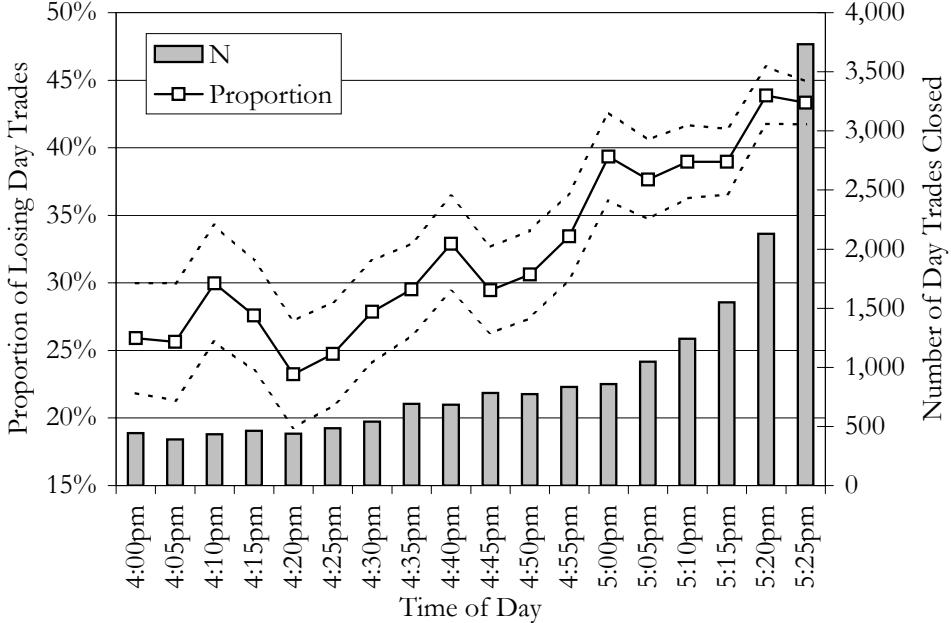


Figure 2: **Delayed Closing of Losing Day Trades.** This figure shows what fraction of the day trades closed towards the end of the trading day result in a loss. A *losing day trade* is a day trade where the average sale price is strictly lower than the average purchase price. The proportion of losing day trades is plotted for each five-minute interval for the last 90 minutes of the trading day together with the 95% confidence interval. The solid bars report the number day traders closed in each interval. The figure uses data on all day trades from individual investors for the period when the trading day ended at 5:30pm (i.e., before August 31, 2000).

3.4.2 Results

Figure 2 shows that the proportion of losing day traders increases sharply towards the end of the trading day. This proportion is approximately 25% an hour and a half before the trading session ends, but is almost 45% during the last ten minutes of the day. This sharp increase is both economically and statistically highly significant. The simultaneous sharp increase in the number of day trades closed is also worth noting. The overall effect is that investors close 16% of *all* losing day trades during the last five minutes of the trading day. In contrast, this fraction is only 9% for successful day trades.¹⁵

Table 3 shows that this result about delayed closing of day trades is robust to including all the controls. The coefficient estimate of -0.136 for the *Time Remaining*-variable indicates that the odds ratio for “a losing day trade” drops by 2/3 when the *terminating trade* is moved

¹⁵Figure 2 uses data on period when the trading day ended at 5:30 pm. The results for the periods with other trading schedules are very similar. For example, the proportion of losses increases from 27% to 48% in the period when the trading session ended at 6:00 pm.

Table 3: Delayed Execution of Day Trade Transactions

This table estimates a logistic regression where the dependent variable is set to one if the day trade results in a loss ($p_s < p_b$) and to zero otherwise. *Time Remaining* is the log-number of seconds remaining until the close + 1, *Market Order* is a dummy variable set to one if the day trader uses a market order, *Spread* is the log-spread at the time of the trade, *Partial Day Trade* is a dummy variable set to one if the purchase and sale volumes are different, *Ownership* is a dummy variable set to one if the day trader already owns shares in the company, *Sell Dummy* is a dummy variable set to one if the terminating trade is a sale, *Stock Return* is the return from the yesterday's close to today's close, and *Day Trader Experience* is the log-number of the investor's earlier day trades + 1. The sample consists of all terminating trades (see text). The data are the investor trading records uniquely matched against a limit order book data set. These data are from September 1998 to October 2001.

Variable	Coeff.	t-value
Intercept	0.072	0.9
Time Remaining	-0.136	-35.9
Transaction Specific		
Market Order	0.944	56.6
Spread	-23.789	-22.9
Market Order * Spread	25.468	22.0
ln(Trade Size)	0.025	4.1
Day Trade Type		
Partial Day Trade	0.584	34.6
Ownership	-0.845	-33.0
Sell Dummy	-0.079	-3.7
Ownership * Sell	0.565	18.6
Controls		
Stock Return	-1.896	-24.2
100 * (Stock Return) ²	0.001	0.7
Day Trader Experience	-0.062	-11.6
<i>N</i>	104,521	
Nagelkerke <i>R</i> ²	16.8%	

from the very end of the trading day to one hour before ($e^{\ln(3600)*(-0.136)} \approx 0.33$).

