Learning Fast or Slow

Brad M. Barber
Graduate School of Management
University of California, Davis

Yi-Tsung Lee
Guanghua School of Management
Peking University

Yu-Jane Liu
Guanghua School of Management
Peking University

Terrance Odean
Haas School of Business
University of California, Berkeley

Ke Zhang
Business School
Nanjing University

December 27, 2018

This paper was previously titled "Do day traders rationally learn about their ability?" We are grateful to the Taiwan Stock Exchange for providing the data used in this study. Barber appreciates the National Science Council of Taiwan for underwriting a visit to Taipei, where Timothy Lin (Yuanta Core Pacific Securities) and Keh Hsiao Lin (Taiwan Securities) organized excellent overviews of their trading operations. We have benefited from the comments of seminar participants at Cal Tech, Columbia University, HKUST, Peking University, PUC Rio de Janeiro, Santa Clara University, Texas A&M, University of Ghent, University of Illinois, University of Oregon, University of Utah, University of Virginia, USC, UT-Austin, Washington University.

Learning Fast or Slow

Abstract

Rational models claim "trading to learn" can explain widespread excessive speculative trading and challenge behavioral explanations of excessive trading (e.g., overconfidence and entertainment). We argue rational learning models cannot explain much of speculative trading by studying day traders in Taiwan from 1992-2006. Consistent with previous studies of learning, unprofitable day traders are more likely to quit than profitable traders. Consistent with models of overconfidence and biased learning, but not with rational learning, the aggregate performance of day traders is negative, the vast majority of day traders are unprofitable, and a great many day traders persist in trading despite persistent losses.

This paper contributes to a debate about whether individual investors trade excessively and, if so, why. Noting the psychologists find that people tend to be overconfident, several theoretical models of overconfidence show that, in a variety of theoretical settings, investors trade more actively if they are overconfident and thereby lower their utility (Odean, 1998; Daniel, Hirshleifer, and Subrahmanyan, 1998, 2001; Kyle and Wang, 1997; Gervais and Odean, 2001; Caballe and Sakovics, 2003). A series of papers analyzing the trading records of individual investors report findings consistent with these models of investor overconfidence: in the year following stock trades, the stocks purchased by individual investors at a US discount brokerage (1987-1993) underperform the stocks they sold (Odean, 1999); unsophisticated Finnish investors underperform all other investors on the Helsinki Stock Exchange (1995-1996) (Grinblatt and Keloharju, 2000); the portfolios of the most active investors at a US discount brokerage (1991-1996) underperform those of less active investors (Barber and Odean, 2000); and, in this same dataset, women, who tend to exhibit less overconfidence than men in areas such as finance, trade less actively than men and reduce their returns less through active trading than do men (Barber and Odean, 2001). Analyzing all stock trades by all investors on the Taiwan Stock Exchange from 1991-1995, Barber, Lee, Liu, and Odean (2008) find that after transactions costs active trading reduces the return of the aggregate individual investor portfolio by 3.8 percentage points a year, the equivalent of 2.2 percent of Taiwan's GDP. Glaser and Weber (2007) match survey data to trading records for individual investors and document investors-based scores of overconfidence correlate with trading activity. Biais, Hilton, Mazurier, and Pouget (2005) use an experimental asset market to document calibration-based overconfidence reduces trade performance, and Deaves, Luders, and Luo (2008) provide experimental evidence that calibration-based overconfidence induces additional trade.[1]

As many of these papers acknowledge, excessive trading could result for reasons other than overconfidence, such as a desire for entertainment, or from a combination of overconfidence and motives such as a desire for entertainment. Grinblatt and Keloharju

[1] Several papers document individual investors in the U.S. earn strong returns immediately following their trades (Barber, Odean, and Zhu, 2008; Kaniel, Saar, and Titman, 2008; Kelly and Tetlock, 2013). Barber, Odean, and Zhu (2008) document poor returns (reversals) at longer horizons using a 20-year dataset derived from TAQ; Kaniel et al. (2008) and Kelly and Tetlock (2013) use datasets with more limited time series and fail to find evidence of poor returns at longer horizons.

(2009) find that Finnish investors who get more speeding tickets trade more actively which they attribute to sensation seeking. Dorn and Sengmueller (2009) match survey responses and administrative records at a German broker and document investors who report enjoying gambling have twice the turnover rates of their peers. The introduction of a national lottery in Taiwan reduced trading by on the Taiwan Stock Exchange by about one fourth (Barber, Lee, Liu, and Odean, 2008; Gao and Lin, 2014); Dorn, Dorn, and Sengmueller (2015) provide similar evidence using California, multistate U.S., and German lotteries. Finally, Linnainmaa (2010) argues that the underperformance of individual investors results from unmonitored limit orders.

Rational learning models are offered as alternative explanations for excessive speculative trading (Mahani and Bernhardt, 2007; Linnainmaa, 2010; Seru et. al., 2010, Lubensky, 2017). Mahani and Bernhardt (2007, p.1334) write that their model indicates "Most inexperienced traders realize losses, conclude that they are unlikely to be skilled, and leave the markets; survivors expand their trades and make more profits. Learning produces the aggressive trading that is traditionally attributed to psychological biases." Linnainmaa (2011) writes "When agents can learn about their abilities as active investors, they rationally 'trade to learn' even if they expect to lose from active investing…. a stylized learning model approximates household trading decision remarkably well." Similarly, Seru et. al. (2010) argue conclude (p. 733) "an open question in the literature is why there is such high trading volume, particularly among seemingly uninformed individual investors. Our results indicate that such trading may be rational; investors may be aware that they will learn from experience and trade in order to learn." While these papers do not dismiss behavioral motives for trade, they make the case that learning **alone** can explain why individuals engage in so much speculative trading.

In these models, investors do not initially know their own abilities and rationally infer their abilities by observing their trading performance. Investors are aware that most speculators lose money, but believe that these losses are small compared to the profits of successful speculators. Thus risk-averse investors rationally try speculation because they believe expected lifetime profits are positive. Unsuccessful investors conclude that they lack or are unable to acquire ability and quit trading speculatively; successful investors continue to speculate and make profits doing so.

Gervais and Odean (2001) also develop a model in which investors with different abilities do not initially know their own ability (i.e., the precision of their information) and infer ability by observing their trading performance. In this biased learning model, self-attribution bias leads investors to update more on successes than on failures. Thus, successful investors become overconfident and unsuccessful investors are slow to realize that they have low ability and to curtail their trading.

The qualitative predictions of the rational and biased learning models are identical. Investors will increase their trading in response to successful trades and decrease their trading in response to unsuccessful trades. The difference is in the speed with which investors respond to feedback. Biased traders learn slowly.

We test the predictions of rational learning models by analyzing the performance of day traders in Taiwan. Using the complete transaction data for the Taiwan Stock Exchange over 15 years (1992 to 2006), we focus on traders who buy and sell the same stock within a day, as these traders are almost surely speculators.

Previous tests of rational learning models of trading have focused primarily on confirming evidence, e.g., do investors increase (decrease) trading in response to successful (unsuccessful) trades? Seru, Shumway, and Stoffman (2010) document unsuccessful Finnish speculators are more likely to stop trading and that some speculators get better at trading with experience. Campbell, Ramadorai, and Ranish (2014) document that previously successful Indian investors trade more aggressively. Linnainmaa (2010) finds that Finnish investors are more likely to increase trade size after successful trades and more likely to decrease trade size or quit trading after unsuccessful trades. Furthermore, the size and quitting effects are stronger early in an investor's career, when his or her prior beliefs about ability are more diffuse. Using US broker data, Nicolosi, Peng, and Zhu (2009) show that trade intensity increases following signals of strong performance. Analyzing data from the National Stock Exchange in India, De et. al. (2010) document that investors increase trading in response to recent profits and that the sign of profits matters more than their magnitude.

We, too, find evidence of learning among day traders. The majority of day traders quit relatively quickly (more than 75% of all day traders quit within two years), and poor performers are more likely to quit. These results are consistent with the models of both rational and biased learning. However, to properly test a model it is important to look for disconfirming evidence. We document behavior that is not consistent with rational Bayesian learning as modeled by Mahani and Bernhard (2007), Linnainmaa (2010), and Lubensky (2017) for two reasons:

First, if the entry (and exit) of speculators is stable over time, then the sign of the expected lifetime profits of new speculators is the same as that of aggregate speculator profits. Therefore risk-averse (or risk-neutral) potential speculators with no special prior knowledge of their abilities should only "trade to learn" if aggregate speculator profits are positive.[2] In fact, using complete data for the Taiwan market, the aggregate performance of day traders net of fees is negative in each of the 15 years that we study. A profit-maximizing risk-averse Bayesian investor would not enter a market if her expected lifetime profits were negative.

Of course, a few investors with unusual abilities, such as extraordinary mathematical skills, might rationally anticipate becoming successful day traders. We do not dispute that it may be rational for some day traders to enter the market or that some learn rationally. However, we argue that entry and learning for the purpose of profits cannot explain a material portion of the observed day trading since the vast majority of day traders and volume can be traced to those with a history of losses. If most Taiwanese day traders undertake day trading because they believe day trading has positive lifetime expected profits, they are mistaken and overconfident. Conditional on that overconfidence, entry may look reasonable. But this behavior is consistent with the overconfidence theories, not with rational trading to learn.

Second, though performance affects day trader persistence, many traders continue to day trade after extended periods of losses. In one analysis, we measure how past performance affects traders willingness to continue. Specifically, we sort traders based on the number of days they have engaged in day trading as of the end of the previous month and on the past profitability of their intraday trades net of trading costs. Previously unprofitable traders with 50 or more days of past day trading experience have a 95.3% probability of day trading again in the next 12 months, while previously profitable traders

[2] In the rational learning models, as well as Gervais and Odean (2001), the unconditional lifetime profit from trying trading is positive.

with 50 or more days of past day trading experience have a 96.4% probability doing so. So, not only do experienced day traders with histories of losses persist in day trading, they do so at nearly the same rate day traders who have been profitable.

We also document that 97% of day traders and over 90% of day trading volume can be traced to investors who, based on their past experience and performance, are likely to lose money in future day trading. Specifically, we sort day traders into six groups based on their past experience. Within each experience group, we further partition investors based on whether their past profits net of trading costs are positive or negative. Thus, we are left with 12 groups based on the double-sort of past experience and past profitability. Among these 12 groups, only the profitable and most experienced investors predictably earn future profits net of trading costs. However, on average, these investors represent less than 3% of all day traders each day and less than 10% of all day trading volume.[3] Put differently, each day 97% of day traders can expect to lose money from trading and more than 90% of all day trading volume can be traced to investors who predictably lose money.

Given the overwhelming lack of profitability of day traders and the persistence of trading in the face of losses, rational trading to learn cannot account for "high trading volume … among seemingly uninformed individual investors."

I. Learning by Speculators

The suggestion that investors learn from experience is neither novel nor controversial.[4] Learning is a ubiquitous feature of human experience. From a welfare and policy perspective, the question is not whether investors learn, but how well they learn. In this section we develop testable predications that emanate from a rational model of learning and highlight the predictions would discriminate between rational and biased models of learning.

---

[3] Because highly experienced profitable day traders trade more actively than others, they are overrepresented when measured on a daily basis. Employing different methodology, Barber, Lee, Liu, and Odean (2014) estimate that less than 1% of individuals who day trade over the course of a year are predictably profitable.

