RFS Advance Access published February 25, 2011

Why Do (Some) Households Trade So Much?

Juhani T. Linnainmaa
University of Chicago, Booth School of Business

When agents can learn about their abilities as active investors, they rationally "trade to learn" even if they expect to lose from active investing. The model used to develop this insight draws conclusions that are consistent with empirical study of household trading behavior: Households' portfolios underperform passive investments; their trading intensity depends on past performance; and they begin by trading small sums of money. Using household data from Finland, the article estimates a structural model of learning and trading. The estimated model shows that investors trade to learn even if they are pessimistic about their abilities as traders. It also demonstrates that realized returns are significantly downward-biased measures of investors' true abilities. (*JEL* D10, G11)

While most households adjust their stock portfolios only infrequently, some trade very actively and underperform passive investments.[1] I ask whether a model in which investors rationally learn from experience can explain this behavior, and, if so, what the model tells us about households' beliefs.

In my model, investors are uncertain about their abilities and learn as they trade. If the value of observing another signal is high, then an investor trades even if she expects to lose money, thus apparently trading "too much." If a trade is successful, the investor infers skill and subsequently trades more. If an investor loses money, she will infer less skill and subsequently trade less. After enough losses, she stops trading altogether. Investors who are especially uncertain about their abilities trade small amounts early in their careers until

I would like to thank Tobias Adrian, Brad Barber, Robert Battalio, Tony Bernardo, Alan Bester, Michael Brennan, John Cochrane, Gene Fama, John Griffin, Mark Grinblatt, John Heaton, Matti Keloharju, Deborah Lucas, Reza Mahani, Kevin McCardle, Adair Morse, Ľuboš Pástor, Josh Pollet, Jeff Russell, Mark Seasholes, Nick Souleles, Luke Taylor, Pietro Veronesi, Yuzhao Zhang, and seminar participants at the 2008 AFA meetings, New York University, University of Chicago, Notre Dame University, Michigan State University, and Wharton for helpful comments. I thank Matti Keloharju for his help with the FCSD registry and Bobbie Goettler for editorial assistance. I am especially thankful for the detailed comments of an anonymous referee and the Editor, Alexander Ljungqvist. Financial support from the University of Chicago, Booth School of Business, is gratefully acknowledged. Send correspondence to Juhani T. Linnainmaa, University of Chicago, Booth School of Business, 5870 S Woodlawn Ave, Chicago, IL 60637; telephone: (773) 834-3176. E-mail: juhani.linnainmaa@ChicagoBooth.edu.

[1] Barber and Odean (2001), for example, report a mean (median) monthly turnover of 6.4% (2.9%) for men trading through a large U.S. discount brokerage firm. These estimates correspond to an investor turning the portfolio around every 1.3 and 4.8 years, respectively, indicating significant heterogeneity in trading activity. Barber and Odean also find that households' risk-adjusted performance decreases in turnover. Choi, Laibson, and Metrick (2002) report that although the introduction of a Web channel increased the amount of trading in 401(k) accounts, less than a quarter of all sample households trade even once. In the Finnish dataset I study in this article, 47.9% of all 1.1 million individuals with any stockholdings never trade during the eight-year sample period. On the other hand, three fifths of all trading activity originates from just 5% of the most active investors.

© The Author 2011. Published by Oxford University Press on behalf of The Society for Financial Studies. All rights reserved. For Permissions, please e-mail: journals.permissions@oup.com.
doi:10.1093/rfs/hhr009

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

they infer skill or leave. Finally, new signals influence subsequent behavior more early in investors' careers, when the priors are diffuse.

I use trading records of Finnish active ("high-frequency") traders to test the predictions of the learning model. These traders increase their trade sizes after successful trades and decrease trade sizes or quit after unsuccessful trades. Both the exit and trade size effects are stronger early in investors' careers. Some traders initially execute very small trades that seem to be motivated by their desire to learn more about their own trading skills.

I add belief distributions on top of a dynamic trading model to create a structural model of an investor population. I extract from this model the distribution of beliefs that comes closest to fitting the aggregate trading patterns in the data. I identify the structural parameters by matching, between the model and data, the number of investors who quit after each trading date and the sensitivity of investors' exit decisions to returns. Because investors can self-select and not trade at all, I also match the number of investors who trade at least once. Due to this entry moment condition, the model parameters represent belief distributions for the entire population, not just for the active traders.

The learning model matches the initial selection step, the conditional exit rates, and the performance-exit sensitivity in the data. The fact that this model can explain self-selection indicates that active traders' preferences do not need to be unlike the preferences of those who choose not to trade. Those who become traders just happen to reside in particular corners of the belief parameter space. The parameter estimates suggest that many investors start active trading without believing they are skilled. They start because they might be. More than a quarter of all traders begin with a belief that they have less than a 1-in-26 chance of being skilled. Approximately 1.2% of all investors (including non-active traders) have genuine trading skills. This fraction is 28.1% among those who begin to trade. Although some may be able to predict short-term price movements, active traders' order choices indicate that many try to profit by supplying liquidity to other investors.

Similar to the trading behavior of speculative futures traders,[2] many households stop active trading after trading (and losing) just a few times. In the estimated model, these investors are very uncertain about their abilities, which means a small number of losses can push them to quit. Intuitively, the investors in the model know that profiting by short-term trading is difficult. Thus, unless an investor receives some positive signals early on, she infers that she belongs to the unskilled majority and quits.

Although the traders in the model have constant abilities, more experienced traders perform better due to a learning-by-survival mechanism. Because a skilled trader is more likely to survive, a larger fraction of the investors who remain in the market are skilled. By inserting the structural parameter estimates back into the model, I find that the fraction of skilled traders increases

[2] See, for example, Ross (1975), Hieronymus (1977), and Teweles and Jones (1999, chapter 11).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

from 28% at the time of the first short-term trade to 61% after ten trades. Thus, many traders learn quickly that they are unskilled and quit.

The structural model yields an estimate of the size of a reverse survivorship bias that affects investor performance measurement. Because investors are more likely to quit after a series of unsuccessful trades, their realized performance is biased downward relative to true, unobservable skill. This bias is best illustrated with a coin-flip example. Suppose I start flipping a coin to measure its bias, but stop the experiment after I get the first tails. For a fair coin, the expected proportion of heads in such an experiment is

$$\left(\frac{1}{2}\right)^1 \frac{0}{1} + \left(\frac{1}{2}\right)^2 \frac{1}{2} + \left(\frac{1}{2}\right)^3 \frac{2}{3} + \cdots = 0.307.$$

Because I stop the experiment after tails, I oversample tails relative to the truth. Similarly, if investors quit after failed trades, data oversample poor performance relative to true skill. Although active traders' average true skill, expressed as a probability of a successful trade, is 0.377 in the estimated structural model, their observed skill is just 0.32.

Although I use both dynamic programming and Bayesian updating to model investor behavior, my results should not be construed as suggesting that households solve complicated mathematical problems or process information flawlessly. What I suggest is that a stylized learning model approximates households' trading decisions remarkably well. Thus, whatever rules of thumb underlie these decisions, the resultant behavior is observationally close to Bayesian updating and optimal policies.[3] Analogously, the literature on honey bees[4] does not suggest that bees solve optimal departure time problems on the fly, but that they use heuristics that are approximate solutions to these problems. Moreover, even if one does not subscribe to the view that these active investors learn from their experiences in anywhere near a rational manner, my results are still of interest because of their implications on investor overconfidence. Most investors in the estimated model trade because they are very uncertain about their trading abilities, not because they are overconfident about their abilities. Thus, investor overconfidence is not a requirement for empirically explaining household trading behavior. An important message from my results is that observed trading behavior does not necessarily imply that households are behaving irrationally. Given that a learning-based model can generate similar trading patterns, care should be taken before jumping to such a conclusion.

My study is most closely related to that of Mahani and Bernhardt (2007), who incorporate investor learning into a general equilibrium model of the markets. They show that learning reduces bid-ask spreads and the price impact of

---

3 I examine the behavior of households during a period of time when they try to profit from high-frequency trading. These are the investors who appear to learn (correctly) from their experiences. The very same households could rely on naïve reinforcement learning heuristics in other tasks, such as in allocating savings (Choi, Laibson, Madrian, and Metrick 2009) or picking stocks (Barber, Odean, and Strahilevitz 2004).

4 See, for example, Gould and Gould (1994).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

liquidity shocks. Mahani and Bernhardt motivate their modeling assumptions by noting that they are consistent with many empirical facts about investor behavior, whereas a model built on overconfidence alone is not. The main difference between this article and Mahani and Bernhardt's is that whereas they explore the equilibrium implications of learning theoretically, I estimate a structural model to draw inferences about actual investors' beliefs and abilities.

Structural models have not been widely used in finance to estimate learning processes. Only two recent papers, by Sørensen (2007) and Taylor (2010), take this approach. Sørensen (2007) uses a structural matching model to examine why companies funded by more experienced venture capitalists are more likely to go public. Taylor (2010) estimates a model in which the board of directors learns about CEO skill.

The article is organized as follows. Section 1 provides the background for the study and discusses the learning-by-trading assumption. Section 2 presents a stylized trading model and examines its empirical implications. Section 3 describes the data I use to test the trading model's predictions. Sections 4 and 5 present the empirical findings. Section 6 concludes.

1. Background

1.1 Stylized Facts about Household Behavior

This article's model reconciles several stylized facts about household behavior. The trading records from (at least) Denmark, Finland, Norway, Taiwan, and the United States suggest that the average household trades "too much," loses to the market, and alters trading intensity in response to past returns. Some households, however, exhibit superior performance.

Excessive trading and underperformance of the average household. Studies such as those by Odean (1999), Barber and Odean (2000, 2001), Grinblatt and Keloharju (2000, 2009), and Barber et al. (2009) find that the average household trades excessively. These households might be better off by holding the market. Barber and Odean (2001), Grinblatt and Keloharju (2009), French (2008), and Kumar (2009) suggest that both overconfidence and the desire to gamble contribute to these results.

Performance heterogeneity. Studies by Barber et al. (2005), Coval, Hirshleifer, and Shumway (2005), Bauer, Cosemans, and Eichholtz (2007), Goetzmann and Kumar (2008), Ivković, Sialm, and Weisbenner (2008), Nicolosi, Peng, and Zhu (2009), Grinblatt, Keloharju, and Linnainmaa (2010), and Seru, Shumway, and Stoffman (2010) find that a small number of individual investors outperform the market or their peers. Harris and Schultz (1998) show that some individual investors who use Nasdaq's Small Order Execution System for day trading are better than others. In many of these studies, the best-performing investors continue to outperform the worst performers from one period to the next.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Past performance and trading activity. Barber et al. (2005) find a positive correlation between the performance and subsequent trading activity among Taiwanese day traders. Investors in the top performing group increase their trading activity, whereas those in the bottom group decrease it. Glaser and Weber (2009) and Nicolosi, Peng, and Zhu (2009) find a similar relation between past performance and trading activity among the customers of a German online broker and a U.S. discount broker, respectively. Seru, Shumway, and Stoffman (2010) find that "exits" play an important role in this finding: The worst-performing investors quit trading altogether.

1.2 Learning and Paper Trading

If a learning model is to explain investor behavior, investors must learn from their own actions and not just by paper trading. Instead of trading with real money, investors could collect historical data to back-test trading strategies, or they could use software to track the performance of hypothetical investments in real time. Although the first approach is infeasible because no high-frequency datasets are publicly available, Web-based paper trading platforms are nowadays common and often integrated to brokerage accounts. Here, I discuss factors that may drive the learning-by-trading mechanism. Although I cannot empirically distinguish between these explanations, the mechanism itself exists in the data: Investors' own past performance greatly influences their future actions.

Practical difficulties with paper trading may impose limits on how much investors can learn from it. First, trade execution quality is an important consideration with short holding periods. Learning about this by paper trading is difficult because, first, trade and quote databases are not publicly available and, second, these data rarely show the direction of the order flow. Moreover, the traders I study use limit orders extensively. Limit orders further complicate paper trading because "hypothetical limit-order executions, constructed either theoretically from first-passage times or empirically from transactions data, are very poor proxies for actual limit-order executions" (Lo, MacKinlay, and Zhang 2002, p. 31).

Investors may get more out of paper trading by expending more effort. This cost, however, creates a tradeoff. If the (opportunity) costs of paper trading are higher than what the investor expects to lose by executing few small real trades, the investor is better off by real trading. These two considerations, the limits of paper trading and the relative cost comparison, may be important reasons for why investors learn by trading. In Section 5.6, I revisit this issue by measuring the (implied) costs of paper trading.

2. A Simplified Trading Model

I use a simplified trading model to demonstrate how uncertainty influences investor behavior. Uncertainty is important because the investor must consider

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

how actions taken today shape the options available in the future. This model differs from other Bayesian portfolio choice models[5] in two ways. First, I add a friction to generate the learning-by-trading mechanism. Second, the investor in my model is uncertain about her own trading skill. Other models often introduce uncertainty about some objective parameter, such as the size of the equity premium.