The results in Table 3 also show that day traders' ability make money is correlated with their ability to close a trade with a limit order. A day trade closed with a market order is often a failure; the predicted increase in the loss odds-ratio is 2.57 when the investor moves from a limit order to a market order. The interaction with the spread is, of course, important. For example, suppose that the log-spread is 2%. Then, the model predicts that the odds ratio for a loss increases by a factor of 4 if the investor moves from a limit order to a market order ($e^{0.944+25.468*0.02} \approx 4.28$).¹⁶

¹⁶The result that individual day traders are very active towards the end of the trading day is important.

4 Performance

We now examine individual day traders' performance at two horizons: short-term day trading profits and the long-term portfolio performance. If day traders were perfectly disciplined, short-term trading profits would measure day traders' abilities to make (or lose) money. However, Section 3 established that this is not the case: day traders not only abort day trades but also sell other holdings to keep the losers. This is important for the long-term performance if the stocks that investors choose for day trading are different from their previous holdings. If so, the act of day trading itself induces individuals to migrate towards riskier portfolios.

This section first studies whether individual day traders make money by day trading. We then examine these individuals' long-term performance and discuss the correlation between short- and long-term performances. Finally, we examine whether the stocks that investors day trade are fundamentally different from the stocks they previously own.

4.1 Classifying Day Traders for Performance Measurement

The way we identify day traders is important for the performance analysis. We use a two-step procedure to separate identification and performance measurement periods.

1. **The identification step.** We classify each investor as a day trader based on the investor's day trading activity during the previous three months. We reclassify all investors each day. We construct five groups: (1) all investors with any day trades, (2-4) investors with one to nine, ten to 19, and more than 20 day trades, and (5) investors with no day trades during the previous three months. (The fifth "no day trades" group is a reference group for the analysis of long-term performance.)
2. **The performance measurement step.** We measure each day trader's performance each day after reclassifying the day traders into the five groups.

For example, Campbell, Lettau, Malkiel, and Xu (2001) and Barber and Odean (2001) suggest that individual investors' short-term speculation (day trading) may have contributed to the increase in idiosyncratic volatility in the 1990s. Our result that individual speculators' delay closing their trades to the very end of the day is consistent with this possibility.

4.2 Individual Day Traders' Short-Term Performance

4.2.1 Measuring Day Trading Profits

We compute four different measures of day trading performance. First, we compute the daily **gross day trading profit** for each investor/stock/day as

$$\begin{aligned}\Pi_{i,t} = & P_{s,i,t}V_{s,i,t} - P_{b,i,t}V_{b,i,t} \\ & + \max(V_{b,i,t} - V_{s,i,t}, 0)(c_{i,t} - P_{b,i,t}) \\ & + \max(V_{s,i,t} - V_{b,i,t}, 0)(P_{s,i,t} - c_{i,t}).\end{aligned}\quad (3)$$

where i and t are stock and day indices. The first line is the profit for a complete a day trade while the last two lines adjust for the residual position. For example, if the investor bought 300 shares but only sold 200, we mark the residual 100 shares to the market at the closing price. Second, **net day trading profit** adjusts for transaction costs:

$$\hat{\Pi}_{i,t} = \Pi_{i,t} - 8.25 - 0.002(P_{b,i,t}V_{b,i,t} + P_{s,i,t}V_{s,i,t}). \quad (4)$$

This adjustment accounts for a fixed fee of €8.25 and a proportional fee of 0.2%. Finally, we compute **gross** and **net day trading returns** by dividing daily profits by the total turnover, $\sum_i(P_{b,i,t}V_{b,i,t} + P_{s,i,t}V_{s,i,t})$:

$$\pi_t = \frac{\sum_i \Pi_{i,t}}{\sum_i(P_{b,i,t}V_{b,i,t} + P_{s,i,t}V_{s,i,t})}, \quad \hat{\pi}_t = \frac{\sum_i \hat{\Pi}_{i,t}}{\sum_i(P_{b,i,t}V_{b,i,t} + P_{s,i,t}V_{s,i,t})}. \quad (5)$$

We measure day traders' performance by first computing the average of each performance measure separately for each day trader group-day. This generates $4 * 5 = 20$ daily time-series. We use time-series means and standard errors to measure individuals' day trading abilities. We also repeat this analysis at monthly level for trading profits. We drop pre-January 1998 observations from the analysis.