[4] A number of papers document investor learning in various forms including Feng and Seasholes (2005), Seru, Shumway, and Stoffman (2007), Nicolosi, Peng, and Zhu (2009), Chiang, Hirshleifer, Qian, and Sherman (2010), Choi, Laibson, Madrian and Metrick (2010), De, Gondhi, and Pochiraju (2010), and Odean, Strahilesvitz, and Barber (2010).

Mahani and Bernhardt (2007) argue that rational Bayesian learning can explain several empirical regularities: cross-sectionally, most speculators lose money; large speculators outperform small speculators; past performance positively affects subsequent trade intensity; most new traders lose money and cease speculation; and performance shows persistence. With view similar to Mahani and Bernhardt's (2007), Linnainmaa (2010) develops a structural model of rational learning. Using trading data from Finland, he finds investors with poor performance are likely to quit and trading intensity increases following good performance.

The "trading to learn" literature argues that investors rationally try their hands at speculation to make profits. These investors are not trading for entertainment nor are they risk-seeking. Consider a risk-neutral investor with a zero discount rate[5], no utility or disutility from the activity of trading, and no opportunity cost for his capital. At time *t,* he chooses speculate if his expected lifetime profits from speculation are positive conditional on his information set at *t,* Ω_t. Let $t = 0$, be the time at which an investor who has never speculated decides whether or not to do so. The rational investor will only speculate if $\mathrm{E}\left(\text{Lifetime Profits} \mid \Omega_0\right) > 0$. Prior to acquiring speculation experience, the average investors' unconditional expected lifetime profits from speculation are the average lifetime profits of other investors who undertake speculation. While any given investor's conditional lifetime profits from speculation may be higher than average, the average investor who engages in speculation should expect to earn average lifetime profits.

If the speculative market is in equilibrium, that is, speculator profits and the numbers of speculators entering and leaving are stable, and if unsuccessful speculators exit the market faster than successful speculators, then the sign of the unconditional expected lifetime profits from speculation is the same as the sign of the average annual profit from speculation. To see this, consider a market in which N new speculators enter the market each year; n, $0 < n < N$, new speculators either are skilled or have the ability to acquire skill (prior to entry the potential speculator cares about his lifetime expected

[5] Risk averse investors and those with high discount rates will require a higher threshold of expected lifetime profits before choosing to enter due to the uncertainty of profits and that, in expectation, losses are likely to be incurred earlier than gains.

profits, irrespective of whether these result from innate skill or learning and, since we have assumed a zero discount rate, the timing of the profits is not important); *N-n* new speculators are unskilled with no ability to acquire skill. Unskilled speculators trade for 1 year, lose 1 dollar, and quit. Skilled speculators trade for $Y > 1$ years, earn a profit of $S > 0$ dollars each year, and quit (retire). Because skilled speculators stay in the market longer than unskilled speculators (i.e., $Y > 1$), skilled speculators account for a greater fraction of active speculators than of new speculators (i.e., $Y*n \,/\, (Y*n + N - n) > n/N$). The unconditional expected lifetime profits of a new speculator are:

$$\mathrm{E}\big(\text{Lifetime Profits}\big) = \frac{n*Y*S}{N} - \frac{N-n}{N} = \frac{n\big(Y*S+1\big)-N}{N}$$

The average expected profits of an active speculator in any year are:

$$\text{Average Annual Profits} = \frac{n*Y*S-\big(N-n\big)}{n*Y+\big(N-n\big)} = \frac{n\big(Y*S+1\big)-N}{N+n*(Y-1)}$$

Average annual profits, i.e., aggregate annual profits divided by the number of all speculators trading in a year, are of the same sign but lower magnitude than expected lifetime profits. Thus, when average annual profits are negative, they provide an upper bound to the unconditional expected lifetime profits of a speculator. Furthermore, since skilled speculators stay in the market longer than unskilled (i.e., $Y > 1$), in equilibrium the proportion of speculators in the market at any time who have skill is greater than the unconditional probability that a new speculator has skill.

Thus if traders have rational prior beliefs about the unconditional expected lifetime profits from engaging in speculation (and are engaging in speculation for the purpose of earning lifetime profits), then the aggregate performance of speculators should be positive. This leads to our first null hypothesis:

H1: *The aggregate net performance of day traders is positive (non-negative).*

The alternative is that the aggregate net performance of day traders is negative. This is consistent with traders holding biased prior beliefs about the unconditional expected lifetime profits from trying day trading.

Once a trader engages in speculation he begins to acquire feedback about his potential skill. In the rational learning models, unsuccessful traders quit trading after the accumulation of negative signals outweighs their positive initial prior beliefs about their ability. Gervais and Odean (2001) develop a model in which investors take too much credit for their success and thus, relative to a Bayesian, overweight successes when learning about their ability. In their model, too, successful investors have more ability than unsuccessful ones and investors respond to good performance by trading more aggressively. In contrast to the rational Bayesian model, their model can also explain persistent trading by previously unsuccessful traders; these traders put too much weight on successes and too little on failures when updating beliefs about their abilities.

Consider a trader who began speculation believing he had positive expected lifetime profits. Given net positive feedback, he will continue to speculate; but at some threshold of negative feedback, he should quit. Consider a trader who began speculation believing he had negative expected lifetime profits. Entry was not a rational decision, and while there is a threshold of positive feedback at which it may be rational to continue trading, persistent trading in the face of losses is not consistent with the models of rational learning.

If traders began speculation with the rational prior belief that their expected lifetime profits from speculation were positive, an econometrician could specify and calibrate a model of quitting behavior. The econometrician could estimate the probability of negative future profits conditional on trading histories and test whether traders quit at the point at which their trading histories forecast future losses. However, when traders begin speculation with negative expected lifetime profits from speculation, the point at which they should rationally quit speculation is before they even start, irrespective of their learning process. When traders who should have negative priors continue to trade in the face of negative feedback, it is difficult to disentangle whether persistence is due to extremely biased priors or biased learning. We present evidence of remarkable trading persistence in the face of losses, for example, nearly 3/4ths of day-trading volume is generated by unsuccessful day-traders with at least 10 days of day trading experience and, in an average month, only 9% of those day traders who have 400 or more days of day trading experience have earned positive lifetime intraday net returns. This perverse persistence might result from biased prior beliefs, biased learning, or both. However, as feedback grows, rational learning models predict that even speculators who began with positively biased priors should quit in the face of losses.

H2: *Experienced day traders with previous net losses will stop trading speculatively.*

Under Gervais and Odean's biased learning model, it is possible for unsuccessful traders to become overconfident and, thus, more active traders, if their learning bias is sufficient. Thus continuing to trade by experienced unprofitable day-traders is contrary to the rational learning models but consistent with biased learning.

II. Data and Methods

II.A. Day Traders and Speculative Trading

Empirical tests of the learning models must identify traders who trade speculatively. Investors might reasonably trade to save (or consume), to rebalance their portfolios, or to reduce their tax liability. Thus, an important feature of our empirical strategy is to identify a clean sample of speculators. We do so by focusing on day trading on the Taiwan Stock Exchange. Day trading is the purchase and sale of the same stock by an investor on a day. We argue that these intraday trades are almost certainly speculative. Moreover, day trading is common and prevalent in Taiwan.

We are not the first to study day trading, though the sample of day traders we study is much larger and the time-series much longer than those in prior studies.[6] The one exception to this generalization being Barber, Lee, Liu, and Odean (2014) who identify a small subset of day traders (less than 1% of the day trading population) predictably earn profits. None of these prior studies used the empirical setting to test rational and biased models of learning, the focus of our investigation.

II.B. Taiwan Market Rules

Before proceeding, it is useful to describe the Taiwan Stock Exchange (TSE). The TSE operates in a consolidated limit order book environment where only limit orders are accepted. During the regular trading session, from 9:00 a.m. to noon during most of our sample period, buy and sell orders can interact to determine the executed price subject to

[6] Harris and Schultz (1998) study SOES bandits at two brokers. Garvy and Murphy (2002, 2005) analyze 15 and 1,386 day traders at one US broker. Seasholes and Wu (2004) analyze the trades of 10 active traders on the Shanghai Stock Exchange. Linnainmaa (2003) analyzes 7,686 Finnish day traders.

applicable automatching rules.[7] Minimum tick sizes are set by the TSE and vary depending on the price of the security. Generally, orders are cleared using automatching rules one to two times every 90 seconds throughout the trading day. Orders are executed in strict price and time priority. An order entered into the system at an earlier time must be executed in full before an order at the same price entered at a later time is executed. Although market orders are not permitted, traders can submit aggressive price-limit orders to obtain matching priority. During our study period, there is a daily price limit of 7% in each direction and a trade-by-trade intraday price limit of two ticks from the previous trade price.

Since our analysis focuses on day trading, an important consideration is transaction costs. The TSE caps commissions at 0.1425% of the value of a trade. Some brokers offer lower commissions for high-volume traders. Officials at brokerage firms and the TSE indicated to us that the largest commission discount offered is 50% (i.e., a commission of roughly 7 basis points); these same officials estimated the trade-weighted commission paid by market participants to be about 10 basis points. We use the 10 basis points when calculating returns net of fees. Taiwan also imposes a transaction tax on stock sales of 0.3%.

II.C. Trades Data and Descriptive Statistics

We use a unique and remarkably complete dataset, which contains the entire transaction data, underlying order data, and the identity of each trader on the Taiwan Stock Exchange (TSE). With these data, we provide a comprehensive accounting of the profitability of day traders during the period 1992 through 2006.

The trade data include the date and time of the transaction, a stock identifier, order type (buy or sell -- cash or margin), transaction price, number of shares, a broker code, and the identity of the trader. In total, the dataset contains 3.7 billion purchase (or sale) transactions with a value of $NT 310 trillion (approximately $10 trillion US).[8] The trader code allows us to broadly categorize traders as individuals, corporations, dealers,

---

[7] Trading also occurred on Saturdays during most of our sample period. Before December 1997, Saturday trading occurred from 9:00-11:00. From January to March, 1998, stocks were traded only on the second and the fourth Saturday in each month. From April 1998 to December 2000, Saturday trading occurred from 9 am to noon. From 2001 on, there has been no trading on Saturday.

[8] The mean TWD/USD exchange rate from 1992 to 2006 was 30.54 with a low of 24.65 and a high of 35.01.

foreign investors, and mutual funds. The majority of investors (by value and number) are individual investors. Corporations include Taiwan corporations and government-owned firms (e.g., in December 2000 the government-owned Post, Banking, and Insurance Services held over $NT 213 billion in Taiwanese stock).[9] Dealers include Taiwanese financial institutions such as Fubon Securities, Pacific Securities, and Grand Cathay Securities. Foreign investors are primarily foreign banks, insurance companies, securities firms, and mutual funds. During our sample period, the largest foreign investors are Fidelity Investments, Scudder Kemper, and Schroder Investment Management. Mutual funds are domestic mutual funds, the largest of which is ABN-AMRO Asset Management with $NT 82 billion invested in Taiwanese stocks in December 2000.

We define day trading as the purchase and sale, in any order, of the same stock on the same day by an investor. Specifically, if an investor buys and sells the same stock on the same day, we calculate the number of shares bought (S_b), the number of shares sold (S_s), the average purchase price (P_b), and the average sales price (P_s). The value of day trading is defined as half of the total value of sales and purchases $(½*P_b*min(S_b,S_s)+ ½*P_s*min(S_b,S_s))$. Over our sample period, aggregate day trading accounts for more than 19% of the total dollar value of aggregate trading volume.