2.1 Assumptions

I assume that an investor lives for T periods and maximizes power utility over terminal wealth,

$$\mathrm{E}[U(W_T)] = \mathrm{E}\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right], \tag{1}$$

where $\gamma \neq 1$ is the coefficient of relative risk aversion. The investor can trade at dates $t = 1, 2, \ldots, T-1$. The investor receives a positive or negative signal about a single stock each day. This signal can be about a short-term price movement or about an opportunity to supply liquidity as a pseudo market maker. The investor then chooses an amount $\theta_t \geq \overline{\theta}$ to invest in this signal, where $\overline{\theta}$ is the minimum trade size. The signal is genuine with probability p, in which case the investment doubles in value. The signal is false with probability $1 - p$, in which case the investor loses the investment. A riskless asset also exists that pays no interest.

The investor has a beta-distributed prior belief about p. If the investor trades, she observes the outcome and updates her beliefs using Bayes' rule. Starting from a prior distribution Beta(α_1, β_1), the posterior is Beta$(\alpha_1 + 1, \beta_1)$ after a success and Beta$(\alpha_1, \beta_1 + 1)$ after a failure.[6] The investor does not observe the outcome if she does not trade. The minimum trade size $\overline{\theta}$ makes learning costly. If an investor believes $p < \frac{1}{2}$, then she weighs the expected trading loss against the value of a new signal. The appendix describes the solution to this model.

2.2 Empirical Implications

Three considerations shape investor behavior in this model. The first factor is the myopic demand: Given everything else, an investor with a higher prior belief about p is more likely to trade and, conditional on trading, trades more. An empirical implication is that investors' optimal trade sizes change over time as investors revise their beliefs about p. This effect's strength depends on how dispersed the prior distribution is: When priors are informative, new observations matter little and lead to only small changes in trade sizes. Because investors resolve uncertainty about p over time, later signals influence

5 See, for example, Detemple (1986), Brennan (1998), and Xia (2001). Pástor and Veronesi (2009) survey this literature.

6 See, for example, DeGroot (1970, p. 160).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

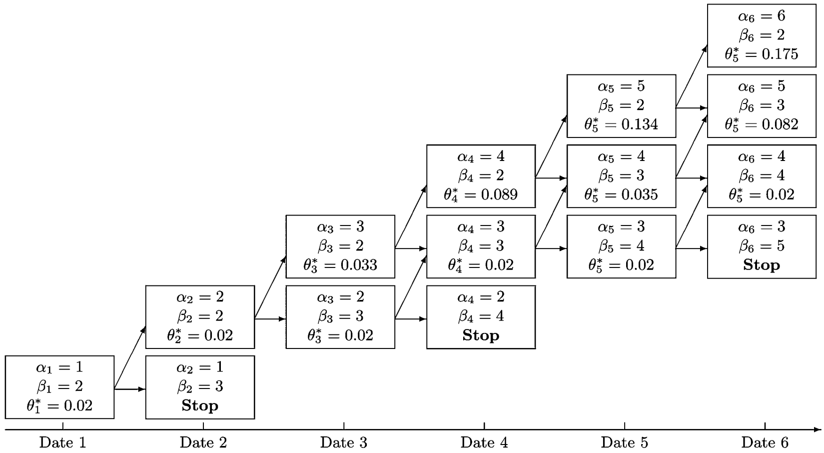

Figure 1
An example of optimal trading behavior in a simplified trading model
An investor has a relative risk-aversion coefficient $\gamma = 2$ and an investment horizon $T = 16$. The investor can trade $\theta_t \geq \overline{\theta}$ at dates $1 \leq t \leq T-1$. The minimum trade size is $\overline{\theta} = 0.02$. A trade is successful with probability p and returns 100%. It is unsuccessful with probability $1 - p$ and returns -100%. The investor's prior distribution is Beta(α_1, β_1) with $\alpha_1 = 1$ and $\beta_1 = 2$; that is, the prior mean belief about p is $\frac{\alpha_1}{\alpha_1+\beta_1} = \frac{1}{3}$. The investor updates her beliefs using Bayes' rule. This figure shows how the beliefs and optimal trade sizes, θ_t^*, change up to the sixth trade.

investor behavior less than early signals. This deceleration-of-learning effect is a common feature of learning models.[7]

The second factor is the intertemporal hedging demand. Although the true ability p is fixed, an investor's subjective investment opportunity set improves after successes and worsens after failures. This positive correlation between outcomes and investment opportunity set changes generates a negative (positive) demand component for investors who are more (less) risk averse than log-utility investors. The third factor is the option value of trading. A trade has value beyond its expected payoff because an investor observes the outcome only if she trades. An investor trades, even if she expects to lose money, if the value of observing another signal is high. Such an investor trades the smallest permissible amount, $\overline{\theta}$, to minimize the expected loss. The empirical implications are that, first, investors with diffuse priors are more likely to trade than those with precise priors and, second, that investors who trade only to learn make small trades.

Figure 1 illustrates how uncertainty influences investor behavior. The investor's prior mean belief about p is $\hat{p}_1 = \frac{1}{3}$. Thus, the investor expects only one-third of her trade signals to be genuine (and profitable). Nevertheless, the investor trades the minimum amount to observe the outcome. If this trade fails, the investor revises her beliefs downward to $\hat{p}_2 = \frac{1}{4}$ and stops trading. If the

[7] See, for example, Harvey (1989) and Pástor and Veronesi (2003).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

trade succeeds, the posterior mean is $\hat{p}_2 = \frac{1}{2}$ and the investor makes another small trade. The investor trades more than $\overline{\theta}$ only after two successful trades. This example illustrates how the option-value-of-trading mechanism weakens as the variance of the posterior distribution decreases. For example, although an investor in the ($\alpha_4 = 2$, $\beta_4 = 4$) node has the same mean belief about p as in the beginning, the investor does not trade because new signals are not as valuable.

3. Data

3.1 Investor Trading Records and Active Traders

This section describes the data I use to test the previous section's predictions. The Finnish Central Securities Depository dataset I employ records the portfolios and trades from January 1, 1995, through November 29, 2002, of all individual investors in Finland. The electronic records I use are exact duplicates of the official certificates of ownership and are hence reliable.

I use a simple algorithm to identify investors who begin trading actively, possibly for a short period of time. I then follow these investors for as long as they remain active. I use intraday round-trip trades to identify the beginning of an investor's active trading career.[8] (Barber et al. 2005 call these investors "day traders.") After an investor makes the first round-trip trade, I include in the career the same investor's other short-term trades. I classify as short-term trades all other intraday round-trip trades, regardless of when they occur, as well as purchases made during a two-week period following any other short-term trade. Thus, the active trading career ends when the investor no longer makes any round-trip trades and also does not make any purchases for at least two weeks. Although the 22,529 investors identified as active traders represent only 4.1% of all households in the FCSD data, they account for 64% of all household trading volume.

Figure 2 examines this algorithm's performance in identifying individual investors' active trading careers. Panel A shows that although the average investor's trading activity already increases before the first round-trip trade (that marks the beginning of the career), the level of activity jumps to a significantly higher level immediately after this day. Panel B shows that the date of the last short-term trade is similarly associated with a significant and long-lasting drop in trading activity. The run-up in trading activity toward the last trade arises because, first, investors by definition trade on the last day of the active trading career and, second, because investors' trades are temporally clustered. Thus, the trading activity run-up is an artifact of how the endpoint is defined. It does not imply that investors increase their trading activity just before they stop trading.

[8] Because the data are from the official registry of stockholdings and trades, these are genuine round-trip trades and not, for example, canceled trades. The registry uses 38 transaction-type categories, including a category for cancellations, to classify movements of securities between accounts. No atypical event would appear as a round-trip trade in these data.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

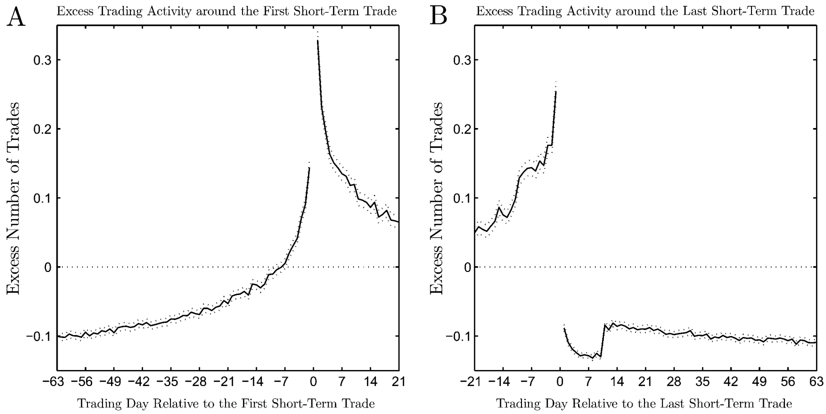

Figure 2
Excess trading activity before and after individual investors' active trading careers
This figure plots the excess trading activity for individual investors identified as active traders. The excess number of trades is the actual number of trades per day minus the average number of daily trades by the same investor during the combined periods of Panels A and B. Panel A starts three months before the first short-term trade and ends one month after it. Panel B starts one month before the last short-term trade and ends three months after it. The dashed lines show the 99% confidence interval. I omit the day of the first (Panel A) and last (Panel B) short-term trade. The data are the complete trading records of all active traders in Finland from January 1995 through November 2002.

An investor who stops "active trading" may continue stock picking after the end of the active trading career but at a frequency that is, by definition, lower than what it is during the career. Moreover, if investors can learn about their short-term trading abilities, they should also be able to learn about their stock-picking abilities, albeit at a possibly slower pace. For the reasons I discuss next, I focus on investors' active trading careers.

Active traders are well suited for testing the learning model because their trades form a clean sequence of choices and outcomes. By contrast, the lack of natural sequencing would complicate an analysis of buy-and-hold investors. Saying what these investors learn about their abilities and when is difficult. Also, risk adjustment is not an issue with short-term trades because intraday volatility swamps the daily risk premium, which is less than two basis points even if the annualized risk premium is 5%. Thus, active traders can profit only by predicting price movements, or by supplying liquidity, and not by capturing the risk premium.

Active traders can also learn about their abilities much faster than buy-and-hold investors. The difficulty with buy-and-hold investors is that we do not know their exposures to the (possibly unknown) risk factors. When both the exposures and factors have to be estimated from the data, disentangling luck from skill is difficult. Harris (2003) advises that "[i]n practice, more than 20 years of returns data are typically required to obtain useful results for a given investment manager." In contrast, learning about short-term trading skill is akin to measuring the bias of a coin. Each trade is just like one flip of a coin. The

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

precision at which we can estimate a person's ability increases in the number of short-term trades. Harris and Schultz (1998), for example, draw inferences about traders' abilities from just five days of trading data.

3.2 Variable Definitions

The gross profit of investor i's short-term trade on day t is

$$\text{gross profit}_{i,t} = \begin{cases} \sum_{s=1}^{n} (p_{i,s,t}^{s} - p_{i,s,t}^{b}) v_{i,s,t}^{b} & \text{if } v_{i,s,t}^{b} = v_{i,s,t}^{s}, \\ \sum_{s=1}^{n} \{(p_{i,s,t}^{s} - p_{i,s,t}^{b}) v_{i,s,t}^{s} + (p_{s,t}^{c} - p_{i,s,t}^{b})(v_{i,s,t}^{b} - v_{i,s,t}^{s})\} & \text{if } v_{i,s,t}^{b} > v_{i,s,t}^{s}, \\ \sum_{s=1}^{n} \{(p_{i,s,t}^{s} - p_{i,s,t}^{b}) v_{i,s,t}^{b} + (p_{i,s,t}^{s} - p_{s,t}^{c})(v_{i,s,t}^{s} - v_{i,s,t}^{b})\} & \text{if } v_{i,s,t}^{b} < v_{i,s,t}^{s}, \end{cases} \tag{2}$$

where n is the number of stocks investor i trades on day t, $p_{i,s,t}^{b}$ and $p_{i,s,t}^{s}$ are the investor's average purchase and sale prices in stock s, $v_{i,s,t}^{b}$ and $v_{i,s,t}^{s}$ are the number of shares purchased and sold, and $p_{s,t}^{c}$ is the stock's closing price on the same day. If an investor buys and sells different amounts, the remaining position, $x_{i,s,t}^{s} - x_{i,s,t}^{b}$, is marked to market at the same-day closing price.