4.2.2 Results

Table 4 shows that the average individual day trader in our sample does not earn day trading profits. This is mostly true even *before* transaction costs: the typical day trader ($N_{DT} > 0$

Table 4: Day Trading Profits and Returns

This table reports time-series averages for day traders' trading profits and trading returns. The gross trading profit for an investor-stock-day is computed as

$$\Pi_{i,t} = P_{s,i,t}V_{s,i,t} - P_{b,i,t}V_{b,i,t} + \max(V_{b,i,t} - V_{s,i,t}, 0)(c_{i,t} - P_{b,i,t}) + \max(V_{s,i,t} - V_{b,i,t}, 0)(P_{s,i,t} - c_{i,t})$$

where P_b and P_s denote average purchase and sale prices, V_b and V_s denote the amounts purchased and sold, and c denotes the closing price. The net trading profit adjusts for transaction costs by subtracting a fixed fee of €8.25 and a proportional fee of 0.2%. Gross and net trading returns are computed by dividing each investor's daily profits by the turnover for the day. A daily time-series is constructed for each measure by averaging across investors. This table reports the time-series means and standard errors for investor groups based on the number of day trades during the past three months (N_{DT}); t -values are reported in parentheses. N is the average number of day traders in each group per day in the time-series. $\% Positive$ reports the average proportion of day traders with positive gross or net profits per day (z -value reported in parentheses). Panel B reports pair-wise t -tests between the groups. Panel C shows the average monthly profits.

Panel A: *Daily Profits and Returns*

Group	N	Profit (€)		Return		% Positive	
		Gross	Net	Gross	Net	Gross	Net
$N_{DT} > 0$	590.3	5.3 (0.5)	-102.4 (-8.9)	-0.21% (-1.3)	-0.76% (-4.6)	51.0% (4.4)	36.8% (-57.0)
$1 \leq N_{DT} < 10$	424.7	-15.9 (-1.8)	-88.1 (-9.8)	-0.25% (-1.3)	-0.86% (-4.5)	50.0% (-0.1)	35.6% (-61.4)
$10 \leq N_{DT} < 20$	80.9	57.1 (1.0)	-96.5 (-1.7)	-0.04% (-2.1)	-0.46% (-22.8)	52.7% (6.7)	39.4% (-26.2)
$20 \leq N_{DT}$	87.3	132.4 (2.4)	-216.7 (-3.9)	0.04% (2.4)	-0.32% (-17.7)	56.3% (14.3)	42.0% (-18.4)

Panel B: *Pair-Wise Differences (Daily)*

Comparison	Return	
	Gross	Net
Group 2 – Group 3	0.21% (1.1)	0.39% (2.0)
Group 3 – Group 4	0.08% (4.1)	0.14% (6.8)

where N_{DT} is the number of day trades during the past three months) gains €5.3 before transactions but loses €102.4 after these costs. The proportion of day traders with profits drops from 51.0% to 37% after adjusting for the costs. Hence, the typical day trader loses money by a considerable margin after adjusting for transaction costs. For the least active group of day traders ($0 < N_{DT} < 10$), even the gross profit is negative (-€15.9).

The most active group of day traders comes closer to being profitable: their average daily

Table 4: (cont'd)

Panel C: *Monthly Profits*

Group	N	Profit (€)		% Positive	
		Gross	Net	Gross	Net
$N_{DT} > 0$	1992.7	-13.0 (-0.2)	-652.3 (-9.0)	47.9% (-3.8)	30.3% (-34.7)
$1 \leq N_{DT} < 10$	1676.4	-128.4 (-2.2)	-519.1 (-8.6)	47.6% (-4.5)	30.7% (-34.7)
$10 \leq N_{DT} < 20$	177.2	527.2 (1.0)	-932.7 (-1.8)	48.5% (-1.4)	28.4% (-20.8)
$20 \leq N_{DT}$	139.2	1,916.3 (1.8)	-2,350.5 (-2.6)	52.0% (1.2)	29.9% (-10.2)

gross day trading profit is €132.4 which translates to a marginally positive gross return, 0.04% ($t = 2.4$). The proportion of day traders who make money before transaction costs is 56.3%. However, after accounting for the brokerage fees, these gains turn into statistically significant losses. The fact that the gross return is very close to zero is, however, more important. It shows that even if some day traders were able to negotiate lower commissions, our conclusion about their poor performance would not change. The results on monthly day trading profits in Panel C are similar. For example, the average day trader in the most active group generates a before-cost profit of almost €2,000 per month. However, after accounting for brokerage firm fees, this turns into a sizable loss. The estimates show that only 3 out of 10 day traders make money after transaction costs in a typical month.