Virtually all day trading can be traced to individual investors. In the average month, individual investors account for over 99% of all day traders (and 95% of day trading volume). Individuals and corporations are free to short sell, though dealers, mutual funds, and foreigners are prohibited from doing so on the TSE. These short sale restrictions might partially explain the tendency for day trading to concentrate among individual investors. In contrast to U.S. markets, dealers are not active providers of liquidity. TSE rules state that dealers are required to "efficiently adjust the demand and supply in the market depending on the market situation, and ensure that the formation of fair price and its sound operation are not harmed," yet dealers face no specific penalties for failing to meet this requirement. Dealer trades emanate from their proprietary trading activity. Based on our discussions with dealers in the TSE, the majority of this

[9] Many corporations are small firms that are majority or wholly owned by an individual. Thus, the corporate category of trader also includes thousands of individual investors who trade under the label of corporation.

proprietary trading is not necessarily intended to provide liquidity. Chae and Wang (2003) also report that TSE dealers are not net providers of liquidity. In the remainder of the paper, we focus on individual investors.

In Figure 1, we plot day trading as a percentage of total trading volume and the number of individuals who day trade by month. While day trading was somewhat less prevalent in the early part of our sample period, the share of volume traced to day trading has been consistently around 20% of total trading volume from 1995 to 2006. In the average month, almost 140,000 individuals day trade[10]. With an adult population of about 16 million (total population about 22 million), this means just shy of 1% of the adult population day trades in the average month. In terms of both a percentage of total trading volume and numbers of traders, day trading is an equilibrium feature of the Taiwan stock exchange with no apparent trend over the period 1997 through 2006.

II.D. Performance Measurement

Our primary performance measurement focuses on the intraday profits of all trades made by day traders and on trade-weighted intraday returns. This analysis focuses on returns on the day that positions are opened. In Section IV C, we also analyze the event time profitability of purchases and sales over longer holding periods to ensure the inferences we draw from the analysis of intraday profits are accurate.

We calculate the intraday returns to day trading, by identifying all trades made by day traders. We calculate the profits on round-trip day trades *and other position opening trades that remain open at the close of the trading da*y. The other trades are either purchases to open a long position or sales that open a short position. The profits for trades that lead to an open position are calculated relative to closing prices on the date of the trade (i.e., mark-to-market at the day's closing price). To calculate the daily return earned by a day trader, we sum the proceeds from stocks sold to close long positions and bought to close short positions (or their mark-to-market equivalent at the close of the trading day) and divide by the cost of initiating the position (i.e., the value of stocks bought or sold short at the time of the purchase or sale). We refer to this return as the gross intraday

[10] Most of the 140,000 individuals who day trade in a month do not do so for all months of the year. Analyzing these same data, Barber, Lee, Liu, and Odean (2014) estimate that approximately 450,000 individuals day trade in a typical year.

return from day trading. To calculate the net intraday return to day trading, we assume a 5 basis points (bps) commission on purchases, a 5 bps commission on sales, a 30 bps transaction tax on sales. (See appendix for details.)

We calculate a trader's past intraday returns to see whether quitting is sensitive to past performance. When calculating these returns we do not market adjust. We believe that when evaluating their past performance day traders are unlikely to market adjust because of their short holding periods and the low market betas of their portfolios. When evaluating future performance we do adjust for market returns because an obvious alternative to day trading would be to hold the market portfolio. Market adjusting does not qualitatively change our findings.

It is important to include both round-trip and one-sided trades to measure the performance of day trading. Focusing only on round-trip trades would yield a biased measure of performance if investors sell winners and hold losers (i.e., exhibit the disposition effect). For example, assume some day traders randomly buy and sell (random traders), while others close only winning investments while riding losers (disposition traders). Were we to analyze only the profits of intraday round-trip trades, it is clear that the disposition traders would have better intraday round-trip returns than the random traders merely because they postpone closing losing positions. Since the disposition effect is prevalent among Taiwanese investors and among day traders elsewhere,[11] it is important to include both round-trip and other trades when analyzing performance.

III. Rational and Behavioral Learning Models: Confirming Evidence

We begin by estimating the survival rate of day traders. Our trading data starts in 1992. To reasonably ensure that we are analyzing new day traders, we restrict our analysis to those who begin day trading after 1992. Our data ends in 2006 and thus is right-censored. We consider a trader to have quit day trading if we observe no day trading for 12 consecutive months. As a result of this requirement, we do not analyze day traders who begin day trading in 2006 since we cannot reliably observe whether they have quit.

---

[11] Barber, Lee, Liu, and Odean (2007) and Linnainmaa (2005) document, respectively, that individual Taiwanese investors and Finnish day traders exhibit the disposition effect.

In Figure 2, we present a five-year Kaplan-Meier survival function. We consider the survival of day traders who have day traded for at least 10 days and designate the first month when they hit the minimum 10-day threshold as their entry month. For many of these traders, day trading is a persistent activity. Only 2.5% drop out within one month, while survival rates at one, two, and three years are 44%, 24% and 15% respectively.

To test whether magnitude of past profitability affects the decision to quit day trading, we estimate the following Cox proportional hazard rate model,

$$h(t,x) = h_0(t)e^{\{XB\}}, \tag{1}$$

where X is a matrix of independent variables, B is a vector of coefficient estimates, $h_0(t)$ is the baseline hazard rate (i.e., the hazard rate when all covariates are equal to zero), and $h(t,x)$ is the hazard rate conditional on a set of covariates (x) at time t. In our application, a trader becomes at risk of quitting once he begins day trading.

Again, we restrict our analysis to day traders who have day traded for a minimum of 10 days. To assess the impact of past performance on quitting, we use net intraday returns. To estimate the impact of past returns on the propensity to quit day trading, we construct a series of 26 dummy variables corresponding to the following ranges in basis points of performance net of trading costs: (-∞, -90], (-90, -85], … , (25, 30], (30, ∞). When estimating the Cox proportional hazard rate model, we set the range (0, 5 bps] as the default category and include the remaining 25 dummy variables as covariates in our estimation. As control variables we include measures of past day trading activity: the log of the number of days with day trading activity, the log of the number of days since a trader's first day trade, and the log of the total volume of day trading. In the event history analysis, all independent variables are updated monthly.

The results of this analysis are presented in Figure 3. The horizontal axis of the figure represents net intraday return categories, while the vertical axis represents the hazard rate relative to the omitted profit category (net intraday return in the interval (0, 0.05]). As predicted by the learning models, the net intraday return is negatively related to the hazard rate. More profitable day traders are less likely to quit.[12]

[12] The relation between profits and the propensity to quit is similar regardless of whether we include control variables. However, the control variables are all reliably related to hazard ratios at the 1% significance level. Traders with more days of day trading experience are less likely to quit; a one standard deviation increase in the log of number of days of past day trading reduces the base case hazard rate to 0.43. Day

However, the effect is not linear. The propensity to quit is relatively insensitive to differences past net intraday returns in the -30 bps to 30 bps range. Traders with past net intraday returns in the (30 bps, infinity) range are more likely to quit than those with returns in the -30 bps to 30 bps range, though not as likely as traders with more negative returns.

In the domain of negative returns, the propensity to quit is quite sensitive to the magnitudes. For example, consider the impact on hazard rates of moving across four equidistant profit categories: (0, 0,05], (-0.30, -0.25], (-0.60, -0.55] and (-0.90, -0.85]. The first move, from just profitable to the low range of losses increases the hazard rate six percentage points (from 1.00 to 1.06); the second move, from the low range of losses to mid range losses, increases the hazard rate by an additional 30 percentage points (from 1.06 to 1.36), the third move, from the mid range of losses to a high range increases the hazard rate by an additional 36 percentage points (from 1.36 to 1.72).

In summary, these analyses provide strong evidence that traders learn about their own ability by trading. Those who profit are less likely to quit, though the effect is most pronounced for those with steep losses.

IV. Rational Learning Models: Disconfirming Evidence

To this point, we find support for the rational and behavioral learning models of investor behavior. Poor performers are more likely to quit day trading. The confirming evidence indicates learning is an important feature of financial markets. In this section, we argue that rational learning does not explain behavior of the large population of speculative investors for three reasons: aggregate performance is negative, experienced speculators lose money, and unprofitable speculators persist.

A. Aggregate performance is negative

To evaluate the performance of day traders relative to their opportunity cost of holding the market portfolio, we estimate abnormal returns by regressing the portfolio excess return (portfolio return less the risk-free rate) on the excess return on a value-

traders who have been trading longer are more likely to quit; a one standard deviation in the log of the number of days since a trader's first day trade increases the hazard rate to 1.24. Heavy day traders are more likely to quit; a one standard deviation increase in the log of past day trading volume increases the hazard rate to 1.11.

weighted market index (market return less the risk-free rate). We construct our own market index using market capitalization from the Taiwan Economic Journal (TEJ) and individual stock returns calculated from the TSE data. The intercept of this regression is our measure of abnormal returns.

In Table 1, we present the gross and net performance of all day traders from a calendar time regression of daily aggregate day trading returns on daily market returns. Though our performance analysis weights investors by the investments they make, we do not distinguish occasional day traders from active day traders in this preliminary analysis. We analyze the day trades and other trades of these investors in the months in which they day trade.

There are two reasons that including all trades in the month of day trading might positively bias our performance analysis. First, due to the disposition effect, day traders are more likely to close profitable positions. Thus, months in which we observe day trading are more likely to be profitable months. Second, it is possible that good investment performance leads to day trading (i.e., reverse causation). (In subsequent analyses, we identify day traders ex-ante to avoid these issues.) We are not concerned by these biases in this preliminary analysis since we document poor performance in aggregate.

In the second column of Table 1, we present the gross abnormal intraday returns of day traders. On average, day traders lose 7 basis points on their day trading before costs (t=-10.2). In the fourth column we see that trading costs more than triple the losses to 23.9 basis points per day. Moreover, we observe reliably negative gross and net performance in all years but 1992, when gross returns are indistinguishable from zero but net returns are reliably negative.

In aggregate, day trading is a losing proposition; day trading is an industry that consistently and reliably loses money. From an industrial organization perspective, it is difficult to understand how such an industry survives. For people to knowingly day trade, most must either be overconfident about their prospects of success or derive non-financial utility from the activity and knowingly suffer losses as a result. Finally, the poor aggregate performance of day trading is not consistent with the learning model of Mahani and Bernhardt (2007). In their model, novice speculators lose while the experienced profit, but aggregate performance should be positive and represent the equilibrium return to day trading. We discuss this issue in detail and explicitly test rational learning models after presenting results on cross-sectional variation in performance.

B. Experienced Day Traders Continue To Lose Money

A central feature of the learning model is the observation that bad traders quit. This raises two natural questions. First, do traders with a long history of losses continue to trade? Second, what fraction of day trading can we trace to traders with a history of losses?

Do traders with a long history of losses continue to trade? To investigate the relation between experience, past performance, and continued trading we categorize day traders into groups based upon past day trading experience and whether or not they have been profitable. Each month from 1993-2005, we sort traders based on the number of days they have engaged in day trading and on the cumulative profitability of their past intraday trades as of the end of the previous month. To have an accurate measure of day trading experience, we exclude traders who day traded in 1992. We exclude traders with fewer than 10 days of past day trading experience and sort the remaining traders into those with 10 to 49 days, 50 to 99 days, 100 to 199 days, 200 to 399 days, and 400 or more days of day trading experience. We further partition these experience groups into those with positive cumulative past intraday profits and those with negative or zero cumulative past intraday profits. We then measure the propensity for traders in each experience/profitability group to stop day trading for one month or for twelve consecutive months.