I subtract commissions from gross profits to compute investors' net profits. Following Grinblatt and Keloharju (2009), I use the lowest commission rate available to new customers, 8.42 euro plus 0.15% of trade value. The investor's day t return on short-term trading is then $r_{i,t} = \frac{\text{net profit}_{i,t}}{x_{i,t}}$, where $x_{i,t}$ is the total size of the trade. The trade size is the maximum of the number of shares bought and sold multiplied by the volume-weighted average transaction price across all trades,

$$x_{i,t} = \sum_{s=1}^{n} \left(\max(v_{i,s,t}^{b}, v_{i,s,t}^{s}) \frac{p_{i,s,t}^{b} v_{i,s,t}^{b} + p_{i,s,t}^{s} v_{i,s,t}^{s}}{v_{i,s,t}^{b} + v_{i,s,t}^{s}} \right). \tag{3}$$

3.3 Summary Statistics on Active Traders

Table 1 shows that the average active trader (Panel A) is younger and more often male than the average non-active trader (Panel B). Active traders also trade more frequently than other investors even after excluding their short-term trades. Whereas the median active trader trades 56 times, the median non-active trader completes just two trades. The median number of (identified) short-term trades is just four, but this number varies considerably across investors. Whereas the bottom quarter completes only two short-term trades, the top quarter completes at least 15 short-term trades and the top 5% completes at least 112 short-term trades. The large number of (nearly) one-time traders bears similarities to the results on speculative futures traders. Hieronymus (1977), for example, documents that 37% of all futures trading accounts were traded only a few times.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 1
Trading Activity and Demographics of Active Traders

Panel A: Active Traders, $N = 22,529$

Variable	Mean	SD	Percentiles 25	50	75
Age	39.5	13.8	29.0	38.0	49.0
Male	81.3%				
Short-term Trades					
Number of Trades	24.8	78.3	2.0	4.0	15.0
Average Trade Size, EUR	13,231.0	49,228.2	2,074.9	4,493.1	10,553.8
Average Profit, EUR	−0.6	961.6	−75.0	−20.6	31.1
Fraction of Purchases Captured	0.823	0.258	0.667	1.000	1.000
Other Trades					
Number of Trades	88.5	119.3	23.0	52.0	110.0
Average Trade Size, EUR	9,668.2	59,138.9	2,285.3	4,318.2	8,411.7

Panel B: Other Individual Investors, $N = 527,436$

Variable	Mean	SD	Percentiles 25	50	75
Age	46.2	18.6	33.0	47.0	59.0
Male	59.2%				
Other Trades					
Number of Trades	6.4	14.4	1.0	2.0	6.0
Average Trade Size, EUR	6,045.0	238,620.8	1,161.6	2,310.0	4,848.1

This table compares active traders' demographics and trading activity to other individual investors. An active trader is an individual investor who completes at least one intraday round-trip trade during the sample period. The data are the complete trading records of all Finnish individuals from January 1995 through November 2002. I exclude investors with round-trip trades during the first three months of the sample. This table shows the number of trades and trade sizes for short-term trades (i.e., round-trip trades plus other speculative trades) and other trades. Fraction-of-purchases-captured is the number of days with short-term trades divided by the total number of stock-purchase days during an investor's active career.

The rules I use to identify "short-term trades" may miss out on some trades that should be included in the analysis. Nevertheless, I classify as short-term trades most trades investors make during their active trading careers. The fraction-of-purchases-captured row in Table 1 indicates that over 80% of the average trader's active-phase trades are identified as short-term trades. All trades are identified as short-term trades for 55.5% of active traders. These statistics suggest that I successfully identify a (possibly short) period of time when an investor trades frequently. Even if these rules exclude some genuine short-term trades that take place outside the active phase, such omissions would only add noise to the resultant sample and mask the possible effects of learning.

The median active trader loses an average of 21 euro, and the opportunity cost of capital, per a short-term trade. However, the profit distribution is skewed to the right. Because substantial short-term profits are more frequent than significant losses, the average investor's average per-trade loss is just one euro. I note that these profit estimates may be upward biased if some investors in the data exhibit a severe disposition effect.[9] A disposition-effect investor could sell

9 See, for example, Shefrin and Statman (1985), Odean (1998), and Barberis and Xiong (2009).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

shares purchased earlier in the day if the price increased significantly. Although I cannot clean the data of such disposition-effect-induced trades, their presence would not affect my analyses. Even if the average return is biased upward, the covariances between returns, exit decisions, and trade size changes, which I use to test for the learning-based explanation for trading, are unaffected.

Limit orders play a prominent role in active traders' execution strategies. Whereas 47.2% of other individual investors' trades originate from limit orders, this proportion is 52.5% for active traders. This difference suggests that many of them may try to profit by supplying liquidity to other (impatient) investors and not by predicting short-term price movements. When these traders are assigned into number-of-short-term-trades deciles, the proportion of limit-order trades increases monotonically from 50% (decile 1) to 56% (decile 10). Thus those active traders who remain in the market the longest use more limit orders. If longevity correlates with success, these numbers indicate that liquidity-providing individuals exhibit better performance than those who try to use market orders to profit from short-term price movements.

4. Testing the Implications of the Learning Model

4.1 Performance Heterogeneity

If learning is to explain investor behavior, some investors must believe they could be skilled. If investors have reasonable prior beliefs, it follows that some investors actually must be skilled. Otherwise, everyone would also know that they must be unskilled and would not trade.

Table 2, Panel A, reports performance persistence regressions that use data on active traders' short-term trades.[10] The first specification regresses the return on the tth trade against the average career return over all earlier trades. The second specification regresses the tth-trade success dummy on the career success rate. If no performance differences exist, or if such differences are transitory, the slopes in these regressions will be zero. I include, as additional regressors, monthly fixed effects as well as dummies for the number of short-term trades an investor has executed.

Both specifications indicate that some investors systematically outperform others. The coefficients for lagged performance are positive and both statistically and economically significant. The first specification indicates that an investor whose average career return is 1% expects to earn a return that is 0.151% higher than the return an investor with a career return of zero earns. The results are similar in the second specification. An investor with a 100% success rate has a 0.271 higher probability of yet another success in comparison with an investor with no past successes.

[10] Table 2, Panel B, as well as other analyses in this section, replicates the analysis using data simulated from a structural learning model. I discuss both the model and the results in Section 5.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 2
Heterogeneity and Persistence in Active Traders' Performance

Panel A: Investor Trading Records

	Full Sample		Early Career		Late Career	
Explanatory Variable	(1)	(2)	(1)	(2)	(1)	(2)
Career Average Return	0.151		0.095		0.316	
	(0.010)		(0.008)		(0.030)	
Career Success Rate		0.271		0.130		0.531
		(0.008)		(0.005)		(0.014)
Adjusted R^2	0.011	0.017	0.013	0.016	0.013	0.023
SD(Explanatory Variable)	0.017	0.184	0.032	0.322	0.010	0.124
N	472,904		99,135		373,769	

Panel B: Structural Learning Model

	Full Sample		Early Career		Late Career	
	(1)	(2)	(1)	(2)	(1)	(2)
Career Average Return	0.072		0.037		0.164	
	(0.004)		(0.005)		(0.009)	
Career Success Rate		0.048		0.026		0.107
		(0.004)		(0.005)		(0.008)
Adjusted R^2	0.002	0.001	0.001	0.001	0.001	0.001
SD(Explanatory Variable)	0.026	0.164	0.049	0.305	0.016	0.095
N	530,361		112,423		417,938	

Panel A reports regressions that measure persistence in active traders' performance. Each observation is a single short-term trade. The data in Panel A are the complete trading records of all active traders in Finland from January 1995 through November 2002. I exclude investors with round-trip trades during the first three months of the sample. Panel B uses data simulated from Section 5's structural model. Specification (1) regresses the return in trade t on the investor's average return from all previous trades; specification (2) regresses the success dummy in trade t on the investor's average success rate over all previous trades. Panel A regressions include both monthly fixed effects and dummy variables for investors trading for the first, second, ..., 50th time ($t > 50$ trades are the omitted category). Panel B includes only tth-trade fixed effects. The regressions are estimated using the full sample, investors' early ($t \leq 10$) trades, and investors' late ($t > 10$) trades. Standard errors, clustered by investor, appear in parentheses.

The last two columns in Table 2 split the sample into investors' early ($t \leq 10$) and late ($t > 10$) trades. The slopes are positive and statistically significant in both samples. In the first specification, a one-standard-deviation shock to the average career return increases today's return by 0.30% in the early-trades sample and by 0.32% in the late-trades sample. In the success-dummy specification, these shocks increase success probabilities by 0.042 and 0.066, respectively.[11]

4.2 Exits after Unsuccessful Trades

I estimate a Cox proportional hazard rate model with time-varying coefficients to model the probability that an investor quits after an unsuccessful trade. This model identifies how different outcomes influence the hazard rate without having to specify and estimate the baseline hazard.

[11] The estimated coefficients are larger in the late-career sample because average career return (and success rate) estimates become more precise as the number of per-investor observations increases. The table reports the standard deviation of the explanatory variable for each specification and sample.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 3
Cox Proportional Hazard Model of Active Traders' Exit Decisions

	Data			Model		
Explanatory Variable	(1)	(2)	(3)	(1)	(2)	(3)
Trade t Return	−1.836	−1.827		−4.304	−4.386	
	(0.245)	(0.245)		(0.092)	(0.092)	
Market Return on Trade t Day		−1.712			5.067	
		(0.414)			(0.757)	
Trade t Success Dummy			−0.488			−0.567
			(0.015)			(0.016)
Career Average Return	−1.572	−1.593		−5.217	−5.214	
	(0.275)	(0.275)		(0.188)	(0.188)	
Career Success Rate			−0.360			−0.671
			(0.027)			(0.032)
$(\ln(\text{Wealth}_t) - \ln(\text{Wealth}_0))/100$	−0.080	−0.080	−0.083	−0.018	−0.017	−0.119
	(0.009)	(0.009)	(0.009)	(0.036)	(0.036)	(0.036)
Pseudo R^2	0.014	0.014	0.016	0.009	0.009	0.006
N		437,111			552,890	

This table reports Cox regressions that measure how sensitive active traders' exit decisions are to performance. The data are the complete trading records of all active traders in Finland from January 1995 through November 2002. The dependent variable is the number of trades an investor makes before quitting. The time-varying covariates are the investor's performance in trade t, the HEX market return on the day of the trade, the average career performance, and the logged change in the investor's wealth from the first trade to one day before the current trade, divided by 100. Columns "Data" use investor trading records; columns "Model" use data simulated from Section 5's structural model. Standard errors, clustered by investor, appear in parentheses.

The time-varying coefficients are the investor's performance in trade t, the market return on the day of the trade, and the average career performance. Because investors who lose enough capital may have to quit even in the absence of learning, I include the change in the investor's portfolio value between the first trade and trade $t-1$ as an additional control. I also include monthly fixed effects based on each investor's start date. The model I estimate is

$$h(t|\mathbf{x}) = h_0(t)\exp(a + b_1\text{Date } t \text{ Return} + b_2\text{Market Return} \\ + b_3 \text{ Career Return} + \text{Controls}), \quad (4)$$

where $h_0(t)$ is the unspecified baseline hazard. I exclude all observations that take place during the last three months of the sample. This buffer ensures that investors who take a break from active trading are unlikely to appear as having stopped trading. Given that the average (median) waiting time between an investor's two short-term trades is 11.6 (2) days, this three-month cutoff should eliminate most false exits. Moreover, if any such misclassifications remain, they add noise to the performance-exit relation.

The estimates in Table 3 indicate that both the current and average career performance are significant determinants of the exit decision. In the success-dummy specification, a successful trade lowers the hazard rate by 38.6% relative to the unsuccessful-trade benchmark. The market-return coefficient is significantly negative, indicating that investors are less likely to quit after a market up day. This result suggests that investors may be unable to disentangle luck from skill perfectly: They seem to attribute positive market return partly

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

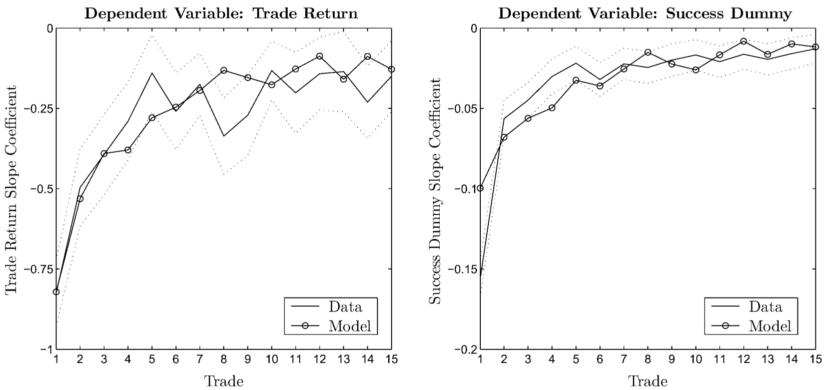

Figure 3
Cross-sectional regressions of active traders' exit decisions against performance
This figure plots the slope coefficients from a regression of active traders' exit decisions against the return or success dummy in trade t. A separate regression is estimated for investors trading for the first time, the second time, and so forth. The regressions include monthly fixed effects. The dashed lines indicate the 95% confidence intervals. The data are the complete trading records of all active traders in Finland from January 1995 through November 2002. The "Model" line uses data simulated from Section 5's structural model.

to their own skill, even after controlling for the trading return they achieve.[12] Although the change in wealth is statistically significant in Table 3, its effect on survival is economically small. Even if an investor's portfolio doubles in value, the hazard rate decreases by only 8%. This result should not, however, be interpreted as proving that exits are driven only by changes in beliefs and not by changes in wealth. Since changes in wealth are highly correlated with changes in beliefs about skill, these regressions cannot disentangle wealth and learning effects.