4.3 Individual Day Traders' Long-Term Performance

4.3.1 Measuring Portfolio Performance

We measure the impact of day traders' disposition effect on portfolio performance with a variant of the Grinblatt and Titman (1993) self-benchmark measure (henceforth: the *GT* measure). This measure quantifies how *active* decisions affect performance. The idea is that if investors benefit from trading, there is a positive correlation between changes in portfolio weights and *future* returns. We modify the *GT* measure for our purposes and compute it as the difference between two returns: the return on the actual portfolio (*post-trade portfolio*) and the return on a lagged portfolio (*pre-trade portfolio*; i.e., the portfolio the investor held

k trading days ago). Individual day traders' post-trade portfolios outperform their pre-trade portfolios if individuals make good active trading decisions.

We measure daily portfolio returns as follows. First, we define the portfolio return as the close-to-close return for shares actually held from $t - 1$ to t . This means that we ignore the direct effects of trading to focus on the *portfolio* performance. More formally, let $x_{i,t}$ denote the number of shares in stock i the investor's portfolio at the end of day t . We compute the post-trade portfolio return as

$$r_t = \frac{\sum_i \min(x_{i,t-1}, x_{i,t})(c_{i,t} - c_{i,t-1})}{\sum_i \min(x_{i,t-1}, x_{i,t})c_{i,t-1}} \quad (6)$$

where $c_{i,t}$ is adjusted for any dividends so that the t and $t - 1$ closing prices are comparable. We compute the pre-trade, k -lagged portfolio return as

$$r_t^k = \frac{\sum_i x_{i,t-k-1}(c_{i,t} - c_{i,t-1})}{\sum_i x_{i,t-k-1}c_{i,t-1}}. \quad (7)$$

The GT measure is then the difference between these two returns,

$$GT_t^k \equiv r_t - r_t^k. \quad (8)$$

We drop two types of investor-day observations at this point from further analysis:

- Investor-day observations where the pre- and post-trade portfolios are identical (i.e., no trades or only complete day trades between dates $t - k$ and t).
- Investor-day observations where the investor is out of the market at date $t - k$ or $t - 1$.¹⁷

We compute the performance measure for one day, one week, one month, and three month horizons as follows. First, we compute the GT measure for each investor-day and then take the average GT measure for each day across investors. We use the time-series means and their standard errors to measure investors' long-term performance. We again drop pre-January

¹⁷This GT measure is conditional on there being portfolio changes; the appropriate interpretation is, "how much better or worse did the investor perform on date t compared to the (different) portfolio he held k trading days ago?" The reason for using a conditional measure is simple. If all investor-day GT observations (i.e., also those where the investor has not made any changes in the portfolio) were included, the inactive day traders' group average would be drawn towards zero. Our definition circumvents this problem: the conditional measure is comparable across investor groups independent of their trading activity.

Table 5: Self-Benchmark Measure of Day Traders’ Performance

This table estimates a variant of the Grinblatt and Titman (1993) measure of performance for day traders. The GT measure is computed as the difference between the return on the actual portfolio that the investor held from $t - 1$ to t and the return on the portfolio that the investor held $k = \{1, 5, 21, 63\}$ trading days ago. This measure is computed for each investor-day. Observations where the investor is out of the market or where $t - 1$ and $t - k$ portfolios are identical are dropped (see text). We create daily time-series from January 1998 to November 2002 for each day trader group by averaging across day traders. This table reports the time-series means and standard errors for investor groups based on the number of day trades during the past three months (N_{DT}); t -values are reported in parentheses. Panel B reports the pair-wise t -tests between the groups.

Panel A: Average GT_t^k -Measure

Group	Lag in the GT_t^k -Measure			
	$k = 1$	$k = 5$	$k = 21$	$k = 63$
$N_{DT} > 0$	-0.115% (-17.8)	-0.056% (-8.8)	-0.033% (-4.5)	-0.026% (-2.7)
$1 \leq N_{DT} < 10$	-0.087% (-14.1)	-0.043% (-7.2)	-0.026% (-3.7)	-0.017% (-1.8)
$10 \leq N_{DT} < 20$	-0.184% (-12.4)	-0.097% (-7.7)	-0.075% (-5.3)	-0.073% (-4.3)
$20 \leq N_{DT}$	-0.260% (-15.5)	-0.151% (-8.9)	-0.116% (-6.1)	-0.159% (-7.2)
$N_{DT} = 0$	-0.017% (-3.3)	-0.012% (-2.5)	-0.006% (-1.0)	0.003% (0.4)

Panel B: Pair-Wise t -Tests

Comparison	Lag in the GT_t^k -Measure			
	$k = 1$	$k = 5$	$k = 21$	$k = 63$
Group 3 – Group 2	-0.096% (-7.0)	-0.054% (-5.3)	-0.049% (-4.6)	-0.056% (-4.6)
Group 4 – Group 3	-0.074% (-3.8)	-0.055% (-3.4)	-0.042% (-2.4)	-0.085% (-4.2)

1998 observations.