Table 2 Panel A reports the rate at which traders stop day trading for at least one month. Not surprisingly, traders with less experience are more likely to stop. For example, 41.98% of profitable traders and 43.19% unprofitable traders with 10 to 49 days of experience stop trading for at least one month while only 7.58% of profitable traders and 9.30% of unprofitable traders with 400 or more days of experience do so. As predicted by both rational and behavioral learning models, unprofitable traders are more likely to stop trading than profitable traders. Across the five experience bins, the unprofitable traders quit at rates ranging from 1.21 to 2.73 percentage points higher than profitable traders. What is more surprising is that the magnitude of the differences in the quitting rates of profitable and unprofitable traders, though statistically significant, is tiny.

Table 2 Panel B reports the rates at which traders stop day trading for at least twelve consecutive months. Naturally, these rates are lower than those for stopping for a month, but a similar pattern emerges as in Panel A. For example, 5.43% of profitable traders and 6.55% of unprofitable traders with 10 to 49 days of experience stop trading for at least twelve months. Of traders who have day traded 400 or more days, profitable traders have a 1.58% probability of quitting for twelve or more months and unprofitable traders have a 2.05% probability quitting. Using columns 2 – 5 and the last four rows of Panel B one can calculate that all previously unprofitable traders with 50 or more days of past day trading experience have a 95.3% probability of day trading again in the next 12 months, while previously profitable traders with similar experience have a 96.4% probability doing so. Traders are sensitive to losses, but not very sensitive.

What fraction of day trading can we trace to traders with a history of losses? To answer this question, each month we sort traders who day trade that month into three groups: novice day traders (traders with 1 to 9 days of day trading experience), unprofitable day traders (traders with 10 or more days of day trading experience and negative or zero life time net intraday profits prior to the month of sorting), and profitable day traders (traders with 10 or more days of day trading experience and positive life time net intraday profits prior to the month of sorting). In Figure 4 we graph the relative sizes of these three groups from 1995 through 2006 when measured by number of traders (Figure 4a) and trading volume in dollars (Figure 4b). Throughout this twelve-year period the fraction of profitable day traders is consistently about 5%. The fraction of unprofitable traders grows to over 2/3rds as more and more novice day traders are reclassified as profitable or unprofitable. As one would expect, profitable and unprofitable day traders account for proportionately more trading volume than novice day traders. In aggregate over the sample period, unprofitable traders account for 72% of trading volume and in the last several years of the sample they consistently account for about 80% of trading volume.

Thus, most day traders are unprofitable, previously unprofitable traders generate most day trading, and previously unprofitable day traders with considerable experience persist at day trading at almost the same high rate as profitable experienced day traders. These observations are not consistent with models of rational learning.

We sort day trader based on past profitability and experience. A reasonable concern is that if these sorts are not good indicators of future performance, traders might reasonably ignore them. However as we demonstrate in the next section, sorting day traders on experience and past profitability is remarkably effective at forecasting their future performance.

C. Event Time Analysis and the Persistence of Performance

So far we considered only intraday returns. This is appropriate given our focus on day trading. However, not all positions initiated by day traders are closed the day they are opened. Do day traders earn profits on positions that they hold beyond the close of trading? We address this question with an event time analysis of returns subsequent to purchases and sales.

On each day (event day t=0), we sort traders on profitability and experience over the last 365 calendar days (t-1 to t-365). Profitable traders earn a cumulative intraday net profit over the previous year that are positive, and unprofitable traders earn negative or zero profits. Experience bins are based on number of days of day trading during the previous 365 days with: Low experience (0-5 days), Exp. 2 (6-10), Exp. 3 (11-20), Exp. 4 (21-40), Exp. 5 (41-80), Hi Exp. (> 80). Thus, each day trader falls into 1 of 12 groups that represent the independent sorting into two profit groups and six experience groups.

To assess the profitability of trades for each group, we aggregate all purchases on event day t=0 for each group. We calculate the market-adjusted abnormal return on event day $t=0$ weighted by the value of stocks bought using the purchase prices and closing prices to calculate the intraday return. This step provides us with one buy return for event day $t=0$ for each calendar day during our sample period. We then calculate the mean market-adjusted abnormal return across all days over our sample period, $\left(MA_0^{buy}\right)$. For subsequent event days ($t=1,10$), we weight daily stock returns by the value of purchases on $t=0$. There is a similar calculation for the sales of each group.

Finally, we calculate the cumulative (market-adjusted) abnormal return on stocks bought less the cumulative (market-adjusted) abnormal return on stocks sold as:

$$CAR_t = \sum_{\tau=0}^{t}\left(MA_\tau^{buy} - MA_\tau^{sell}\right). \qquad (2)$$

The results for a 10 trading day event horizon are displayed in Figure 5. The most striking results in the graph are: 1) irrespective of experience, traders with past profits outperform those with past losses and 2) for traders with positive past profits, cumulative abnormal returns increase with experience, and 3) profitable traders with more than 40 days of day trading experience in the last year earn more than enough to cover 40 basis points of round trip transaction costs. Most of this profit is earned on the day of the transaction.

We look at the persistence of profitability from another angle by calculating, as in Table 1, gross and net abnormal intraday returns. We partition day traders daily using the same criteria as for Figure 5. Results are reported in Table 3.

As in Table 1, these results are based on daily calendar time regressions. Thus we assume that each day—but not each trader or each security—is a separate independent observation. The reader may notice that many of the reported t-statistics are quite large. For example, we reject the null hypothesis that the alpha of the net returns of previously unprofitable day traders with more than 80 days of experience is 0 with a t-statistic of -57.90. This may strike some readers as a remarkably large t-statistic for a sample of only 3,728 observations. However, the aggregate performance of this group of day traders is remarkably persistent. An alternative, non-parametric, approach to testing the performance of this group of traders is to employ a binomial test. Previously unprofitable day traders with more than 80 days of experience, earn positive aggregate net returns on only 920 out of 3,728 days. For p = 0.5, the probability of 920 or fewer successes in 3,728 trials is less than 6.2×10^{-220}.

In Table 3, we see that 1) Irrespective of experience, traders with past profits outperform those with past losses, 2) for traders with positive past profits, cumulative abnormal returns increase with experience, and 3) profitable traders with more than 40 days of day trading experience in the last year earn more than enough to cover their transaction costs.

These results confirm that an extensive history of profitability is a strong predictor of future profitable. However, very few traders are predictably profitable. In the last column of Table 3, we see that only 9.81% (3.20%+6.61%) of day trading volume is generated by predictably profitable day traders. From column 8, we can calculate that these predictably profitable traders constitute less than 3% of all day traders active on an average day.[13]

V. New Day Trader Entry

We argue that it is not rational for the average investor who begins day trading to believe doing so will be profitable. Furthermore, it is not rational for the average day trader who is motivated by profits to continue to trade actively despite a history of poor performance. Of course, if an investor mistakenly believes that day trading has positive expected lifetime profits, entry may be rational conditional on that mistaken belief. But it is not rational for hundreds of thousands of investors to have materially mistaken beliefs. We believe that many investors begin day trading because they are overconfident in their ability to earn profits day trading. We conjecture that two factors might influence beliefs and the entry decisions of new day traders (and the beliefs and decisions of existing traders who continue trading): recent market returns and recent day trader returns.

Strong market returns may attract the attention of and increase the overconfidence of investors. Statman, Thorley, and Vorkink (2006) document that across countries, volume increases following strong market returns; they attribute increased volume to increase overconfidence consistent with the theoretical predictions of Gervais and Odean (2001). Potential day traders may also, mistakenly,[14] believe that day trading profits are highly correlated with market returns. Thus strong recent market returns may lead to increased day trader entry because 1) investors are paying more attention to the market, 2) investor overconfidence increases with strong market returns, and 3) investors may believe that strong market returns are indicative of strong day trading opportunities.

[13] Because highly experienced profitable day traders trade more actively than others, they are overrepresented when measured on a daily basis. Employing different methodology, Barber, Lee, Liu, and Odean (2014) estimate that less than 1% of individuals who day trade over the course of a year are predictably profitable.
[14] As reported in Table 1, the beta of aggregate day trading returns is only 0.26.

Recent day trader performance may also influence entry decisions because potential day traders learn of the recent successes of others and decide to try day trading themselves. Han and Hirshleifer (2016) argue this social transmission channel can explain investors' willingness to invest in active strategies. We thus expect the aggregate performance of day traders will also draw new entrants into the market above and beyond general market returns.

To test these conjectures, we estimate the following time series regression using 712 weekly observations:

$$ND_t = a + \sum_{\tau=t-1}^{t-12} b_\tau R_{m\tau} + \sum_{\tau=t-1}^{t-12} c_\tau DP_\tau + dt + \mu_{dow} + \mu_{mth} + e_t \tag{3}$$

ND_t is the number of new day traders in week t, R_{mt} is the value-weighted Taiwan market return for week t, and DP_t is the aggregate performance of day traders in week t. We standardize these variables to facilitate the interpretation of regression results. New day traders (ND_t) are traders who execute a round-trip day on day t but never did so prior to day t. We drop the first year (1992) from this analysis since we would misidentify many 1992 day traders as new. Day trader performance (DP_t) is measured as the residual from a univariate regression of weekly day trader performance on the weekly market return (and is thus orthogonal to market returns). As controls we include fixed effects for the number of trading days in the week (μ_{dow}) and calendar month (μ_{mth}) and a time trend (t); results are not sensitive to the inclusion of these control variables. To address concerns regarding serial correlation, we estimate these regressions using Newey-West correction assuming six lags.[15]

In addition to analyzing the entry of new day traders, we estimate two analogous regressions where the independent variable is either the number of repeat day traders (day traders who trade on day t but also traded prior to day t, RP_t) or the ratio of new day traders to total day traders on day t, $ND_t/(ND_t+ RP_t)$.

We present the key results of this regression in the six graphs of Figure 6. Each column of the figure presents two graphs, one for the 12 estimated coefficients on the lagged market return (top graph) and one for 12 estimated coefficients on the lagged

[15] We also estimate models including 12 lags of the dependent variable and find qualitatively similar results for new day traders and repeat day traders.

performance of day traders (bottom panel). The three columns of the figure represent the three regressions where the independent variable is alternatively new day traders (blue lines in left column), repeat day traders (red lines in middle column), or the ratio of new to all day traders (green lines in right column).

New day trader entry and repeat day trader activity are positively correlated with both recent market returns and recent aggregate day trader performance. The influence of recent market returns on entry and repeat activity is more persistent than that of recent aggregate day trader performance. However at a lag of only one week, aggregate day trader performance is more influential. This suggests that if lagged aggregate day trader performance is influencing new entry and repeat activity through social channels, those channels focus on very recent performance. This is consistent with day traders bragging to others about their most recent successes.