Figure 3 examines how return sensitivity changes over time. This figure plots slope coefficient estimates from cross-sectional linear regressions. The first regression uses data on investors trading for the first time; the second uses data on investors' second trades; and so forth. The dependent variable takes the value of one if the investor exits after the current trade, and zero otherwise. The independent variable in these regressions is the return (left panel) or success

[12] The fact that we do not observe true signals complicates this interpretation. Suppose, for example, that observed trading return is $r_{i,t} = \text{ability}_i + r_t^{\text{mkt}} + \varepsilon_{i,t}$, so that the signal about ability is $r_{i,t} - r_t^{\text{mkt}}$. A regression of some measure of behavior $X_{i,t}$ on the trading return and market return, $X_{i,t} = b_0 + b_1 r_{i,t} + b_2 r_t^{\text{mkt}} + e_{i,t}$, should then yield $\hat{b}_1 \approx -\hat{b}_2$ if investors filter out market returns and $\hat{b}_1 \approx 0$ if they do not. (I assume that $\varepsilon_{i,t}$ is uncorrelated with r_t^{mkt}.) The finding $\hat{b}_1 \approx \hat{b}_2$ is puzzling because it is not close to either of these alternatives. It indicates that investors revise their beliefs upward when the market goes up, even after controlling for their own trading return. It appears to suggest that $\varepsilon_{i,t}$ is negatively correlated with market returns. Such a correlation could arise, for example, from the marking of the market assumption and asynchronous trading. If an investor leaves a position open overnight and the underlying stock does not trade during the last m minutes of the trading day, market return is informative about the trade's true value at the close. If the market goes up, the investor's trade performed better than what the recorded return $r_{i,t}$ suggests, and vice versa. Nevertheless, I interpret Table 3's negative market-return coefficient cautiously. In addition to possibly not filtering out market returns perfectly, investors appear to learn (rightfully or not) something from market returns themselves.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 4
Pooled Regressions of Active Traders' Trade Size Changes on Performance

	Data			Model		
Explanatory Variables	Full Sample	Early Career	Late Career	Full Sample	Early Career	Late Career
Trade $t-1$ Success Dummy	0.083	0.102	0.077	0.041	0.045	0.039
	(0.004)	(0.006)	(0.004)	(0.002)	(0.003)	(0.002)
Trade $t-1$ Log(Trade Size)	0.652	0.612	0.661	0.976	0.953	0.980
	(0.006)	(0.004)	(0.007)	(0.000)	(0.001)	(0.001)
Adjusted R^2	0.501	0.463	0.500	0.957	0.889	0.966
SD(Trade t Success Dummy)	0.497	0.496	0.498	0.499	0.500	0.499

This table reports estimates from regressions that measure how sensitive active traders' trade size choices are to past performance. Each observation is a single short-term trade. The data are the complete trading records of all active traders in Finland from January 1995 through November 2002. I exclude investors with round-trip trades during the first three months of the sample. I regress the log-size of trade t against trade $t-1$'s log-size and success dummy. The regression also includes the change in the investor's portfolio value from one day before trade $t-1$ to one day before trade t. I exclude observations where an investor quits between trades $t-1$ and t. I include monthly fixed effects and dummies that represent investors trading for the first time, those trading for the second time, and so forth. Trades $t > 50$ are the omitted category. Regressions are estimated in the full sample, in a sample consisting of investors' early trades, and in a sample consisting of investors' late trades. Columns "Data" use investor trading records; columns "Model" use data simulated from Section 5's structural model. Standard errors, clustered by investor, appear in parentheses.

dummy in trade t. The regressions also include monthly fixed effects. If this regression included only a constant, the trade t regression would measure the probability that an investor who is still trading after trade $t-1$ remains alive after trade t. The slope estimates for the return variable measure how positive and negative outcomes shift the exit rate up or down.

The slope coefficient estimates increase rapidly at first before tending more slowly toward zero. This concave coefficient pattern, which is observed in both the return and success-dummy specifications, indicates that the speed of learning decreases over time. The changes are economically significant. For example, whereas the difference in exit probabilities between successful and unsuccessful traders is 0.156 in the first trade, it is only 0.053 in the second trade.

4.3 Trade Size Changes

If investors learn about their abilities, short-term trading becomes more attractive after successes and less attractive after failures. I estimate an AR(1) type model with trade t's log-size as the dependent variable and trade $t-1$'s log-size and success dummy as regressors. I exclude investors who exit after trade $t-1$. This restriction is important because Table 3 already establishes that investors are more likely to quit after unsuccessful trades. Here, I measure whether outcomes affect trade sizes even after ignoring exits.

The estimates in Table 4 indicate that past outcomes influence trade sizes significantly. The full sample estimates, for example, show that an investor with a successful trade increases the trade size by $e^{0.083}-1 = 8.6\%$ relative to an investor with an unsuccessful trade. The results also support the

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

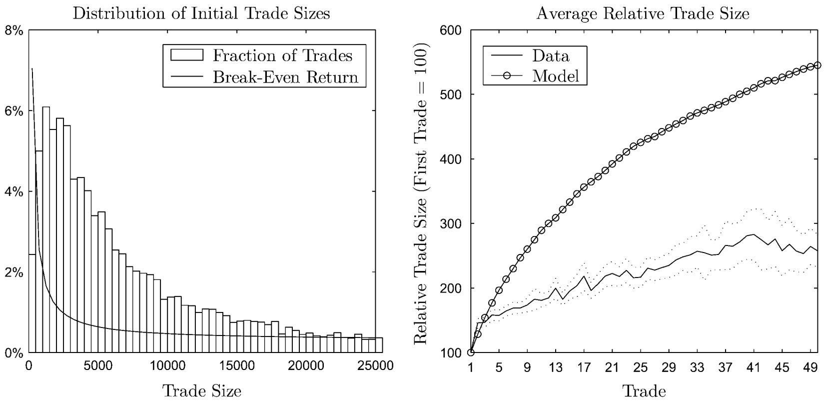

Figure 4
Distribution of initial trade sizes, break-even returns, and surviving traders' trade sizes
Panel A shows the distribution of trade sizes for active traders who are trading for the first time. Each category is 500 euro wide, and the x-axis labels give the midpoint of each category. The line shows the implied break-even return, which is the percentage difference between the purchase and sale price required for the investor to recover the 8.42 euro + 0.15% commission. Panel B shows the average trade size of an investor still alive at trade t. Trade sizes are standardized by the size of the first trade. The dashed lines show the 95% confidence interval. The "Model" line uses data simulated from Section 5's structural model.

deceleration-of-learning argument: More experienced investors pay less attention than inexperienced investors to new data.[13]

4.4 Exploratory Trades

In the model, an investor trades even if she expects to lose money when the value of observing another signal is high. Such an investor trades only the smallest possible amount. The left panel in Figure 4 shows the distribution of trade sizes for investors trading for the first time. It suggests that this mechanism may also be present in the data. Many of these initial trades are very small given their associated costs. To emphasize the significance of these costs, the figure overlays the implied break-even returns. An investor has to earn at least this high a return to recover the commissions. For example, an investor who buys shares for 1,000 euro must sell them at a 1.99% higher price to break even after the 8.42 euro + 0.15% commission.

The high break-even returns in Figure 4 are difficult to explain in a model without skill uncertainty. For example, 11.6% of the initial trades are so small (<1,406 euro) that the implied break-even return is at least 1.5%. If an investor expects to predict correctly a greater than 1.5% one-day price movement, why does she not trade more? These high break-even returns are less of a puzzle if some investors trade to learn about their abilities despite expecting to lose money.

[13] If market return is added to Table 4's regression, the success-dummy coefficient is 0.078 (t = 21.1) and the market-return coefficient is 1.307 (t = 12.8). This finding, which parallels that in Table 4, revoices the possibility that investors may not filter out market returns perfectly as they learn.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Consistent with the interpretation that some initial trades are exploratory, investors who continue to trade increase their trade sizes significantly. The right panel shows that the average investor who goes on to trade for the second time increases the trade size by 46%. The average investor who is still alive by the 25th trade makes a trade that is 2.2 times as large as her first trade. This increasing trade size pattern, however, holds only for successful investors: The average trade size change across all observations is negative, −1.5%. This estimate indicates that unsuccessful investors decrease their trade sizes until they ultimately quit.

Non-learning models could explain some of the trading patterns this section investigates. An investor could, for example, set aside a "gambling account" and trade until this account runs out of money. Such an investor's exit decisions would correlate with performance, much in the same way as they do in the data and in the learning model. Nevertheless, I believe it is difficult for an alternative story to simultaneously fit multiple patterns in the data. If an investor's beliefs are fixed, why would trade size changes correlate with past performance, why are so many initial trades exceptionally small, and why does the exit-performance sensitivity decrease over time? Although non-learning mechanisms may amplify some of the patterns in the data, I show in the next section that a learning model is simultaneously consistent with many empirical patterns.

5. Simulated Moments Parameter Estimation

5.1 A General Trading Model

I solve a general trading model that relaxes some of the assumptions of the binomial model. The next section adds belief distributions on top of this model to create a structural model of the investor population. I assume that an infinite-horizon investor with a subjective discount factor β maximizes power utility over consumption:

$$\mathrm{E}\left[\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma}\right]. \tag{5}$$

The investor can invest $\theta_t \geq \overline{\theta}$ in a risky trading opportunity that returns r_{t+1}. The remainder, $1-\theta_t$, earns a riskless rate of r_f. The wealth dynamics are

$$W_{t+1} = (W_t - c_t)\big(1 + r_f + \theta_t(r_{t+1} - r_f)\big). \tag{6}$$

Investor ability μ drives the mean of the trading opportunity return. The investor's date t belief about μ is normally distributed with mean m_t and variance υ_t. The investor observes signals $s_{t+1} \equiv \mu + \varepsilon_{t+1}^s$, where ε_{t+1}^s is normal with mean zero and variance σ_s^2. This signal translates into the trading opportunity return r_{t+1} as follows. First, the mean is bounded to rule out *a priori*

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

unreasonable beliefs. This bounded mean is $\mu^* = \min(\max(\mu, \underline{\mu}), \overline{\mu})$. Second, a market-wide shock (or an otherwise unpredictable return component), ε^m_{t+1}, is added to the return.[14] Third, the return is left-censored at $x = -1$ to give the trading opportunity limited liability. With these assumptions, the realized return is $r_{t+1} \equiv \max(\mu^* + \varepsilon^s_{t+1} + \varepsilon^m_{t+1}, -1)$.

If the investor does not trade, she does not observe the return and her beliefs remain unchanged. If the investor trades, she uses the signal s_{t+1} to update her beliefs. Her date $t+1$ beliefs are

$$m_{t+1} = (1 - w_t)\, m_t + w_t\, s_{t+1} \quad \text{and} \tag{7}$$

$$v_{t+1} = w_t\, \sigma_s^2, \tag{8}$$

where $w_t \equiv \frac{v_t}{v_t + \sigma_s^2}$. Because the investor maximizes isoelastic utility, the indirect utility function is separable in wealth:

$$V(W_t, m_t, v_t) = \frac{W_t^{1-\gamma}}{1-\gamma} B(m_t, v_t). \tag{9}$$

The Bellman equation is then

$$\begin{aligned} &V(W_t, m_t, v_t) \\ &= \max\left\{ \max_{c_t, \theta_t \geq \overline{\theta}} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \beta \mathrm{E}\left[\frac{((W_t - c_t) R_{\theta_t, t+1})^{1-\gamma}}{1-\gamma} B(m_{t+1}, v_{t+1}) \right] \right\}, \right. \\ &\qquad \left. \max_{c_t} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \beta \frac{((W_t - c_t)(1 + r_f))^{1-\gamma}}{1-\gamma} B(m_t, v_t) \right\} \right\}, \end{aligned} \tag{10}$$

where $R_{\theta_t, t+1} \equiv 1 + r_f + \theta_t (r_{t+1} - r_f)$. If the investor trades, her beliefs evolve to (m_{t+1}, v_{t+1}); if not, her beliefs remain unchanged at (m_t, v_t). The outer maximization problem in Equation (10) compares the utility from trading with that from quitting.