4.3.2 Results

The Grinblatt and Titman (1993) measure estimates in Table 5 indicate that the day traders’ portfolio changes hurt their performance. The results highlight three distinct regularities. First, all day trader groups experience a significantly negative impact on their long-term performance. (The “no day trades” group estimates show that this finding is *specific* to day traders: the underperformance of these investors is statistically indistinguishable from

zero at one- and three-month horizons.) For example, the typical day trader ($N_{DT} > 0$) experiences a performance gap of -0.115% ($t = -17.8$) one day after a change in holdings. This underperformance does not disappear as the horizon increases. The three-month GT estimate—i.e., the average underperformance *per day* during the first three months after a holdings change—is -0.026% . This translates to a total performance gap of $(1 - 0.00026)^{63} - 1 \approx -1.6\%$ in three months using the point estimate.

Second, the GT estimates are monotonically decreasing in the day trading activity. The underperformance of the most active day traders is dramatic: the one-day estimate is -0.26% and the three-month estimate is -0.159% . For example, suppose conservatively that we have overestimated the three-month GT -measure by a factor of 1.5 and that the true gap is “only” -0.1% per day. Even with this estimate, the active day traders’ daily difference translates to a (raw) performance gap of -6% in three months! (Note that this measure is *conditional* on there being changes in holdings. Hence, the result does not say that these day traders continuously hurt their performance by this amount.)

Third, the GT estimates dampen as the horizon increases. This shows that *most* of the gap occurs soon after a change in holdings. This may arise from the reason *why* day traders change their holdings—i.e., when they keep losing shares intended for day trading. The results show that when individuals’ hang on to same-day losers, they are also on the wrong side of the market in the coming days.

The finding that more active day traders perform worse is important because these are precisely the investors with high day trading profits. Our results show that poor long-term performance may offset (relatively) high short-term day trading profits. The data confirms this relation: for the average day trader ($N_{DT} > 0$), the correlation between the $t - 1$ net trading return and GT_t^1 is negative 58.9% of the time. For the most active day traders, this correlation is negative 59.4% of the time. The implication of these aggregate- and individual-level results is that high day trading returns predict poor subsequent portfolio performance. Hence, an analysis of day trading profits would give a misleading picture of active day traders’ performance.

4.4 Day Traders' Migration Towards Riskier Portfolios

4.4.1 Measuring Changes in Factor Loadings

The Grinblatt and Titman (1993) measure is a risk-adjusted measure of stock picking skills if an investor's risk-profile does not change over time. We expect individual day traders to violate this assumption. We have shown that individuals' portfolios change as they day trade because of the disposition effect. If the stocks individuals pick for day trading are "different" from the stocks they previously own, day traders' exposure to market-wide shocks changes over time. This section addresses two questions about these portfolio changes. First, we examine whether individual day traders migrate towards stocks that are different from the ones they used to hold. Second, we examine whether this migration explains any of the long-term performance results of the previous section.

Our method of evaluating changes in factor loadings is straightforward. Suppose that returns r_t and r_t^k (from Eq. 6 and 7) are generated by M factors:

$$\begin{aligned} r_t - r_f &= \beta_1 \tilde{f}_{1,t} + \cdots + \beta_M \tilde{f}_{M,t} + \varepsilon_t \\ r_t^k - r_f &= \beta_1^k \tilde{f}_{1,t} + \cdots + \beta_M^k \tilde{f}_{M,t} + \varepsilon_t^k \end{aligned} \quad (9)$$

where β_j is the old portfolio's loading on factor j and β_j^k is the new loading. The difference $GT_t^k \equiv r_t - r_t^k$, is then:

$$GT_t^k = (\beta_1 - \beta_1^k) \tilde{f}_{1,t} + \cdots + (\beta_m - \beta_m^k) \tilde{f}_{m,t} + (\varepsilon_t - \varepsilon_t^k). \quad (10)$$

Hence, if an investor's portfolio profile changes over time, a regression of his GT measure against the factor portfolios produces slope estimates different from zero. These slope coefficients indicate the direction where the investor is moving his portfolio.

We now specify the (observable) factors and then formulate the test. We use the Fama and French (1993) three-factor model but also add a fourth factor: the excess return on the technology industry portfolio, $r_{techrf,t} \equiv r_{tech,t} - r_{f,t}$.¹⁸ We include this portfolio for two

¹⁸We construct equally-weighted factor portfolios and use a 50% cut-off point to create the SMB and HML factors because of the size of the market. The trade-off is between being able to span the risk factor (i.e., to maximize distance between the portfolios with respect to the variable of interest) and being able to estimate the top- and bottom-portfolio returns reliably. The latter is the more important concern in a small market.

reasons. First, the stocks that individuals prefer to day trade are probably from this sector (Jordan and Diltz 2002). Second, the “tech boom” is a popular explanation for the rise and fall of day traders. Hence, it is interesting to include the technology portfolio and see if the data supports this proposition. (We *do* show that the results are robust to replacing this factor with momentum portfolios.)