VI. Discussion

Our data are remarkably well suited for testing models of rational learning such as Mahani and Bernhard (2007). Mahani and Bernhard write that their "prototypical novice speculator is the Japanese hairdresser Kiyoshi Wakino" who day trades between giving haircuts (p. 1317). We observe the day trading of hundreds of thousands of investors over a seventeen-year period. And while our day traders are Taiwanese and certainly not all hairdressers, it is probable that—like Kiyoshi Wakino—many of our novice day traders pursue trading in addition to, if not during, a regular job. Despite the size and appropriateness of our data, our results do not support the rational learning models. In Mahani and Bernhard's model, day trading is, in aggregate, profitable because skilled day traders are able to take advantage of the insensitivity of liquidity traders to price and the willingness of competitive market-makers to forego a profit. In Taiwan, day traders, in aggregate, lose money. It is not rational for a risk-averse investor with no special claim to superior ability to undertake day trading in hopes of discovering that he is amongst the chosen few. Furthermore, it is not rational for day traders who have incurred persistent losses to continue day trading for the purpose of learning about their ability.

So why do investors take up day trading and why do so many persist in the face of losses? We consider three broadly defined answers to this question.

First, it could be the case that day traders do not have standard risk-averse preferences; they may be risk-seeking or, as suggested by Kumar (2009), attracted to investments with highly skewed investments, such as lotteries, that have negative expected payoffs but a small probability of a large payoff. However, the day trading profits that we document are similar in magnitude to, and far less prevalent than, the losses. Unlike lottery winners, day traders must succeed on repeated gambles in order to achieve overall success. Such repeated gambles do not tend to generate highly skewed distributions. Furthermore, daily day trading returns have a negative mean, and yet lower variance and less right-hand skewness than the average returns of Taiwanese stocks. Define the annual day trading return as the sum of the returns earned on each day of day trading. For traders with a minimum of ten days of day trading, the skewness of the annual return is -0.22 (i.e., modestly negatively skewed). In contrast, when we calculate the skewness of annual returns across individual stocks listed on the TSE from 1981 to 2009, the coefficient of skewness is positive in all but one year and averages 2.36. Thus, a risk or lottery seeker could better maximize his utility, with far less effort, by simply buying and holding a single volatile stock.

Second, day traders may be overconfident in their prior beliefs about their abilities and biased in the way they learn. Several papers (e.g., Odean (1998, 1999), Barber and Odean (2000, 2001)) argue that overconfidence causes investors to trade more than is in their own best interest. Overconfident day traders may simply be bearing losses that they did not anticipate. While novice day traders undoubtedly realize that other day traders lose money, stories of successful day traders may circulate in non-representative proportions, thus giving the impression that success is more frequent than it is. Once investors undertake day trading, their prior overconfidence may be reinforced through biased learning as in Gervais and Odean (2001). Furthermore, day traders who earn gross profits but net losses may not fully consider trading costs when assessing their own ability.

Third, day traders may trade for non-financial motivations including entertainment, a taste for gambling, and the desire to impress others (see, e.g. Grinblatt and Keloharju (2009)). Some investors may enjoy the process of day trading so much that they are willing to persist in the face of regular losses. Some investors may be attracted to the casino like qualities of day trading with its frequent bets, wins, and losses.[16] Some investors may choose to day trade in hopes of impressing others.[17]

We are unable to explicitly test whether day traders are motivated by overconfidence rather than the desire for entertainment, gambling, or to impress others. Nor is there reason to believe that overconfidence and non-financial motivations are mutually exclusive. Quite to the contrary, entertainment, gambling, and the desire to impress others are all likely to be more attractive reasons to trade if one is overconfident about one's likelihood of success.

In Mahani and Bernhard's model, "*all* speculators are made worse off if some speculators are *slightly* overconfident" (p. 1315). Our results are consistent with this prediction. If heavy day traders persist in trading due to overconfidence, then that overconfidence is detracting from their own welfare and that of other speculators. The welfare of the heavy traders themselves is diminished because, on average, they earn net losses; the welfare of other speculators is diminished because, on average, heavy traders earn gross profits thereby reducing the average returns of other investors. The beneficiaries of this overconfidence are brokerage firms—through commissions—and the government—through the transaction tax.

VII. Conclusion

We test predictions of models of learning by rational traders and find clear evidence that the decision to continue day trading is influenced by their previous day trading returns. Nevertheless, rational models of learning do not explain all or even most day trading. Only the most experienced previously profitable day traders—less than 3% of active day traders on an average day--earn predictably positive net returns, and yet most experienced unprofitable day traders continue to trade and continue to reap losses.

[16] Kumar (2009) shows a correlation between the propensity to gamble and the types of investment decisions U.S. investors make. Barber, Lee, Liu, and Odean (2009) document that the introduction of a National Lottery in Taiwan coincided with a significant drop in trading volume on the Taiwan Stock Exchange. Grinblatt and Keloharju (2009) document that investors prone to sensation seeking trade more frequently.

[17] Several papers argue that investment decisions are influenced by social concerns, for example, Barber, Heath, Odean (2003), Statman (2004), and Hong and Kacperczyk (2009).

Nearly 3/4ths of day trading can be traced to traders with a history of losses. Persistent trading in the face of losses is inconsistent with models of rational learning. So, too, is the decision to try day trading when ex-ante expected lifetime profits are negative. For prospective day traders, “trading to learn” is slow learning.

References

Barber, B.M., Heath, C., Odean, T., 2003. Good rationales sell: reason-based choice among group and individual investors in the stock market. Management science 33, 1636-1652.

Barber, B.M., Lee, Y.T., Liu, Y.J., Odean, T., 2007. Is the aggregate investor reluctant to realize losses? Evidence from Taiwan. European Financial Management 13(3), 423-447.

Barber, B.M., Lee, Y.T., Liu, Y.J., Odean, T., 2008. Just how much do individual investors lose by trading? Review of Financial Studies 22(2), 609-632.

Barber, B.M., Lee, Y.T., Liu, Y.J., Odean, T., 2014. The cross-section of speculator skill: Evidence from day trading. Journal of Financial Markets 18, 1-24.

Barber, B.M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. Journal of Finance 55(2), 773-806.

Barber, B.M., Odean, T., 2001. Boys will be boys: Gender, overconfidence, and common stock investment. Quarterly journal of economics 116(1), 261-292.

Barber, B. M., Odean, T., & Zhu, N. (2008). Do retail trades move markets? *The Review of Financial Studies*, *22*(1), 151-186.

Benos, A., 1998, Overconfident speculators in call markets: Trade patterns and survival. Journal of Financial Markets, 1, 353-383.

Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. Review of economic studies 72(2), 287-312.

Caballé, J., Sákovics, J., 2003. Speculating against an overconfident market. Journal of Financial Markets 6(2), 199-225.

Campbell, J.Y., Ramadorai, T., Ranish, B., 2014. Getting better or feeling better? How equity investors respond to investment experience. NBER Working Paper 20000.

Chae, J., Wang, A., 2003. Who makes markets? Do dealers provide or take liquidity. Unpublished working paper, Sloan School of Management, MIT.

Chiang, Y.M., Hirshleifer, D., Qian, Y., Sherman, A.E., 2011. Do investors learn from experience? Evidence from frequent IPO investors. The Review of Financial Studies 24(5), 1560-1589.

Choi, J.J., Laibson, D., Madrian, B.C., Metrick A., 2009. Reinforcement learning and savings behavior. The Journal of finance 64(6), 2515-2534.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor psychology and security market under-and overreactions. Journal of Finance 53(6), 1839-1885.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 2001. Overconfidence, Arbitrage, and Equilibrium Asset Pricing. Journal of Finance 56(3), 921-965.

De, S., Gondhi, N.R., Pochiraju, B., 2010. Does sign matter more than size? An investigation into the source of investor overconfidence. Available at SSRN: http://ssrn.com/abstract=1657926.

Deaves, R., Lüders, E., Luo, G.Y., 2008. An experimental test of the impact of overconfidence and gender on trading activity. Review of finance 13(3), 555-575.

Dorn, A.J., Dorn, D., Sengmueller, P., 2014. Trading as gambling. Management Science 61(10), 2376-2393.

Dorn, D., Sengmueller, P., 2009. Trading as entertainment? Management Science 55(4), 591-603.

Feng, L., Seasholes, M.S., 2005. Do investor sophistication and trading experience eliminate behavioral biases in financial markets? Review of Finance 9(3), 305-351.

Gao, X., Lin, T.C., 2014. Do individual investors treat trading as a fun and exciting gambling activity? Evidence from repeated natural experiments. Review of Financial Studies 28(7), 2128-2166.

Garvey, R., Murphy, A., 2001. How profitable day traders trade: An examination of trading profits. Working paper, University College Dublin, Dublin, Ireland.

Garvey, R., Murphy, A., 2005. The Profitability of Active Stock Traders. Journal of Applied Finance 15(2), 93.

Gervais, S., Odean, T., 2001. Learning to be overconfident. Review of financial studies 14(1), 1-27.

Glaser, M., Weber, M., 2007. Overconfidence and trading volume. The Geneva Risk and Insurance Review 32(1), 1-36.

Grinblatt, M., Keloharju, M., 2000. The investment behavior and performance of various investor types: a study of Finland's unique data set. Journal of financial economics 55(1), 43-67.

Grinblatt, M., Keloharju, M., 2001. What makes investors trade? Journal of Finance 56(2), 589-616.

Grinblatt, M., Keloharju, M., 2009. Sensation seeking, overconfidence, and trading activity. Journal of Finance 64(2), 549-578.

Harris, J.H., Schultz, P.H., 1998. The trading profits of SOES bandits. Journal of Financial Economics 50(1), 39-62.

Hong, H., Kacperczyk, M., 2009. The price of sin: The effects of social norms on markets. Journal of Financial Economics 93(1), 15-36.

Kaniel, R., Saar, G., & Titman, S. (2008). Individual investor trading and stock returns. *The Journal of Finance*, *63*(1), 273-310.

Kelley, E. K., & Tetlock, P. C. (2013). How wise are crowds? Insights from retail orders and stock returns. *The Journal of Finance*, *68*(3), 1229-1265.

Kumar, A., 2009. Who gambles in the stock market? Journal of Finance 64(4), 1889-1933.

Kyle, A., F.A. Wang, 2007 Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?, Journal of Finance 52(5), 2073-2090.

Linnainmaa, J.T., 2003. The anatomy of day traders. AFA 2004 San Diego Meetings. Available at SSRN: http://ssrn.com/abstract=472182 or doi:10.2139/ssrn.472182

Linnainmaa, J.T., 2005. The individual day trader. University of California, Berkeley, working paper.

Linnainmaa, J.T., 2010. Do limit orders alter inferences about investor performance and behavior? Journal of Finance 65(4), 1473-1506.

Linnainmaa, J.T., 2011. Why do (some) households trade so much? Review of Financial Studies 24(5), 1630-1666.

Lubensky, D., Smith, D., 2016. Learning from speculating and the no trade theorem. Available at SSRN: https://ssrn.com/abstract=2049580.

Mahani, R., Bernhardt, D., 2007. Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity. Journal of Finance 62(3), 1313-1340.

Nicolosi, G., Peng, L., Zhu, N., 2009. Do individual investors learn from their trading experience? Journal of Financial Markets 12(2), 317-336.

Odean, T., 1998. Volume, volatility, price, and profit when all traders are above average. Journal of Finance 53(6), 1887-1934.

Odean, T., 1999. Do investors trade too much? American economic review 89(5), 1279-1298.

Seasholes, M.S., Wu, G., 2007. Predictable behavior, profits, and attention. Journal of Empirical Finance 14(5), 590-610.