Equation (10) can be simplified by replacing c_t-values with the optimal choices from the first-order conditions. Function $B(m_t, v_t)$, which determines the optimal behavior, becomes[15]

$$\begin{aligned} &B(m_t, v_t) \\ &= \min\left\{ \left[1 + \left(\beta \min_{\theta_t \geq \overline{\theta}} \left\{ \mathrm{E}[R^{1-\gamma}_{\theta_t, t+1} B(m_{t+1}, v_{t+1})] \right\} \right)^{\frac{1}{\gamma}} \right]^{\gamma}, \right. \\ &\qquad \left. \left[1 + \left(\beta (1 + r_f)^{1-\gamma} B(m_t, v_t) \right)^{\frac{1}{\gamma}} \right]^{\gamma} \right\}. \end{aligned} \tag{11}$$

---

[14] An investor may lose money even if she makes the "right" trade: Even if a stock experiences a positive idiosyncratic shock, an investor with a long position loses money if the market falls enough. More generally, ε^m_{t+1} represents any return component unrelated to trading skill. The assumptions state that the investor correctly filters out the unpredictable return component before updating.

[15] I assume $\gamma > 1$. The solution for $\gamma < 1$ replaces the minimization problems with maximization problems.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

The second argument in the outer minimization problem is the payoff the investor receives if she does not trade. In this case, the investor is "stuck" with the same beliefs (m_t, υ_t) and thus always makes the same do-not-trade choice. The solution to this consumption-savings problem is

$$B^{\text{no trade}}(m_t, \upsilon_t) = \left[1 - (\beta(1+r_f)^{1-\gamma})^{\frac{1}{\gamma}}\right]^{-\gamma}. \tag{12}$$

Although this value is independent of the prior distribution, m_t and υ_t have to be such that the optimal choice is to refrain from trading. Equation (11) can then be written as

$$\begin{aligned} B(m_t, \upsilon_t) \\ = \min\Bigg\{ & \left[1 + \left(\beta \min_{\theta_t \geq \overline{\theta}} \left\{\mathrm{E}[R^{1-\gamma}_{\theta_t,t+1} B(m_{t+1}, \upsilon_{t+1})]\right\}\right)^{\frac{1}{\gamma}}\right]^{-\gamma}, \\ & \left[1 - (\beta(1+r_f)^{1-\gamma})^{\frac{1}{\gamma}}\right]^{-\gamma} \Bigg\}. \end{aligned} \tag{13}$$

Similar to Equation (12), $B(m_t, \upsilon_t)$ can also be solved in closed form when the investor has resolved all uncertainty about her abilities. In these $\upsilon_t = 0$ states, $B(m_t, 0)$ is

$$\begin{aligned} & B(m_t, 0) \\ & = \min\left\{\left[1 - \left(\beta \min_{\theta_t \geq \overline{\theta}} \left\{\mathrm{E}\left[R^{1-\gamma}_{\theta_t,t+1}\right]\right\}\right)^{\frac{1}{\gamma}}\right]^{-\gamma}, \left[1 - \left(\beta(1+r_f)^{1-\gamma}\right)^{\frac{1}{\gamma}}\right]^{-\gamma}\right\}. \end{aligned} \tag{14}$$

All values of $B(m_t, \upsilon_t)$ can be computed recursively on a mean-variance grid. Equation (14) gives the end-of-the-grid solution for $\upsilon_t = 0$. In every grid point (m^j, υ^k) with $\upsilon^k > 0$, $B(m^j, \upsilon^k)$ depends only on the period $k+1$ values $B(m^{j'}, \upsilon^{k+1})$. There is no feedback from k to $k+1$ because the investor always transitions to a lower variance state. Hence, if $B(m^{j'}, \upsilon^{k+1})$ is known for all $m^{j'}$, the expectation $\mathrm{E}\left[R^{1-\gamma}_{\theta_t,t+1} B(m_{t+1}, \upsilon_{t+1})\right]$ in Equation (13) is approximately

$$\mathrm{E}\left[R^{1-\gamma}_{\theta_t,t+1} B(m_{t+1}, \upsilon_{t+1})\right] \approx \sum_{j'=1}^{N} \Pr(m^{j'} \mid m^j, \upsilon^k) R^{1-\gamma}_{\theta_t, m^{j'} \mid m^j, \upsilon^k} B(m^{j'}, \upsilon^{k+1}), \tag{15}$$

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

where N is the number of grid points for m, $\Pr(m^{j'} \mid m^j, v^k)$ is the probability of transitioning from m^j to $m^{j'}$, and $R_{\theta_t, m^{j'} \mid m^j, v^k}$ is the implied return realization associated with this transition. Both the transition probabilities and implied returns can be computed from Equation (7). One of the benefits of this model is its computational efficiency. Because each grid point needs to be visited only once, the entire problem can be solved in a reasonable amount of time even for large grids. I use a grid with 1,000 points in both the mean and variance dimensions for a grid with a total of one million elements.

5.2 Structural Model of the Population

I add belief distributions on top of the trading model to create a structural model of the investor population. The first part of the model is the life-cycle model detailed above. The second part consists of distributions for investors' prior beliefs. The mean of each investor's prior distribution is drawn from a normal distribution with mean M_m and standard deviation S_m. The variance is drawn independently from an exponential distribution with mean M_v. I assume that investors' returns are (possibly) observed with noise. I assume that this noise is mean zero with variance σ_y^2. I estimate σ_y^2, which does not influence investor behavior, from the data along with the three belief distribution parameters.

This noise assumption allows for two mechanisms. First, it could be that we indeed observe a noisy version of the investor's actual signal. For example, the realized return could be 0.1% but, at the same time, the investor knows that she received abnormally poor execution. The investor corrects for "bad luck" by using a value higher than 0.1% as the signal. From an econometrician's viewpoint, returns are thus observed with noise. Similarly, when an investor keeps a position open over night, we incur a measurement error between the trade's true value (which the investor knows) and the marked-to-market value computed using the stock's closing price. This error arises because the closing trade takes place at either the bid or the ask and not at the true fundamental price. Also, if the stock does not trade during the last m minutes of the trading day, the trade's true value (which the investor knows) differs from that implied by the closing price. Second, σ_y^2 allows for the possibility that investors use additional information to revise their beliefs. For example, if the investor receives an additional signal s_t' that is not in the data, it will appear as if the original signal s_t is noisier than it actually is. When s_t' is high, the investor appears to revise her beliefs too much upward and vice versa. From the econometrician's perspective, signals will thus seem noisier than what they really are. I return to these interpretations in Section 5.4.

I set the other parameters of the model as follows. I assume that each model period has a length of two weeks and set each investor's risk aversion to $\gamma = 2$ and the subjective discount factor to $\beta = 0.997$ (0.925 per year).[16]

16 This β estimate is between the values used in, e.g., Krebs (2003) and Storesletten, Telmer, and Yaron (2007).

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

I set the total volatility of the trading-opportunity return, conditional on μ^*, to 2%.[17] I also set $r_f = 0.31\%$ (8% per year) and the minimum trade size to $\overline{\theta} = 2\%$. With this assumption, an investor with a wealth of \$25,000 needs to commit at least \$500 to the trading opportunity. This choice appears to be close to the actual practical minimum trade size observed in the data because only 1.9% of initial trades are this small.

I bound the mean trading opportunity return between $\underline{\mu} = -10\%$ and $\overline{\mu} = 2.5\%$. If investors try to profit by providing liquidity (and not by trading on private information), the upper bound on the expected return is a function of the bid-ask spread. I set asymmetric bounds because of transaction costs: If an investor makes a small round-trip trade and the stock price does not change, the investor's percentage loss can be substantial (see Figure 4). The reason for introducing these bounds is to demonstrate that unreasonable positive beliefs do not drive the estimates. No investor trades just because the mean could be, for example, 5% per trade. I discuss the unbounded-mean estimation results in Section 5.9.

I assume that 25% of return variance is due to the unpredictable return component, ε^m_{t+1}. I let this component follow a Bernoulli distribution, $\varepsilon^m_{t+1} = \pm\sigma_m$, with both outcomes equally likely. The lower bound for the size of the unpredictable component is the amount of systematic risk in individual stock returns. In a market model, this fraction is approximately 25%.[18] In addition to systematic risk, other trading-return components may also be unrelated to investor skill. For example, investors always face some execution risk (i.e., uncertainty about prices at which they can trade), and the idiosyncratic component of stock returns may be sensitive to some events the investor views as unpredictable. Any such factors increase the amount of variance that is unrelated to investor skill. I use the 25% assumption as my main specification and discuss higher estimates in Section 5.9.

5.3 Identifying Structural Parameters

Because the structural model does not yield closed-form estimation equations, I use the simulated method of moments (SMM) for indirect inference. I identify M_m, S_m, M_v, and σ^2_y from three types of moment conditions. The first moment is the initial entry decision that takes the value of one for investors who trade at least once and zero for those who never trade. The other moments come from the 15 exit-return regressions shown in the left panel of Figure 2. The first regression is estimated using data on investors trading for the first time;

[17] The volatility estimate in the data over all trades is 0.05. I set conditional volatility to just 0.02 because the unconditional return variance is the sum of the conditional-variance component and the variance of μ^*. In the data simulated from the estimated structural model, the unconditional volatility is 0.08.

[18] Suppose that the market's volatility is 30% and the average stock's volatility is 60%. The market model, $r_{i,t} = a_i + b_i r_{m,t} + e_{i,t}$, then implies that the fraction of systematic risk in the average stock's (with $b = 1$) return is $\left(\frac{0.3}{0.6}\right)^2 = 25\%$.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

the second uses data on investors' second trades; and so forth. The dependent variable in these regressions is an exit dummy variable, and the explanatory variables consist of an intercept, realized return, and monthly fixed effects. (The average fixed effect is constrained to zero to identify the intercept.) I take the intercepts and return slope estimates from these regressions for a total of 30 additional moment conditions. The intercepts trace out an exit curve that measures the fraction of investors who quit after the tth trade conditional on still being alive. The slope coefficients on returns give a return-sensitivity curve. It measures how investor sensitivity to realized returns changes over time.

To understand the identification, first considering how m_0 and v_0 influence investor behavior is useful. An increase in the prior mean m_0 increases the likelihood of the investor trading at least once. It also decreases both exit rates and return sensitivities because if m_0 is higher, then m_t also is higher for any given return path. An increase in the variance of the prior v_0 increases the option value of trading, also inducing more investors to trade. However, because high-prior variance investors learn faster than low-prior variance investors, both exit rates and return sensitivities increase in v_0.[19]

Structural parameters M_m, S_m, and M_v determine the composition of the population. The parameters are identified by finding a mixture that allows the model to match not just the entry moment, but also the shapes of the exit-rate and return-sensitivity curves. Figure 5 illustrates how shocks to the structural parameters perturb the exit-rate and return-sensitivity curves. In this figure, I first evaluate the moment conditions at the estimated parameter values (see Table 5 below). I then shock each structural parameter by half a standard error and plot proportional changes in moment conditions to show how exit-rate and return-sensitivity curves move and change shape.[20]

The effects of the four structural parameters are as follows. First, an increase in M_m has the same effect as that of increasing a single investor's m_0: It increases the number of investors who trade at least once; it lowers and flattens the exit rate curve; and it raises and flattens the return-sensitivity curve. (The solid line in the left panel of Figure 5 shows that all exit rates drop as M_m increases; the fact that this line is upward sloping shows that this drop is larger [relative to the starting point] for initial trades.) The return-sensitivity curve

[19] If $v_t > v_t'$, then $v_t - v_{t+1} > v_t' - v_{t+1}'$ by Equation (8). That is, the amount by which uncertainty decreases after a trade increases in the amount of initial uncertainty. The following example illustrates why exit rates and return sensitivities increase in v_0. Suppose we start from an investor (m_0, v_0) who does not trade and then increase the variance up to such point v_0' that she wants to trade. The expected posterior mean is the same as the prior mean, $\mathrm{E}(m_{t+1} \mid m_t) = m_t$, but the posterior variance is strictly lower. Hence, the investor quits if she realizes the expected outcome (or anything less). Investors who trade because v_0 is high are thus likely to quit and are very sensitive to (in particular initial) returns.