We estimate the following model for each day trader j :

$$GT_{t,j}^k = \alpha_j + \beta_{\text{mkt},j} r_{\text{rmrf},t} + \beta_{\text{tech},j} r_{\text{techrf},t} + \beta_{\text{smb},j} r_{\text{smb},t} + \beta_{\text{hml},j} r_{\text{hml},t} + \varepsilon_{t,j} \quad (11)$$

using the entire sample period from January 1995 to November 2002. We drop day traders with less than 20 GT_t^k observations from the subsequent analysis. Next, we compute average coefficient estimates $(\widehat{\alpha}, \dots, \widehat{\beta_{\text{PR1MO}}})$ for each day trader group. We then examine two questions:

- Are the average intercepts from Eq. 11 significantly different from zero, and are there differences between day trader groups?
- Are the average slope estimates from Eq. 11 significantly different from zero, and are there differences between day trader groups?

The first one asks whether the long-term underperformance is explained by changes in factor loadings. The second one asks whether *all individual day traders migrate towards similar type of stocks*. Note that the average loadings estimates across investors can be zero even if each individual’s intercept dampens significantly. The economic intuition is that the factor model in Eq. 11 can fit “different investors for different reasons.” The average loading difference estimates are significantly different from zero only if all individual day traders’ migrate towards similar type of portfolios.

4.4.2 Results

The multifactor model estimates in Table 6 show that factor portfolios explain part of long-term underperformance and that individual day traders migrate towards similar portfolios.¹⁹

¹⁹We report the results only for lags $k = 1$ and $k = 63$ for the sake of brevity but the results are similar for $k = 5$ and $k = 21$.

Table 6: Systematic Risk Factors in Day Traders' Portfolio Changes

We estimate regression

$$GT_{t,j}^k = \alpha_j + \beta_{mkt,j} r_{rmrf,t} + \beta_{tech,j} r_{techrf,t} + \beta_{smb,j} r_{smb,t} + \beta_{hml,j} r_{hml,t} + \varepsilon_{t,j}$$

for each day trader j with at least 20 observations between January 1995 and November 2002. $RMRF$, $TECHRF$, SMB , and HML are self-financed portfolios. This table reports average coefficient estimates across day traders for groups based on the number of day trades during the past three months (N_{DT}); standard errors are reported in italics under the estimates. Panel A reports the one-day lagged estimates ($k = 1$) and Panel B reports the three-month estimates ($k = 63$). The average R^2 from the regression is reported under \widehat{R}_i^2 . The average GT^k measure across investors is reported under \widehat{GT}_i^k . We first compute the average GT^k for each investor-group (i.e., we include daily observations where day trader j belongs to group g) and then average across individuals in each group.

Panel A: Average Coefficient Estimates for GT_i^1 Regressions

Group	α	β_{mkt}	β_{tech}	β_{smb}	β_{hml}	\widehat{R}_i^2	\widehat{GT}_i^1
$N_{DT} > 0$	-0.103%	-0.059	0.018	-0.001	0.017	13.6%	-0.136%
	<i>0.007%</i>	<i>0.011</i>	<i>0.003</i>	<i>0.004</i>	<i>0.006</i>		<i>0.005%</i>
$1 \leq N_{DT} < 10$	-0.074%	-0.058	0.016	0.001	0.013	14.8%	-0.108%
	<i>0.007%</i>	<i>0.012</i>	<i>0.003</i>	<i>0.004</i>	<i>0.006</i>		<i>0.005%</i>
$10 \leq N_{DT} < 20$	-0.168%	-0.100	0.029	-0.008	0.031	15.0%	-0.215%
	<i>0.022%</i>	<i>0.036</i>	<i>0.010</i>	<i>0.015</i>	<i>0.021</i>		<i>0.016%</i>
$20 \leq N_{DT}$	-0.249%	-0.110	0.010	0.007	0.015	12.2%	-0.309%
	<i>0.027%</i>	<i>0.048</i>	<i>0.011</i>	<i>0.015</i>	<i>0.017</i>		<i>0.020%</i>
$N_{DT} = 0$	-0.013%	-0.019	0.024	0.003	-0.004	15.1%	-0.022%
	<i>0.003%</i>	<i>0.006</i>	<i>0.002</i>	<i>0.002</i>	<i>0.003</i>		<i>0.002%</i>