Seru, A., Shumway, T., Stoffman, N., 2009. Learning by trading. Review of Financial Studies 23(2), 705-739.

Statman, M., 2004. What do investors want? Journal of Portfolio Management 30(5), 153-161.

Strahilevitz, M.A., Odean, T., Barber, B.M., 2011. Once burned, twice shy: How naïve learning, counterfactuals, and regret affect the repurchase of stocks previously sold. Journal of Marketing Research 48(SPL), S102-S120.

APPENDIX: Details of Return and Profit Calculations

We calculate the intraday return from day trading on day t for a particular group (g) of investors weighted by the value of investors' trades:

$$r_{gt} = \frac{\sum_i \sum_j \left(S^L_{ijt} - B^L_{ijt}\right) + \left(S^S_{ijt} - B^S_{ijt}\right)}{\sum_i \sum_j \left(B^L_{ijt} + S^S_{ijt}\right)}, \tag{A1}$$

where *B* and *S* denote the value of buys and sells (with superscripts L and S for long and short transactions, respectively) on day *t* in stock *i* by investor *j*. For long positions, the sales value $\left(S^L_{ijt}\right)$ is the value based on the actual transaction price or the closing price if the long position is not closed out prior to the end of trading. For short positions, the purchase value $\left(B^S_{ijt}\right)$ is the value based on the actual transaction price or the closing price if the short position is not closed out prior to the end of trading. The numerator in equation A1 is the intraday gross profit of trader *i* on day *t*.

Consider a concrete example where an investor buys a stock for \$100 and sells later in the day for \$102. On the same day, the investor shorts a stock (the same stock or a different stock) for \$100 and later covers the short with a purchase at \$97. The investor makes profits of \$5 = (102-100) + (100-97). We scale the dollar profits by the total value of the opening positions, \$200 = \$100 + \$100. Thus, we assume the investor put \$200 of capital at risk and earned an intraday return of \$5/\$200 = 2.5%. This is an accurate representation of the returns if the investor trades in parallel (i.e., both positions are open at the same time). For investors who trade sequentially, we correctly calculate dollar profits of \$5, but the capital at risk would be \$100 rather than \$200 as the \$100 would be deployed sequentially. Thus, we always estimate the correct sign of returns, but for day traders who trade sequentially our return estimates are biased toward zero. In addition, we do not know the extent to which traders use leverage, which would increase the magnitude of returns for both gains and losses, but again the sign of the gains and losses would be the same as those in our calculations. In summary, the sign of the day trading returns that we calculate is accurate, though the magnitudes may differ because of sequential trading or the use of leverage.

When we calculate net returns and net profits, we deduct a 5 bps commission for all trades (10 bps round-trip commission) and a 30 bps transaction tax for sales. Put differently, buys cost 5 bps (C_b) and sells cost 35 bps (C_s). We also increase the capital requirements to reflect the total cost of the opening positions:

$$r_{gt}^{net} = \frac{\sum_i \sum_j \left(S_{ijt}^L - B_{ijt}^L\right) + \left(S_{ijt}^S - B_{ijt}^S\right) - c_b * \left(B_{ijt}^L + B_{ijt}^S\right) - c_s\left(S_{ijt}^L + S_{ijt}^S\right)}{\sum_i \sum_j \left(B_{ijt}^L + S_{ijt}^S\right) + c_b B_{ijt}^L + c_s S_{ijt}^S}, \tag{A2}$$

Continuing our example from above, the net return for the trader would be:

$$\frac{(102-100)+(100-97)-0.0005(100+97)-0.0035(102+100)}{(100+100)+0.0005*100+0.0035*100} = \frac{4.19}{200.40} = 2.09\% .$$

Note the net return (2.09%) is roughly 40 bps (the total round-trip trading costs of 10bps in commissions and 30 bps in transaction tax) less than the gross return (2.50%). The shortfall is slightly greater than 40 bps because we also increase the capital required to open the positions. The numerator in equation A2 is the intraday net profit of trader *i* on day *t*.

Table 1. Gross and Net Abnormal Intraday Returns from Day Trading: 1992 to 2006

This table presents the daily percentage alpha from aggregate day trading of day traders. Day trading is defined as round-trip trades by the same stock/investor/day. Each day all investor's trades are included in this analysis if the trader made at least one day trade in the calendar month. The alphas are estimated using the following regression of daily returns: (Rpt-Rft)=α_p+β_p(Rmt-Rft)+ε_{pt}, where Rpt, Rmt, and Rft are the portfolio return, market return, and risk-free return (respectively). The gross day trading return is calculated from daily round-trip trades plus the intraday returns on open trades; an open trade is a trade made during the day that results in an outstanding position at the close of the day. The net day trading return assumes a 10 bps round-trip commission and a 30 bps transaction tax on sales.

	Gross		Net			
	α(%)	t-stat	α(%)	t-stat	Beta	R-Sq
All Years	-0.070	-10.20	-0.239	-34.78	0.26	43%
1992	0.018	0.57	-0.091	-2.91	0.37	49%
1993	-0.089	-3.40	-0.206	-7.80	0.33	59%
1994	-0.119	-5.03	-0.275	-11.65	0.32	58%
1995	-0.047	-2.18	-0.221	-10.23	0.24	45%
1996	-0.088	-4.90	-0.245	-13.58	0.26	48%
1997	-0.089	-3.27	-0.271	-10.00	0.26	39%
1998	-0.067	-3.13	-0.256	-11.99	0.21	43%
1999	-0.053	-1.91	-0.224	-8.04	0.25	42%
2000	-0.007	-0.18	-0.190	-5.23	0.23	36%
2001	-0.088	-2.43	-0.279	-7.68	0.24	37%
2002	-0.104	-3.47	-0.289	-9.67	0.23	39%
2003	-0.100	-4.39	-0.280	-12.34	0.22	38%
2004	-0.062	-2.51	-0.244	-9.90	0.22	39%
2005	-0.087	-5.00	-0.280	-16.23	0.28	41%
2006	-0.089	-4.07	-0.277	-12.82	0.26	38%

Table 2. Quitting Day Trading Conditional on Experience and Past Profitability

Day trading is defined as round-trip trades by the same stock/investor/day. In panel A, we present the percentage of traders who day trade in month t-1 but not in month t. In Panel B, we present percentage of traders who day trade in month t-1 but not in months t through t+11 We sort traders on experience (rows) and profitability (columns). Experience is the number of previous days on which a trader has day traded through month t-1. Only traders with 10 or more days of day trading experience are included. Profitable day traders are those with mean daily intra-day returns that are positive from the first month in which they make an intraday roundtrip trade through month t-1. The analyses begin in 1993 and exclude traders who day traded in 1992. N is the number of trader/month observations.

Experience (X)	Unprofitable Traders		Profitable Traders		Difference	*t*-statistic
Partition	Mean	N	Mean	N	in Means	(difference in means)
Panel A. Percentage of Day Traders who Day Trade in Month t-1 but do not Day Trade in Month t						
10 d <= X < 50 d	43.19%	5,578,277	41.98%	540,228	1.21%	17.19
50 d <= X < 100 d	29.82%	2,460,476	27.08%	154,435	2.73%	23.41
100 d <= X < 200 d	22.47%	2,018,002	20.06%	123,009	2.41%	20.45
200 d <= X < 400 d	15.95%	1,344,396	13.67%	90,711	2.28%	19.27
400 d <= X	9.30%	832,988	7.58%	81,713	1.72%	17.54
Panel B. Percentage of Day Traders who Day Trade in Month t-1 but not in Months t to t+11						
10 d <= X < 50 d	11.73%	5,578,277	11.41%	540,228	0.33%	7.18
50 d <= X < 100 d	6.55%	2,460,476	5.43%	154,435	1.12%	18.73
100 d <= X < 200 d	4.62%	2,018,002	3.65%	123,009	0.98%	17.61
200 d <= X < 400 d	3.24%	1,344,396	2.43%	90,711	0.81%	15.18
400 d <= X	2.05%	832,988	1.58%	81,713	0.47%	10.20

Table 3. Day Trader Performance conditional on Prior Year's Day Trading Experience and Profitability

Day Traders are grouped based upon how many days they engaged in day trading and whether their intraday returns were profitbale during the previous 365 days. The alphas are estimated using the following regression of intraday daily returns: (Rpt - Rft) = αp + βp(Rmt - Rft) + εpt, where Rpt, Rmt, and Rft are the portfolio return, market return, and riskfree return (respectively). The gross day trading return is calculated from daily round-trip trades plus the intraday returns on open trades; an open trade is a trade made during the day that results in an outstanding position at the close of the day. The net day trading return assumes a 10 bps round-trip commission and a 30 bps transaction tax on sales. "Day Trade / All Trade" is the fraction of the group's trading that is round-trip day trades. The last column presents the share of all day trading accounted for by each group.

# Days of Day Trading Experience in Previous Year	Gross		Net						
	α(%)	t-stat	α(%)	t-stat	Beta	R-Sq	Average Daily N	Day Trades / All Trades	Share of All Day Trading by Volume
Panel A. Previously Profitable Day Traders									
0 d <= X < 5 d	-0.066	-7.64	-0.140	-16.26	0.33	46%	14936	9.73%	7.81%
6 d <= X < 10 d	0.001	0.18	-0.107	-14.57	0.29	46%	2110	18.90%	2.12%
11 d <= X < 20 d	0.053	7.75	-0.070	-10.32	0.27	47%	1976	22.53%	2.50%
21 d <= X < 40 d	0.120	19.33	-0.019	-3.14	0.24	47%	1732	26.34%	2.90%
41 d <= X < 80 d	0.197	35.39	0.035	6.34	0.21	46%	1455	33.36%	3.20%
80 d <= X	0.333	72.65	0.118	26.03	0.17	45%	1724	50.52%	6.61%
Panel B. Previously Unprofitable Day Traders									
0 d <= X < 5 d	-0.139	-16.10	-0.213	-24.60	0.33	46%	46078	8.95%	20.55%
6 d <= X < 10 d	-0.167	-20.75	-0.270	-33.60	0.31	47%	9468	18.52%	6.81%
11 d <= X < 20 d	-0.171	-21.98	-0.289	-37.26	0.31	47%	10335	23.99%	8.77%
21 d <= X < 40 d	-0.171	-23.05	-0.310	-41.83	0.30	48%	10308	30.47%	10.47%
41 d <= X < 80 d	-0.166	-23.71	-0.332	-47.56	0.28	49%	9184	37.81%	11.34%
80 d <= X	-0.132	-21.88	-0.348	-57.90	0.25	50%	9639	50.72%	16.91%