[20] Figure 5 complements economic intuition but does not replace it. Because the plots are based on shocks to the estimated parameter values, Figure 5 demonstrates local identification. These graphs also do not technically rule out the possibility that, for example, a combination of shocks to parameters 1 and 2 would coincide with the response to shocks to parameters 3 and 4.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

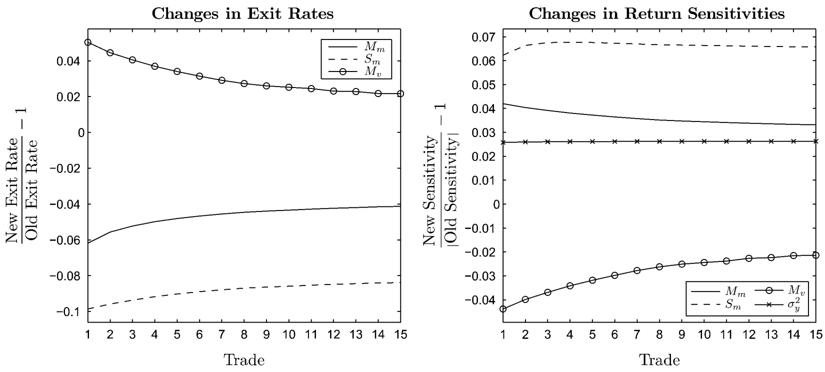

Figure 5
Identifying Structural Parameters
This figure shows how exit-rate and return-sensitivity moment conditions respond to half a standard error shocks to the values of the structural parameters M_m, S_m, M_v, and σ_y^2. These 2 * 15 moment conditions are estimated by regressing date t exit dummy variable against date t realized return, where $1 \leq t \leq 15$. The intercepts measure conditional exit rates, and the coefficients represent sensitivity of exit on returns. The baseline model is evaluated at the estimated parameter values (see Table 5). This figure plots the proportional changes in moment conditions between the shocked and baseline model, $M_k^{\text{shocked}} \Big/ \left|M_k^{\text{baseline}}\right| - 1$.

flattens because an increase in M_m decreases the number of investors who are almost indifferent between trading and not trading. Investors with high m_0s are less sensitive to new data. This increase affects early trades more because future m_ts equal m_0 plus any learning that takes place. Second, an increase in S_m gives investors more extreme m_0s: Investors with a low mean get even lower mean, and vice versa. An increase in S_m lowers and flattens the exit-rate curve and raises and adds curvature to the return-sensitivity curve. From the perspective of identification, an important difference between S_m and M_m is that M_m has more of an effect on the entry decision and that they affect, in particular, the shape of the exit-rate sensitivity curve very differently. An increase in M_m flattens this curve, whereas S_m steepens and adds curvature to it.[21]

The third parameter, M_v, increases both the average v_0 and the variance of v_0s in the population. The intuition about how v_0 influences any one investor's behavior applies: As M_v increases, the number of investors who trade at least once increases; the exit-rate curve rises and steepens; and the return-sensitivity curve falls and steepens, that is, investors become more sensitive to returns, in particular for low values of t. From the perspective of identification, one

---

[21] An increase in S_m does not influence the entry moment condition as much as M_m. A low-mean investor who receives an even lower mean (as S_m increases) still does not trade, and a high-mean investor who receives an even higher mean still trades. The response to an S_m shock is, however, positive when most investors start in the do-not-trade region. In this case, some investors are pushed from the do-not-trade region into the trade region. An M_m shock, on the other hand, unequivocally pushes all investors toward the trade region, resulting in a stronger effect. Thus, even if M_m and S_m influenced the shapes of the exit-rate and return-sensitivity curves similarly (which they do not, as Figure 5 shows), the total effects of an M_m shock could not be replicated by an S_m shock of any magnitude.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

key difference between M_υ and M_m and S_m is that a change in M_υ pushes entry- and exit-rate moments in the same direction, whereas changes in M_m and S_m push these moments in opposite directions. The last parameter, σ_y^2, affects only the return-sensitivity curve; it is unrelated, by construction, to the other moments. An increase in σ_y^2 adds noise to observed returns and thus leads to a parallel upward shift in the return-sensitivity curve. σ_y^2 is thus identified off the level of this curve.

I use SMM to estimate the four structural parameters. For an investor with prior beliefs (m_0, υ_0), I let $H(m_0, \upsilon_0)$ denote a 16×1 vector of the date 0 entry decision and conditional exit rates. The first element of this moment vector is either zero (do not trade) or one (trade at least once); the second element is the probability that an investor who trades at least once exits after the first trade; and so forth. The model implied kth moment is

$$M_k(\Theta) = \frac{\int_{-\infty}^{\infty}\int_0^{\infty} \varphi(m_0; M_m, S_m)\, f(\upsilon_0; M_\upsilon)\, \Pr(\text{alive}_k | m_0, \upsilon_0) H_k(m_0, \upsilon_0) dm_0 d\upsilon_0}{\int_{-\infty}^{\infty}\int_0^{\infty} \varphi(m_0; M_m, S_m)\, f(\upsilon_0; M_\upsilon)\, \Pr(\text{alive}_k | m_0, \upsilon_0) dm_0 d\upsilon_0}, \tag{16}$$

where $\Theta = [M_m \;\; S_m \;\; M_\upsilon \;\; \sigma_y^2]$, φ is the normal density, f is the exponential density, $\Pr(\text{alive}_k \mid m_0, \upsilon_0)$ is the probability that an investor with initial beliefs (m_0, υ_0) reaches the kth moment condition, and k indexes the moment conditions from 0 to 15. Due to the denominator in Equation (16), the exit rates are conditional: $M_k(\Theta)$ is trade k exit rate for investors who survive to make this trade. I obtain the expectations required for the exit-regression moment conditions from integrations similar to Equation (16) and augment vector M_k with these moments. All computations assume that investors' prior beliefs are well calibrated. For example, if an investor's posterior mean crosses the quit-trading threshold for returns $r_{t+1} \leq r_{\text{quit}}$, the investor's exit rate is $\Pr(r_{t+1} \leq r_{\text{quit}} \mid m_t, \upsilon_t)$.

The use of the mean-variance belief grid in the life-cycle model solution simplifies the estimation of the full structural model. I first record the following statistics for each grid point: the date 0 entry decision, the conditional exit rates for dates $1, \ldots, 15$, and the conditional expectations used in the exit regression. Parameters M_m, S_m, and M_υ determine the likelihood of drawing an investor with prior beliefs close to grid point $(m^{(j)}, \upsilon^{(k)})$. Equation (16) can thus be evaluated by weighing each grid point based on Θ. The SMM estimator $\hat{\Theta}$ is then

$$\hat{\Theta} = \arg\min_{\Theta} \left(\hat{M} - M(\Theta)\right)' W \left(\hat{M} - M(\Theta)\right), \tag{17}$$

where $\hat{M}$ is the vector of estimated moments from the data, $M(\Theta)$ is the vector of model-implied moments from Equation (16), and W is an arbitrary

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

positive-definite weighting matrix. I use the optimal weighting matrix for estimation.

5.4 Estimation Results

Table 5 shows the estimation results. The discrepancies between the model and data moments are both statistically and economically small. Most of the simulated moments are within one standard error of the data moment. Although some deviations are larger, such as the return sensitivity for the eighth trade, the overall fit of the model is good. The test of overidentifying restrictions does not reject the model. The J test statistic of 28.5 has a p-value of 0.39.

The M_m and S_m estimates suggest that most individuals are pessimistic about their trading abilities. The mean of each individual's prior distribution is drawn from a normal distribution with a mean of −4.73% and a standard deviation of 2.01%. This finding suggests that just 0.93% of all individuals have a positive prior mean about their abilities. However, because investors in the model always consider what *might* be, the upper-tail probability $\Pr(\mu > 0)$ is what largely determines who trades and who does not. The low mean in the population-wide distribution is the model's way of ensuring that this probability is so low for most individuals that they have no incentive to trade.

Panel C reports traders' abilities based on the estimated model. The first two rows report the probability that an investor is skilled and the probability that an investor receives a positive first-day return. (For non-active investors, the first-day-return computation applies to a hypothetical trade.) The average trader assigns a probability of 0.281 to being skilled and a probability 0.377 to a positive first-day return. These numbers do not coincide because returns reflect two sources of uncertainty: the variance in ability and the variance in returns. Even among the most skilled traders, the probability of a positive one-day return is still just 0.689 because of the randomness in returns; they expect to make money in fewer than seven out of 10 trades.

The last two rows in Panel C show the 95% confidence intervals investors have about μ (the unbounded mean). Here, I rank investors by the precision of their prior beliefs. The 5% of active traders with the most precise prior beliefs are also most optimistic: They believe μ to be between 0.18% and 1.29%. As the dispersion in beliefs increases, traders become more pessimistic about their abilities. For example, the 95% confidence interval for the 5% of traders with the most dispersed beliefs is $-4.54\% < \mu < 0.48\%$. These investors trade because of the possibility that the true μ may be in the far-right tail of their prior distribution.

Most active traders think they will lose money. At the same time, some investors stay out even though they may be skilled. If investors' beliefs are unbiased, the belief distributions coincide with skill distributions. Thus, the belief distributions in Panel C imply that 28.1% of those who try active trading are skilled. The remaining three quarters are unskilled, but they trade because they

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 5
Structural Learning Model

Panel A: Actual Moments versus Simulated Moments

	Conditional Exit Rates			Return Slope Coefficients		
	Data			Data		
	$\hat{b}$	SE	Model	$\hat{b}$	SE	Model
Entry (Date 0)	4.10%	0.14%	4.09%			
Trade 1	20.69%	0.67%	20.33%	−0.825	0.053	−0.816
Trade 2	13.28%	0.78%	13.13%	−0.496	0.061	−0.551
Trade 3	9.03%	0.76%	9.61%	−0.394	0.064	−0.413
Trade 4	7.15%	0.80%	7.52%	−0.289	0.061	−0.328
Trade 5	5.72%	0.71%	6.14%	−0.140	0.060	−0.270
Trade 6	4.99%	0.70%	5.18%	−0.260	0.061	−0.230
Trade 7	3.94%	0.67%	4.47%	−0.175	0.049	−0.199
Trade 8	4.29%	0.66%	3.95%	−0.336	0.060	−0.177
Trade 9	4.18%	0.69%	3.51%	−0.271	0.064	−0.158
Trade 10	2.96%	0.67%	3.18%	−0.132	0.046	−0.143
Trade 11	3.12%	0.71%	2.88%	−0.202	0.064	−0.130
Trade 12	2.34%	0.67%	2.63%	−0.143	0.057	−0.119
Trade 13	2.51%	0.67%	2.42%	−0.135	0.064	−0.110
Trade 14	2.19%	0.69%	2.23%	−0.230	0.057	−0.102
Trade 15	2.19%	0.66%	2.07%	−0.150	0.056	−0.094

Panel B: Structural Parameter Estimates

	Prior Distribution Parameters				
	M_m	S_m	M_v	σ_y^2	J test (p-value)
Estimate	−4.7295	2.0102	0.0034	0.0060	28.47
SE	(0.4736)	(0.2171)	(0.0004)	(0.0003)	(0.387)

Panel C: Trader and Non-Trader Abilities (%)

Population Subset	Percentiles of the Variance of the Prior Distribution					
	Mean	5%	25%	50%	75%	95%
All Investors						
Pr(Skilled)	1.200	0.000	0.000	0.000	0.000	0.788
$\Pr(r_1 > 0)$	5.116	0.003	0.133	1.066	5.264	25.049
$m_{0,p=0.975}$	−3.720	−4.579	−3.234	−5.406	−3.307	−2.651
$m_{0,p=0.025}$	−5.738	−5.096	−4.453	−7.292	−5.982	−6.579
Active Traders						
Pr(Skilled)	28.121	1.295	3.911	12.538	43.708	99.442
$\Pr(r_1 > 0)$	37.675	14.115	24.848	35.534	48.257	68.891
$m_{0,p=0.975}$	0.714	1.286	0.802	0.716	0.618	0.477
$m_{0,p=0.025}$	−2.275	0.184	−1.330	−2.203	−3.175	−4.540
Other Investors						
Pr(Skilled)	0.051	0.000	0.000	0.000	0.000	0.036
$\Pr(r_1 > 0)$	3.726	0.002	0.119	0.926	4.334	17.508
$m_{0,p=0.975}$	−3.909	−4.503	−4.388	−2.967	−2.607	−3.205
$m_{0,p=0.025}$	−5.886	−5.009	−5.589	−4.821	−5.236	−7.066

(*continued*)

know they *might* be skilled. By contrast, only 0.05% of those who stay out are skilled. Given the relative sizes of the trader and non-trader groups, this finding implies that 4.1% of all investors with true $\mu > 0$ never try active trading.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Table 5
Continued

Panel D: Cumulative Exit Rates for Skilled and Unskilled Traders

	Cumulative Exit Rates (%)		
	Skilled	Unskilled	% of Skilled Investors
Trade 0	0.00	0.00	28.12
Trade 1	2.45	27.33	34.43
Trade 2	5.65	45.27	40.28
Trade 3	8.79	56.31	44.96
Trade 4	11.68	63.58	48.69
Trade 5	14.29	68.68	51.70
Trade 6	16.65	72.43	54.19
Trade 7	18.78	75.31	56.27
Trade 8	20.74	77.60	58.06
Trade 9	22.52	79.46	59.61
Trade 10	24.18	81.01	60.97

A structural learning model is estimated using data on all active traders in Finland. An active trader is an individual who completes at least one intraday round-trip trade during the sample period. Panel A compares simulated moments to sample moments. Panel B reports on the structural parameter estimates. Parameters M_m, S_m, and M_v are multiplied by 100. Panel C reports the probabilities that traders assign to being skilled, to realizing a positive first-day return, and 95% confidence intervals that investors have about their skill. Panel D reports on skilled and unskilled traders' cumulative exit rates. The first part of the structural model is a trading model in which a CRRA investor learns about her abilities. The second part consists of belief distributions. The mean of each investor's prior distribution is drawn from a normal distribution Normal(M_m, S_m). The variance is drawn independently from an exponential distribution with mean M_v. Investors' returns are observed with noise; the variance of this mean zero noise component is σ_y^2. Each period in the model has a length of two weeks. The following parameters are fixed: The subjective discount factor is $\beta = 0.997$; the volatility of returns is 0.02; 25% of return variance is due to an unpredictable component; the risk-free rate is 0.31% per period; and the minimum trade size is $\overline{\theta} = 2\%$. The Simulated Method of Moments with the optimal weighting matrix is used to estimate the model. The 31 conditions are the initial entry decision, the conditional exit rates, and the return sensitivities for trades 1, ..., 15 (see text for details).