Panel B: Average Coefficient Estimates for GT_i^{63} Regressions

Group	α	β_{mkt}	β_{tech}	β_{smb}	β_{hml}	\widehat{R}_i^2	\widehat{GT}_i^{63}
$N_{DT} > 0$	-0.005%	-0.008	0.026	0.026	-0.010	12.7%	-0.021%
	<i>0.004%</i>	<i>0.007</i>	<i>0.002</i>	<i>0.003</i>	<i>0.003</i>		<i>0.002%</i>
$1 \leq N_{DT} < 10$	-0.005%	-0.003	0.027	0.026	-0.008	13.1%	-0.019%
	<i>0.004%</i>	<i>0.007</i>	<i>0.002</i>	<i>0.003</i>	<i>0.004</i>		<i>0.002%</i>
$10 \leq N_{DT} < 20$	-0.060%	-0.021	0.023	0.038	-0.026	14.6%	-0.087%
	<i>0.015%</i>	<i>0.027</i>	<i>0.009</i>	<i>0.011</i>	<i>0.014</i>		<i>0.009%</i>
$20 \leq N_{DT}$	-0.077%	-0.032	-0.010	0.057	-0.079	12.3%	-0.143%
	<i>0.025%</i>	<i>0.041</i>	<i>0.014</i>	<i>0.018</i>	<i>0.025</i>		<i>0.018%</i>
$N_{DT} = 0$	-0.014%	0.023	0.018	0.023	-0.014	4.4%	-0.004%
	<i>0.002%</i>	<i>0.004</i>	<i>0.001</i>	<i>0.002</i>	<i>0.002</i>		<i>0.001%</i>

A comparison of the raw GT_k measures (the rightmost column) and the intercepts from the regression shows that the long-term underperformance weakens for all day trader groups. The effect is moderate for the next-day measure but considerable at the three-month horizon. For

example, the change for “all day traders” is from a highly significant -0.021% to a statistically insignificant -0.005 . The model’s ability to explain long-term underperformance is particularly good for the least active group of day traders ($0 < N_{DT} < 10$). Although the model reduces the estimates also for the other day trader groups, the intercepts remain significantly negative. For example, the most active group’s raw three-month measure is -0.143% while it is only about half of that (-0.077%) after controlling for loadings changes. This result is encouraging but it still raises the question about the source of this underperformance. The results show that the idiosyncratic component is strong only at short lags—the model does poorly in explaining the next-day gap but performs well in explaining the long-term gap. Hence, the residual underperformance arises in the few days after a holdings change.

The slope estimates show that individual day traders migrate towards same type of stocks. The one-day measure shows that day traders move away from the general market towards technology stocks. Note that the changes in SMB and HML factor loadings also become significant at the three-month horizon. These estimates show that day traders systematically tilt their holdings towards small technology stocks with low B/M ratios. This is consistent with the common proposition that many individual day traders ended up taking large positions in the technology sector.

We believe these results are not only of statistical but also of economical significance. We only pick up the migration component that is *common* across all investors: all investors must make same type of changes repeatedly. If individuals’ portfolio movements were idiosyncratic, the average estimates across individuals would be zero. Moreover, because individuals’ do not change their entire portfolios, the *magnitude* of the profile changes we expect to observe is small. Collectively, the fact that we pick up significant changes in factor loadings in our test is an important result.

4.4.3 Risk or Mispicing?

Does the technology portfolio in our model specification (Eq. 11) capture movements of a risk factor or does it reflect time-variation in the mispricing of technology stocks?²⁰ We address this concern and estimate an alternative specification of Eq. 11 with RMRF, SMB, HML,

²⁰Such a mispricing argument about is debatable. For example, Pástor and Veronesi (2005) show that rational learning about firms’ profitability can explain the high Nasdaq valuations in the late 1990s.

PR1YR, and PR1MO as the factor portfolios.²¹ This specification gives an average intercept of $\hat{\alpha} = -0.072\%$ (s.e. = 0.022%) for the most active group of day traders. The Fama-French factor slope coefficients are also similar. This is close to our earlier estimate and confirms that the technology portfolio alone does not generate the results.

The question of whether day traders exposed themselves to more risk or to more mispricing (or whether they themselves could tell the difference) is an interesting question but beyond the scope of this paper. (The Fama-French and momentum factor portfolios are open to both interpretations; see, e.g., Hirshleifer 2001.) However, our results do not require a stance on this issue. The results show that market-wide shocks drive *most* of day traders' long-term underperformance. This contrasts with a view that individual day traders' might (somehow) possess negative stock picking skills. We believe that—at least for our purposes—this “shocks versus negative skills” comparison is more important than the “risk or mispricing” question.

5 Conclusions

This paper studies how individual day traders' tendency to hold on to losers affects their performance. The inability to come to terms with losses may be costly when the objective is to profit from short-term price movements alone. For example, suppose an investor purchases shares in the morning with intent of selling them in the afternoon at profit. If the company announces disappointing news, the investor may abort the day trade and hold the shares, possibly selling something else from his portfolio to finance the purchase. The investor has now moved towards a portfolio he would not have chosen under different circumstances. This paper studies the existence and the consequences of this type of behavior.