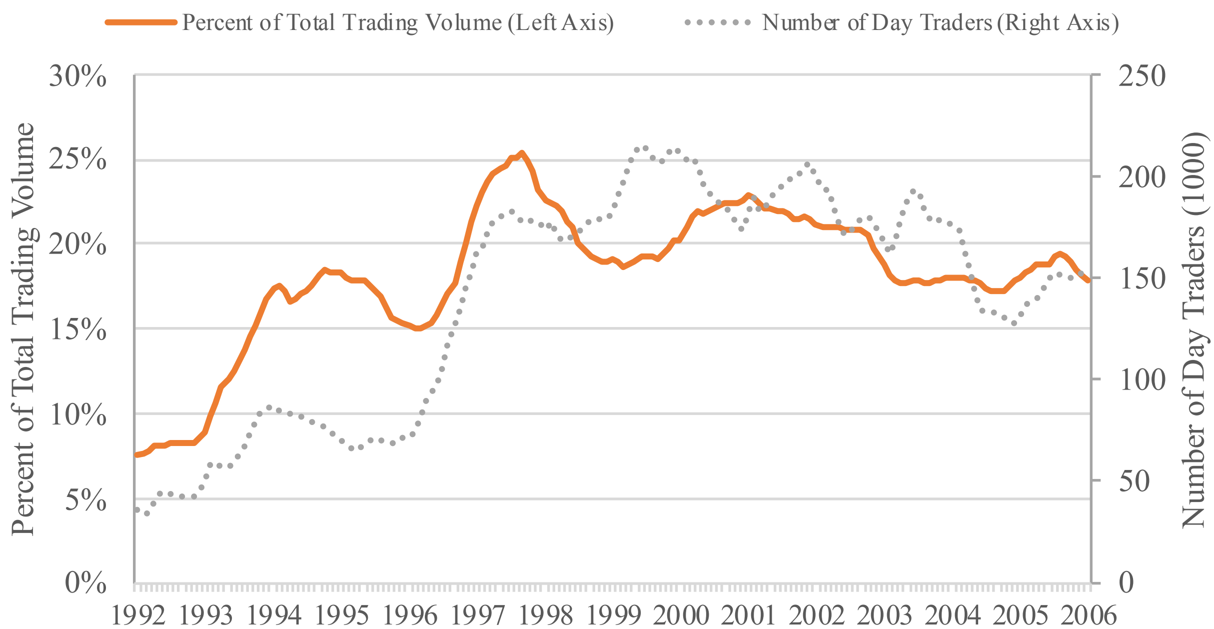

Figure 1. Day Trading as a Percent of Total Volume and Number of Individual Day Traders
The figure presents the 12-month moving average for (1) the number of individual investors who engage in day trading and (2) day trading as a percent of total trading volume.

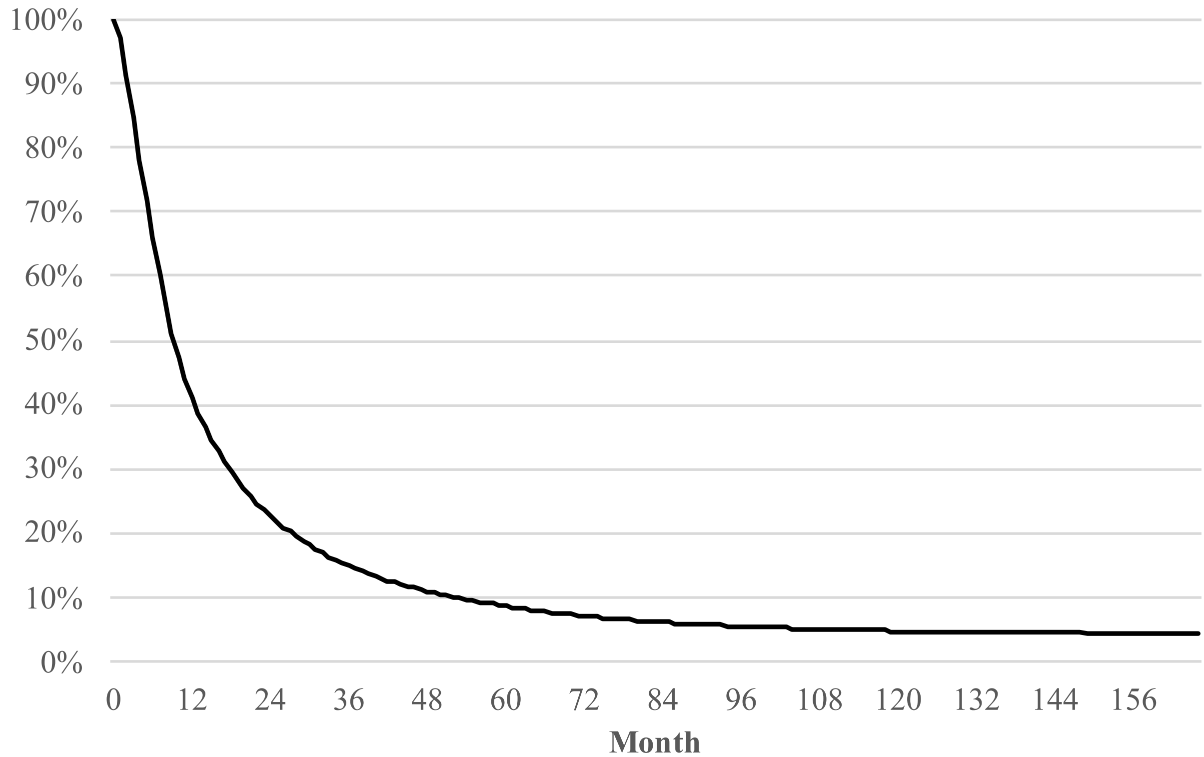

Figure 2. Day Trader Survival Function
Observations are monthly. Entry in month 0 is defined as the month when a day trader has 10 or more days of day trading. Quitting (the failure event) is defined as the first month in which we observe no day trading in 12 consecutive months.

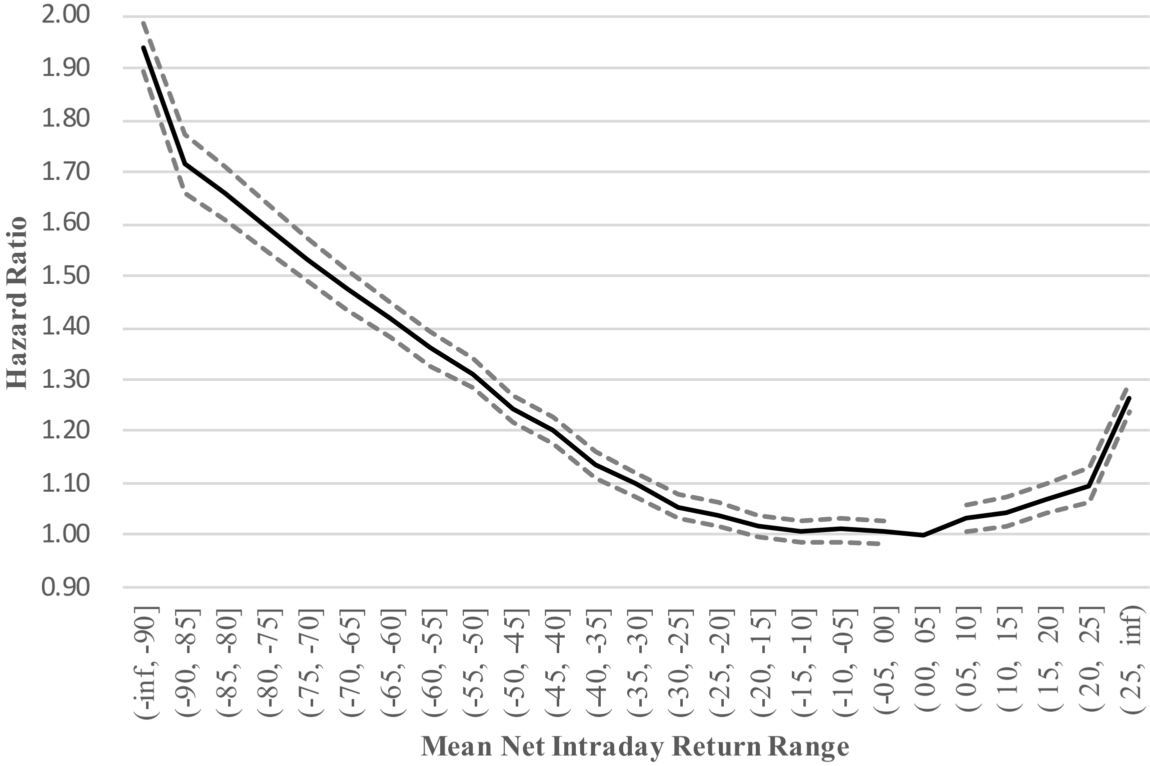

Figure 3. Hazard Ratio for Quitting Day Trading Conditional on Past Performance, 1993-2005. Observations are monthly and exclude traders who day traded in 1992 or started in 2006. Entry in month 0 is defined as the month when a day trader has 10 or more days of day trading. Quitting (the failure event) is defined as the first month in which we observe no day trading in 12 consecutive months. The figure reports the hazard ratio for quitting (solid line) and the 95% confidence interval (dashed lines) for different return categoreis relative to the default category of (0,5] basis points (bps) where the hazard ratio is equal to one by construction. Mean net intraday return range is the investor's average daily return on days with trades from the first month of day trading to the end of the prior month.