The total amount of (genuinely) skilled investors is 1.2% in the calibrated model. This figure is within the estimate by Coval, Hirshleifer, and Shumway (2005), who find persistent superior performance among the investors in the top decile. This estimate is also within the bounds of Grinblatt, Keloharju, and Linnainmaa (2010), who find superior performance among the 4% of individuals who receive the highest score on an IQ test. These estimates must, however, be compared with some caution. First, I probably understate the prevalence of skill because I consider only skill in active trading. I do not try to identify investors who achieve superior performance at a lower trading frequency. Second, the 1.2% estimate includes all investors who *could* profit by trading. The actual number of skilled traders who remain in the market is lower because some never trade and others exit after inferring, by bad luck, that they are unskilled. Finally, skill is far more prevalent among those who try active trading: 28% of first-time traders have at least a slightly positive μ.

The estimate of σ_y^2 is significantly positive in Table 5, Panel B. As discussed above, this measurement noise may reflect several sources. First, the returns investors "see" may not coincide with the returns computed from the data due to, for example, investors' ability to adjust returns for abnormally good or poor execution. Investors may also not perfectly filter out nonskill information, such

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

as market returns, from the returns they experience. The effect of such filtering error from the model's perspective is the same as increasing the measurement error in returns. Second, $\sigma_y^2 > 0$ may also indicate that investors use additional information, beyond their own return, to learn about their abilities. From the estimation viewpoint, it will appear as if the return signals are noisier than what they actually are. The level of σ_y^2 should thus be interpreted cautiously. For example, if investors use multiple signals, it will appear as if the trade signal is far noisier than what it actually is, even if these additional signals are relatively uninformative. I thus only interpret the significantly positive σ_y^2 estimate as suggesting that either investors' signals are observed with noise or that investors use some additional signals to learn about their abilities.

5.5 Learning by Survival

An interesting application of the structural model is the study of the learning-by-survival mechanism. Investors in the model do not learn how to trade. They learn about their fixed skill μ only as they trade. However, because low-ability investors are more likely to lose money by trading, and because investors quit after losses, surviving investors are of higher quality.

I study this selection process by inserting the structural parameter estimates back into the model. Table 5, Panel D, tabulates the cumulative exit rates for the skilled and unskilled investors and shows how the average ability of the trader pool changes over the first ten trades. The calibrated model shows that an unskilled trader quits after the first trade with probability 0.273. The cumulative probability of quitting after the second trade is 0.453. Skilled traders' cumulative exit rates are significantly lower, 0.025 and 0.057, respectively. This estimate indicates that every twentieth skilled investor is unlucky, exiting after the second trade despite being, unbeknownst to her, skilled. (One fourth has quit after the tenth trade.) These exit-rate differences increase the average skill of surviving traders. Whereas 28% of the investors trading for the first time are skilled, this fraction increases to 52% by the fifth trade and to 61% by the tenth trade. Due to this survival mechanism, trade performance correlates positively with trading experience despite the fact that investors' abilities are fixed.

Figure 6 plots the fraction of profitable trades for investors trading for the tth time. In the simulated data, the success rate increases quickly at the onset as many unskilled traders quit. This growth, however, slows down as the success rate begins to asymptote toward 65%. (In the estimated model, when only skilled traders remain, the average skilled trader still loses money in one third of the trades.) The dashed line shows that although many unskilled investors exit the market quickly, some survive for a long time. For example, even after 100 trades, 3.2% of surviving traders have at least slightly negative μ. These traders survive because when μ is just below zero, learning whether μ is positive or negative is difficult; that is, these investors require many observations to make such a call. At the same time, because μ is close to zero, the perceived cost of learning is negligible and the investor finds it optimal to trade. This

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

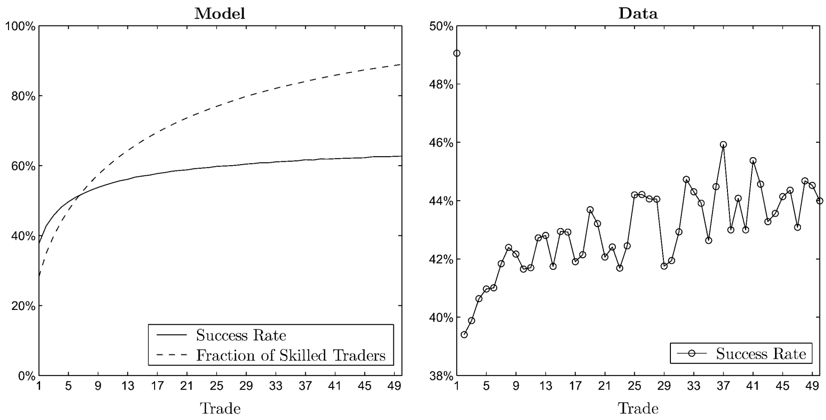

Figure 6
A comparison of success rates between the structural model and data
This figure plots the average success rates over active traders' first 50 trades. The left panel simulates data from the structural model at the estimated structural parameter values. The right panel computes success rates from the data. The dashed line in the left panel is the fraction of skilled investors, $\mu > 0$.

result suggests that the existence of long-lived traders with poor performance is not necessarily inconsistent with a learning model.

Except for the first trade, average success rates show the same pattern in the data: The sharp initial increase is followed by more modest improvements.[22] The success rate for the first trade is abnormally high, almost 50%, but this high estimate is an artifact of how I identify active trading careers. The first trade is always, by definition, a round-trip trade (which is associated with higher returns), whereas many of the subsequent trades are other speculative (nonround-trip) trades.

5.6 Implied Cost of Paper Trading

Because investors in the data make real trades instead of paper trades, actual trading must dominate paper trading in terms of expected utility. Consequently, a trader's expected utility would decrease if she were to paper trade. I use this observation to estimate the implied (lower bound) cost of paper trading. I simulate investors from the model at the estimated structural parameter values and examine investors who trade in spite of expecting to lose money. For each investor who trades with the expectation of losing money, I compute the certainty-equivalent cost: Instead of trading, how much would that

[22] The average success rate remains below 50% even by the 50th trade. In reality, the average surviving trader earns positive profits for two reasons. First, investors' trade sizes correlate positively with realized returns in the data: Large positive outcomes, in terms of dollar value, are more frequent than negative outcomes of the same magnitude. Second, Figure 6 maintains the assumption about a fixed 0.15% + 8.24 euro commission schedule. In reality, long-lived traders pay less in commissions (as little as max[0.06% $*$ trade size, 3 euro]) because the schedule is a function of past trading activity. Because there are no historical records on investor-specific commission schedules, I use the same schedule for all investors. An imputation of lower commissions for surviving traders would steepen the curve in the "data" panel, thus improving its match with the curve in the "model" panel.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

investor be willing to pay to observe the first outcome without trading? This certainty-equivalent estimate is almost the same as the expected loss except that it adjusts for the concavity of the utility function.

The average implied cost of paper trading is $\mathrm{CE}_1 = 0.031\%$ (of wealth) in the estimated model for investors who are about to trade for the first time. The average investor with $m_0 < 0$ would prefer paper trading if the cost of doing so is less than CE_1. However, negative-expectation investors continue to trade until they either believe they are skilled or become convinced they are unskilled. In expectation, the average investor in this group trades 5.6 times before reaching one of these outcomes. Accounting for the expected evolution in beliefs, I find an implied lump-sum certainty-equivalent cost of $\mathrm{CE}_T = 0.112\%$ for the average investor with $m_0 < 0$. That is, the average $m_0 < 0$ investor would be willing to pay at most CE_T to learn her "type" without trading. Both CE_1 and CE_T are economically quite small: Trading to learn does not appear to be very costly in certainty-equivalent terms.

5.7 Reverse Survivorship Bias

Another application of the model is to measure the size of a reverse survivorship bias that influences investor performance measurement. This bias arises because investors quit after unsuccessful trades. If an investor never quits, then the average skill measured from the data would be an unbiased estimate of investor skill. However, because investors quit after unsuccessful trades, the skill estimates from the data are biased downward.

This observation follows from optional stopping-time theorems.[23] If ε_t is mean zero and a martingale, and thus $\mathrm{E}\left[\sum_{t=1}^{T} \varepsilon_t\right] = 0$, the average ε_t is different from zero if ε_t is correlated with the (random) stopping time T:

$$\mathrm{E}\left[\frac{1}{T}\sum_{t=1}^{T} \varepsilon_t\right] = \mathrm{cov}\left(\frac{1}{T}, \sum_{t=1}^{T} \varepsilon_t\right). \tag{18}$$

If the survival probability increases in ε_t, the sample average of ε_t is negative (instead of zero). Analogously, the skill estimates are too low relative to investors' true skills because investors quit after low returns and so the data oversample poor performance.

An alternative explanation for this bias follows from considering what must have happened to an investor who stops trading. An investor stops trading only if her posterior mean falls below some threshold, say, m_{quit}. The key observation is that because each realization is a noisy signal about true μ, the posterior mean can drop from m_0 to m_{quit} only if the realized ability is strictly lower than m_{quit}. The reason is that the posterior mean is always between the prior mean and signal. Thus, at the time of the exit, the best estimate about an investor's

[23] See, for example, Williams (1991). Linnainmaa (2010) measures mutual fund managers' abilities and discusses the reverse survivorship bias in greater detail.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

ability is m_{quit} but her observed ability is $\hat{\mu} < m_{\text{quit}}$. The difference between the quitting boundary, which is endogenous in the model, and the realized performance represents the reverse survivorship bias.

The gap between the true and observed skill is economically significant. Although the average investor's actual success probability is 0.377 in the model, her realized success rate is just 0.32. Thus, this reverse survivorship bias is an important consideration when measuring traders' abilities. Equation (18) shows that the magnitude of this bias depends on how sensitive investors' exit decisions are to poor performance. Because active traders are highly sensitive to poor performance, the bias in their performance estimates is large.

5.8 Additional Measures of Model Fit

I simulate data from the structural model to examine how closely simulated and actual traders resemble each other. I generate a matched-size sample of 22,529 traders and use these data to repeat the tests I described in Section 4. To approximate the length of the Finnish Central Securities Depository dataset, each simulated trader generates at most 100 observations.

These additional tests suggest that the model is consistent with the data also in dimensions the moment conditions do not directly represent. First, Table 2, Panel B, shows that simulated traders' performances also persist. The slope coefficient is 0.07 in a regression of a trade t return against career average return; this estimate is 0.15 in the data. Similar to the data, both the model and data slope coefficients are higher in the late-career sample. Second, Table 3 and Figure 3 show that investor survival follows similar patterns in the model and data. The fit in Figure 3 is almost perfect because the model is estimated by matching, among other things, exit-return sensitivities.[24] The key difference is the role of market return (taken as ε^m_{t+1} in the model). In the model, this variable enters with a positive sign because the investor cares only about the signal s_{t+1}, but the observed return r_{t+1} is the sum of this component and ε^m_{t+1}. Hence, if ε^m_{t+1} is positive, s_{t+1} is less than r_{t+1}. The change in wealth has the same negative sign in both the model and data. Despite the fact that wealth change is inconsequential for model investors, it correlates with exit decision because it is a transformation of the career average return. Although a model investor's exit decision depends only on the posterior distribution, all four independent variables inform about this distribution through different functional forms.