We find that individual day traders' behavior supports all aspects this story. First, individuals are very reluctant to complete their day trades at loss. They often sell other stocks from their portfolios to finance their “unintended” purchases. Second, an analysis of individual day traders' short- and long-term performance shows that this behavior is costly. The average day trader does not earn day trading profits. Moreover, although the most active day traders are profitable *before* brokerage firm fees, even they lose significantly after accounting for these costs. Only 3 out of 10 of our day traders earn positive profits in a typical month.

²¹PR1YR and PR1MO are momemtum portfolios (see, e.g., Carhart (1997)). The first one is based on stock returns from -1 year to -1 month; the latter captures the spread in -1 month to the previous day returns.

The long-term performance results are more striking, reflecting the costs associated with the investors' portfolio changes. Individual day traders' post-trade portfolios significantly underperform their pre-trade portfolios. This drift is *specific* to day traders: investors who have either quit day trading or are yet to begin do not hurt their performance. We demonstrate two strong patterns in this long-term performance. First, more active day traders hurt their long-term performance more: hence, there is a negative relation between short-term day trading profits and long-term portfolio performance. Second, *most* of the long-term performance arises soon after a change in holdings.

The long-term underperformance is linked to the Fama and French (1993) factors and the movements of the technology sector. First, we find that the typical individual day trader migrates away from the general market towards small technology stocks with low B/M ratios. Second, after accounting for these changes in the exposure to market-wide shocks changes, *most* of the long-term underperformance disappears. Only the more active day traders continue to exhibit residual underperformance during the first few days after a change in portfolio holdings. We do not attempt to distinguish between risk and mispricing explanations to why day traders lose in the long-run—the result we emphasize is that it is market-wide shocks that drive underperformance, not day traders' (somehow) negative stock picking skills.

Our findings pertain to all individuals who try to earn a quick profit in the stock market. We find that behavioral biases—e.g., the reluctance to take losses—significantly affect not only investors' immediate welfare but may also push them towards *too risky*, suboptimal portfolios. This mechanism for portfolio changes suggests that individuals might be reluctant to hold the resulting portfolios under different circumstances. For our individual day traders, both consequences of this mechanism—the immediate day trading losses and the migration towards riskier portfolios—are very significant.

References

- Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean (2004). Do individual day traders make money? Evidence from Taiwan. University of California, Berkeley, working paper.
- Barber, B. M. and T. Odean (2001). The Internet and the investor. *Journal of Economic Perspectives* 15, 41–54.
- Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance* 56, 1–43.
- Carhart, M. (1997). On persistence in mutual fund performance. *Journal of Finance* 52, 57–82.
- Fama, E. and K. French (1993). Common risk factors in the returns of stocks and bonds. *Journal of Financial Economics* 33, 3–56.
- Garvey, R. and A. Murphy (2001). How profitable day traders trade: an examination of trading profits. University College Dublin, working paper.
- Grinblatt, M. and M. Keloharju (2000). The investment behavior and performance of various investor-types: A study of Finland's unique data set. *Journal of Financial Economics* 55, 43–67.
- Grinblatt, M. and M. Keloharju (2001). What makes investors trade? *Journal of Finance* 56, 589–616.
- Grinblatt, M. and S. Titman (1993). Performance measurement without benchmarks: An examination of mutual fund returns. *Journal of Business* 66, 47–68.
- Harris, J. H. and P. H. Schultz (1998). The trading profits of SOES bandits. *Journal of Financial Economics* 50, 39–62.
- Hirshleifer, D. (2001). Investor psychology and asset pricing. *Journal of Finance* 56, 1533–1597.
- Jordan, D. J. and J. D. Diltz (2002). The profitability of day traders. *Financial Analysts Journal* 59, 85–94.

- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Linnainmaa, J. (2003). Who makes the limit order book? Implications for contrarian strategies, attention-grabbing hypothesis, and the disposition effect. University of California, Los Angeles, working paper.
- Linnainmaa, J. (2005). Learning from experience. University of California, Los Angeles, working paper.
- Locke, P. R. and S. C. Mann (2005). Professional trader discipline and trade disposition. *Journal of Financial Economics* 76, 401–444.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.
- Pástor, L. and P. Veronesi (2005). Was there a Nasdaq bubble in the late 1990s? *Journal of Financial Economics*,. Forthcoming.
- Schlarbaum, G. G., W. G. Lewellen, and R. C. Lease (1978). Realized returns on common stock investments: The experience of individual investors. *Journal of Business* 51, 299–325.
- Seasholes, M. and G. Wu (2004). Profiting from predictability: smart trades, daily price limits, and investor attention. University of California, Berkeley, working paper.
- Shefrin, H. and M. Statman (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40, 777–790.