Figure 4a. Percentage of Day Traders

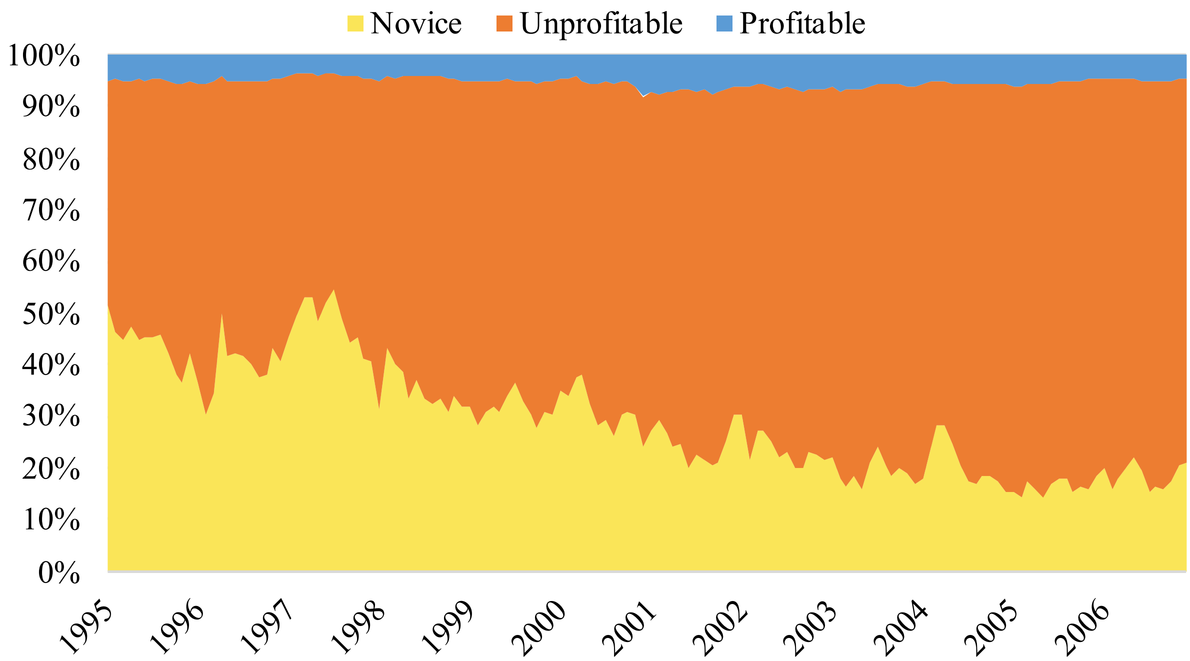

Figure 4b. Percentage of Day Trading Volume

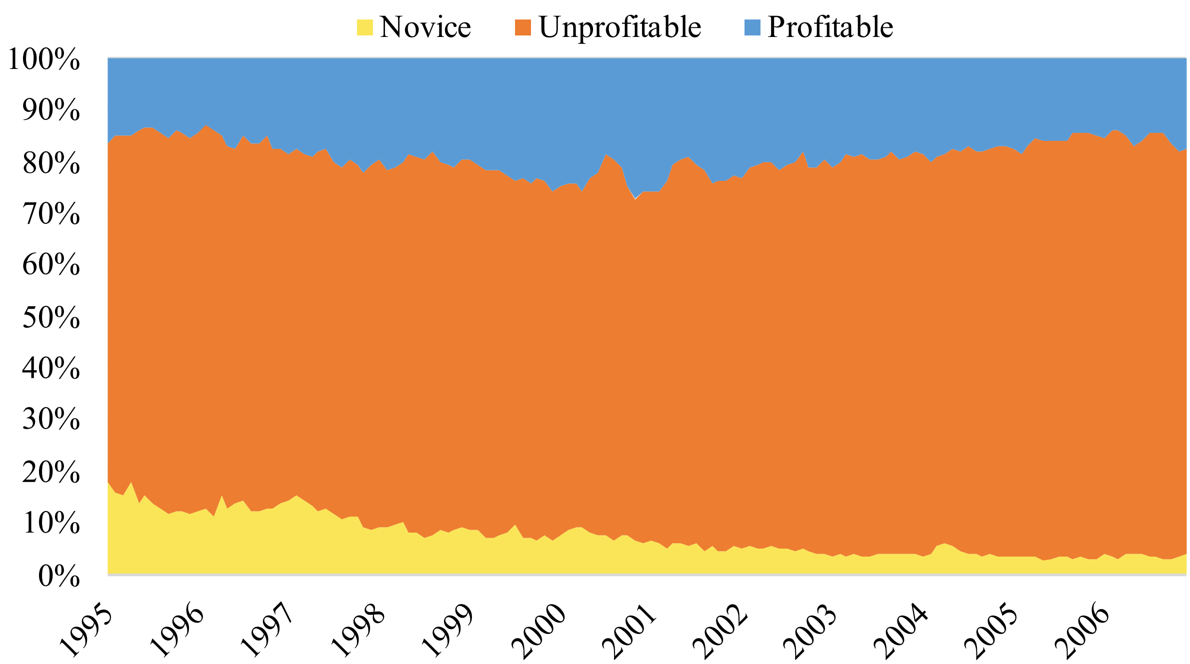

Figure 4. Day Trading by Novice, Profitable, and Unprofitable Day Traders
Novice day traders are those with less than 10 days of day trading.

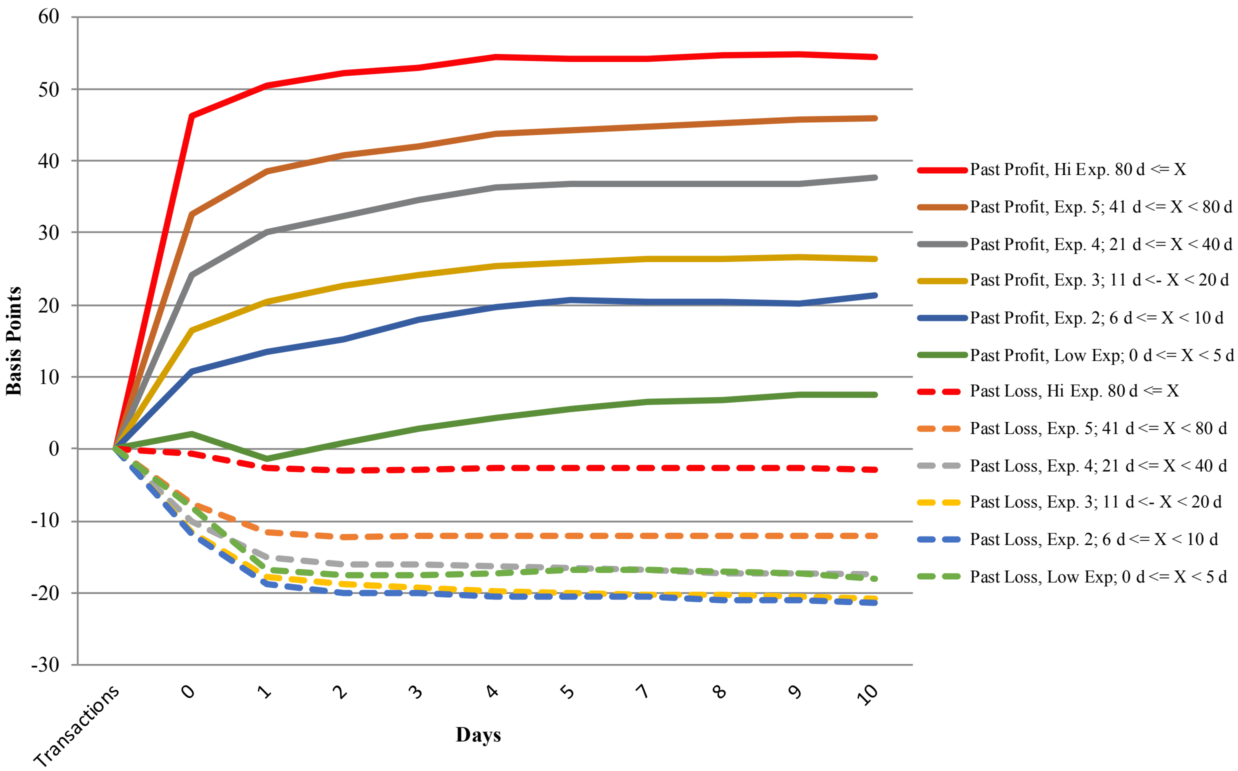

Figure 5. Cumulative abnormal returns (CARs) in Event Time conditional on Past Profitability and Experience, 1993-2005

On event day 0, traders who day trade that day are sorted based on net profitability of intraday trades and day trading experience over previous year (*t-1* through *t-365*). CARs are calculated from the time of each transaction (on event day 0) through the close of trading on event days 0 through 10. Profitable (unprofitable) traders are those with cumulative intraday net profits over the previous 365 calendar days that are positive (negative or zero). Experience bins are based on number of days of day trading during the previous 365 days with: Low experience (0-5), Exp. 2 (6-10), Exp 3 (11-20), Exp 4 (21-40), Exp 5 (41-80), Hi Exp (> 80).

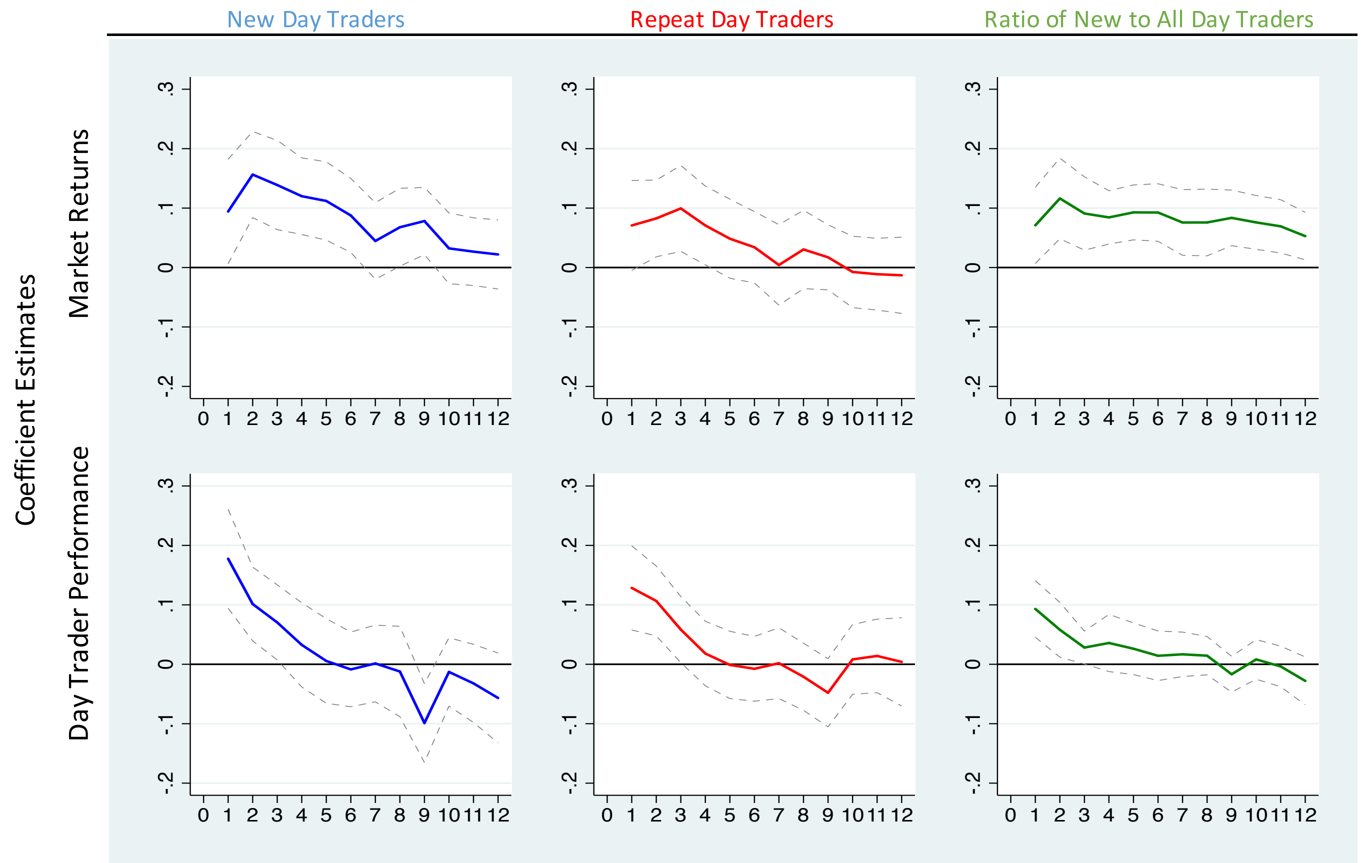

Figure 6: The Effect of Market Returns and Day Trader Performance on Day Trading Activity

The figure depicts the results of three regressions using weekly observations from 1993-2006. The left column (blue lines, labeled New Day Traders) presents the key results of the following regression in two graphs:

$$ND_t = a + \sum_{\tau=t-1}^{t-12} b_\tau R_{m\tau} + \sum_{\tau=t-1}^{t-12} c_\tau DP_\tau + \mathrm{dt} + \mu_{dow} + \mu_{mth} + e_t$$

where ND_t is the number of new day traders in week t, R_{mt} is the value-weighted Taiwan market return for week t, and DP_t is the aggregate performance of day traders in week t. We standardize these variables to facilitate the interpretation of regression results. The top graph plots the estimated coefficients on lagged market returns, and the bottom graph plots the estimated coefficients on lagged day trader performance. As controls we include fixed effects for the number of trading days in the week (μ_{dow}) and calendar month (μ_{mth}) and a time trend (t). Standard errors are calculated using Newey-West correction assuming six lags.

The center column (red lines, labeled Repeat Day Traders) presents the key results of a similar regression with repeat day traders (RDt) as the dependent variable.

The right column (green lines, labeled Ratio of New to All Day Traders) presents the key results of a similar regression with the ratio of new to all day traders as the dependent variables, $ND_t/(ND_t+RD_t)$.