Third, Table 4 shows that trade size choices in the model also are sensitive to successes and failures. The average simulated trader increases her trade size

[24] In Table 3 and in the regressions underlying Figure 3, R^2s are low both in actual and in simulated data. In the model, investors' exit decisions depend on three factors: returns, beliefs, and the interaction between the two. If investors held identical priors, returns would correlate perfectly with exit decisions. However, because investors priors differ, realized returns become poor predictors of exit decisions. Investors with sufficiently "good" priors remain in the market no matter how low the return is, and those with sufficiently "bad" priors quit after the smallest disappointment.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

by 4.1% after a success in the AR(1) model; this increase is 8.6% in the data. Similar to the data, the sensitivity of trade sizes on returns is lower in the late-career sample. Surviving traders' trade sizes in Figure 4 increase in both the model and data. This increase, however, is more pronounced in the model. For example, the average θ_{25} is 4.3 times as large as θ_1. This ratio is 2.2 in the data.[25]

Fourth, in terms of exploratory trades, 89.8% of simulated investors' first trades are constrained by $\overline{\theta}$. The average first-time trader trades the smallest possible amount because she is pessimistic about her abilities. The empirical value analogous to this estimate depends on what we view as being too small a trade in the data. If we consider an implied same-day return of 1% as being "unreasonable," then 22.2% of actual initial trades are exploratory (see the discussion in Section 4.4). However, if we use 0.5% as the implied return cutoff, this proportion is 59.6%. Although the model and data do not line up perfectly in every dimension, I consider these similarities satisfying given the parsimony of the structural model.

5.9 Robustness Tests

The fit of the model is almost the same when investors' risk aversion increases from $\gamma = 2$ to $\gamma = 5$. The key difference is that because high risk-aversion investors are less willing to assume risk, fewer potentially skilled investors begin trading. Although 1.4% of all individuals are now skilled, one tenth of them never start trading. The reverse survivorship bias, expressed as a gap between true and observed success rates, increases from 0.057 to 0.067. Second, the results are nearly identical when skill μ is unbounded (i.e., $\mu^* = \mu$). This similarity is not surprising given that the $\underline{\mu} = -10\%$ and $\overline{\mu} = 2.5\%$ bounds do not bind investors' beliefs in Table 5. Third, the results are largely the same if I omit trade-1 exit-rate and return-sensitivity moment conditions. Fourth, the results remain similar when I increase the fraction of return variance attributed to the market-wide component from 25% to 50% and 75%. The fraction of skilled individuals is slightly higher in these specifications: 1.3% and 1.6%. The main difference is that the σ_y^2 estimate decreases as investors' signals become noisier. In the 50% specification, $\hat{\sigma}_y^2 = 0.0048$; in the 75% specification, it is 0.0034. These robustness tests suggest that even though extracted beliefs are slightly different for alternative parametrizations of the model, the main conclusions about skill, learning by survival, and reverse survivorship bias are robust.

[25] The effect may be stronger in the model than in the data because actual investors may not always be able to trade as little as $\overline{\theta}$. They must choose an integer number of shares to trade, and this choice also is often subject to "lot-size" restrictions (i.e., the exchange may consider an order valid for the limit order book in stock X only if the number of shares is a multiple of, for example, ten). If early trade sizes are sometimes rounded upward but later trades are not (because surviving investors' optimal trade size is often greater than $\overline{\theta}$), we would expect the data curve to be less steep than the model curve.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

6. Conclusion

I show that a model in which investors learn about their trading skills over time is consistent with several empirical regularities in household trading behavior. Households appear to trade excessively given that their returns do not cover trading costs; trading intensity increases in past performance; and some households quit trading altogether after poor performance. The key mechanism in the learning model is the option value of trading. Because investors learn by trading, they have an incentive to trade even if they think they are unskilled. Even a small probability of being skilled can outweigh expected trading losses.

A structural trading model comes close to matching the number of investors who begin active trading and the decrease in both conditional exit rates and sensitivity of exit to lagged returns. The key findings are that, first, almost three-fourths of the investors who start active trading are unskilled; second, over one-fourth of all traders believe the chance of being skilled is less than 1-in-26; third, approximately 1.2% of all individuals have skill; fourth, because of the learning-by-survival mechanism, performance correlates positively with experience; and fifth, realized returns are significantly downward-biased measures of investors' true abilities.

Although I assume that investors are rational and have well-calibrated prior distributions, the results neither prove nor disprove these assumptions. For example, I do not know if investors learn by trading because there are limits to paper trading or if some other reason is behind this mechanism. Investors' prior distributions may be more dispersed (or tighter) than what they should be, given the sum of all investors' knowledge. However, my main conclusions are robust to these reservations. Investors learn from their own experiences, and the resulting evolution of beliefs is an important driver of investor behavior.

Appendix

Solution to the Simplified Trading Model

The investor's date t value function takes the following form given the isoelastic-utility assumption:

$$V(W_t, \alpha_t, \beta_t) = \frac{W_t^{1-\gamma}}{1-\gamma} B(\alpha_t, \beta_t), \tag{19}$$

where W_t is date t wealth and α_t and β_t are the parameters of the beta-distributed prior about p. The Bellman equation is then

$$\begin{aligned} V(W_t, \alpha_t, \beta_t) &= \max\left\{\max_{\theta_t \geq \overline{\theta}} \left\{ \mathrm{E}\left[\frac{(W_t(1+\theta_t r_t))^{1-\gamma}}{1-\gamma} B(\alpha_{t+1}, \beta_{t+1})\right]\right\}, \frac{W_t^{1-\gamma}}{1-\gamma}\right\} \\ &= \frac{W_t^{1-\gamma}}{1-\gamma} \underbrace{\min\left\{\min_{\theta_t \geq \overline{\theta}} \left\{\hat{p}_t (1+\theta_t)^{1-\gamma} B_u + (1-\hat{p}_t)(1-\theta_t)^{1-\gamma} B_d\right\}, 1\right\}}_{=B(\alpha_t, \beta_t)}, \end{aligned} \tag{20}$$

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

where $\hat{p}_t \equiv \frac{\alpha_t}{\alpha_t+\beta_t}$, $B_u \equiv B(\alpha_t + 1, \beta_t)$, and $B_d \equiv B(\alpha_t, \beta_t + 1)$. Equation (20) assumes that $\gamma > 1$; the $\gamma < 1$ solution replaces the minimization problems with maximization problems. The outer minimization problem in $B(\alpha_t, \beta_t)$ compares the indirect utility from trading to that from quitting. The inner minimization problem is solved by evaluating it at the constrained optimal trade size,

$$\theta_t^* = \max \left\{ \frac{(\hat{p}_t B_u)^{\frac{1}{\gamma}} - ((1-\hat{p}_t) B_d)^{\frac{1}{\gamma}}}{(\hat{p}_t B_u)^{\frac{1}{\gamma}} + ((1-\hat{p}_t) B_d)^{\frac{1}{\gamma}}}, \ \overline{\theta} \right\}. \tag{21}$$

The entire problem can be solved recursively using a recombining binomial tree. Each node j on day t is characterized by three values: parameters α_t^j and β_t^j of the posterior distribution, which depend only on the location of the node, and $B(\alpha_t^j, \beta_t^j)$, which can be computed recursively by starting from date $T-1$. On date $T-1$, $B_d \equiv 1$ and $B_u \equiv 1$, and so all values $B(\alpha_{T-1}^j, \beta_{T-1}^j)$ can be computed from Equations (20) and (21). After computing $B(\alpha_{T-1}^j, \beta_{T-1}^j)$ for every j, values $B(\alpha_{T-2}^j, \beta_{T-2}^j)$ can be computed, and so forth.

References

Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean. 2005. Do Individual Day Traders Make Money? Evidence from Taiwan. Working Paper, University of California–Berkeley.

———. 2009. Just How Much Do Individual Investors Lose by Trading? *Review of Financial Studies* 22:609–32.

Barber, B. M., and T. Odean. 2000. Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *Journal of Finance* 55:773–806.

———. 2001. Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment. *Quarterly Journal of Economics* 116:261–92.

Barber, B. M., T. Odean, and M. Strahilevitz. 2004. Once Burned, Twice Shy: Naïve Learning, Counterfactuals, and the Repurchase of Stocks Previously Sold. Working Paper, University of California–Berkeley.

Barberis, N., and W. Xiong. 2009. What Drives the Disposition Effect? An Analysis of a Long-standing Preference-based Explanation. *Journal of Finance* 64:751–84.

Bauer, R., M. Cosemans, and P. Eichholtz. 2007. The Performance and Persistence of Individual Investors: Rational Agents or Tulip Maniacs? Working Paper, Maastricht University.

Brennan, M. J. 1998. The Role of Learning in Dynamic Portfolio Decisions. *European Economic Review* 1: 295–306.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick. 2009. Reinforcement Learning and Savings Behavior. *Journal of Finance* 64:2515–34.

Choi, J. J., D. Laibson, and A. Metrick. 2002. How Does the Internet Affect Trading? Evidence from Investor Behavior in 401(k) Plans. *Journal of Financial Economics* 64:397–421.

Coval, J. D., D. A. Hirshleifer, and T. Shumway. 2005. Can Individual Investors Beat the Market? Working Paper, University of Michigan.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill College.

Detemple, J. B. 1986. Asset Pricing in an Economy with Incomplete Information. *Journal of Finance* 41: 383–92.

French, K. R. 2008. Presidential Address: The Cost of Active Investing. *Journal of Finance* 63: 1537–73.

Glaser, M., and M. Weber. 2009. Which Past Returns Affect Trading Volume? *Journal of Financial Markets* 12:1–31.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Goetzmann, W. N., and A. Kumar. 2008. Equity Portfolio Diversification. *Review of Finance* 12:433–63.

Gould, J. L., and C. G. Gould. 1994. *The Honey Bee*. New York: Scientific American Library.

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.

———. 2009. Sensation Seeking, Overconfidence, and Trading Activity. *Journal of Finance* 64: 549–78.

Grinblatt, M., M. Keloharju, and J. T. Linnainmaa. 2010. IQ and Stock Market Trading. Working Paper, University of Chicago.

Harris, J. H., and P. H. Schultz. 1998. The Trading Profits of SOES Bandits. *Journal of Financial Economics* 50:39–62.

Harris, L. 2003. *Trading and Exchanges: Market Microstructure for Practitioners*. New York: Oxford University Press.

Harvey, A. C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, MA: Cambridge University Press.

Hieronymus, T. A. 1977. *Economics of Futures Trading*. New York: Commodity Research Bureau.

Ivković, Z., C. Sialm, and S. Weisbenner. 2008. Portfolio Concentration and the Performance of Individual Investors. *Journal of Financial and Quantitative Analysis* 43:613–56.

Krebs, T. 2003. Human Capital and Economic Growth. *Quarterly Journal of Economics* 118:709–44.

Kumar, A. 2009. Who Gambles in the Stock Market? *Journal of Finance* 64:1889–933.

Linnainmaa, J. T. 2010. Reverse Survivorship Bias. Working Paper, University of Chicago.

Lo, A. W., A. C. MacKinlay, and J. Zhang. 2002. Econometric Models of Limit-order Executions. *Journal of Financial Economics* 65:31–71.

Mahani, R. S., and D. Bernhardt. 2007. Financial Speculators' Underperformance: Learning, Self-selection, and Endogenous Liquidity. *Journal of Finance* 62:1313–40.

Nicolosi, G., L. Peng, and N. Zhu. 2009. Do Individual Investors Learn from Their Trading Experience? *Journal of Financial Markets* 12:317–36.

Odean, T. 1998. Are Investors Reluctant to Realize Their Losses? *Journal of Finance* 53:1775–98.

———. 1999. Do Investors Trade Too Much? *American Economic Review* 89:1279–98.

Pástor, Ľ., and P. Veronesi. 2003. Stock Valuation and Learning about Profitability. *Journal of Finance* 58: 1749–90.

———. 2009. Learning in Financial Markets. *Annual Review of Financial Economics* 1:361–81.

Ross, R. L. 1975. Financial Consequences of Trading Commodity Futures Contracts. *Illinois Agricultural Economics* 15:27–31.

Seru, A., T. Shumway, and N. Stoffman. 2010. Learning by Trading. *Review of Financial Studies* 23: 705–39.

Shefrin, H., and M. Statman. 1985. The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence. *Journal of Finance* 40:777–90.

Sørensen, M. 2007. How Smart Is Smart Money: A Two-sided Matching Model of Venture Capital. *Journal of Finance* 62:2725–62.

Storesletten, K., C. I. Telmer, and A. Yaron. 2007. Asset Pricing with Idiosyncratic Risk and Overlapping Generations. *Review of Economic Dynamics* 10:519–48.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011

Taylor, L. A. 2010. Why Are CEOs Rarely Fired? Evidence from Structural Estimation. *Journal of Finance* 65:2051–87.

Teweles, R. J., and F. J. Jones. 1999. *The Futures Game: Who Wins, Who Loses, & Why*. New York: McGraw-Hill, 3rd ed.

Williams, D. 1991. *Probability with Martingales*. Cambridge, MA: Cambridge University Press.

Xia, Y. 2001. Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation. *Journal of Finance* 56:205–46.

Downloaded from rfs.oxfordjournals.org at Serials Department on March 14, 2011