

Learning By Trading

Amit Seru, Tyler Shumway, and Noah Stoffman*

This version: February 15, 2009

First version: March 15, 2006

Forthcoming, *Review of Financial Studies*

Abstract

Using a large sample of individual investor records over a nine-year period, we analyze survival rates, the disposition effect and trading performance at the individual level to determine whether and how investors learn from their trading experience. We find evidence of two types of learning: some investors become better at trading with experience, while others stop trading after realizing that their ability is poor. A substantial part of overall learning by trading is explained by the second type. By ignoring investor attrition, the existing literature significantly overestimates how quickly investors become better at trading.

*We thank Brad Barber, Dan Bernhardt, James Choi, Juhani Linnainmaa, Uday Rajan, Mark Seasholes, Morten Sorensen, Matthew Spiegel (the editor), Ning Zhu, two anonymous referees, and seminar participants at the 2007 Western Finance Association meeting, the 2007 Wharton Household Portfolio Choice conference, Carnegie Mellon University, Nanyang Technical University, National University of Singapore, Singapore Management University, University of California at Irvine, University of Edinburgh, University of Manchester, University of Toronto, and the University of Michigan for helpful comments. Any remaining errors are ours. We are grateful to Jussi Keppo for helping us acquire the data used in this study, and to the Mitsui Life Financial Research Center at the University of Michigan for funding. We also thank Jarkko Heinonen and Monica Bergström at the Helsinki Exchange for helping us with the data. Seru (amit.seru@ChicagoBooth.edu) is at Chicago GSB, Shumway (shumway@umich.edu) is at the Ross School of Business at the University of Michigan, and Stoffman (nstoffma@indiana.edu) is at the Kelley School of Business at Indiana University. Address correspondence to Tyler Shumway at Ross School of Business, University of Michigan, 701 Tappan Street, Ann Arbor, Michigan, 48109. Phone: (734) 763-4129.

Academics have recently shown an interest in the investment behavior and performance of individuals, a field that has been called ‘household finance’ by Campbell (2006). Over the past decade, several researchers have documented a number of behavioral biases among individual investors. More recently, researchers have found evidence that some individual investors are more informed or skilled than others.¹ Considering these findings, it is natural to ask how skilled or informed investors acquire their advantage. For example, do investors learn by trading? If so, to what extent do investors improve their ability and to what extent do they learn about their inherent ability? And how quickly do investors learn? In this paper, we exploit trading records to study both average investor performance and the strength of the behavioral bias known as the disposition effect.² We correlate performance and disposition with investor experience and investor survival rates to determine whether and how investors learn by trading.

Motivated by the existing economics literature on learning, we consider two specific ways in which investors can learn. First, in the spirit of classical learning-by-doing models (Arrow 1962, Grossman, Kihlstrom, and Mirman 1977), investors might improve their ability as they trade (‘learning by doing’). Second, as investors trade they might realize that their inherent level of ability is low and decide to stop trading (‘learning about ability’). Although these types of learning are different, they are not mutually exclusive, and the primary contribution of our paper is to separate these types of learning empirically and estimate the importance of each. Our results provide robust evidence for both types of learning but the magnitudes of the learning estimates

¹Coval, Hirshleifer, and Shumway (2005) document significant performance persistence among individuals. Ivković and Weisbenner (2005) find that individuals place more informed trades in stocks of companies located close to their homes, and Ivković, Sialm, and Weisbenner (2008) show that individuals with more concentrated portfolios tend to outperform those who are more diversified. Linnainmaa (2007) finds that individuals who trade with limit orders suffer particularly poor performance.

²The disposition effect is the propensity of investors to sell assets on which they have experienced gains and to hold assets on which they have experienced losses. The effect was first proposed by Shefrin and Statman (1985), and was subsequently documented in a sample of trading records from a U.S. discount brokerage firm by Odean (1998). The effect has been found in other contexts, including in Finland (Grinblatt and Keloharju 2001), China (Feng and Seasholes 2005, Shumway and Wu 2006), and Israel (Shapira and Venezia 2001); among professional market makers (Coval and Shumway 2005), mutual fund managers (Frazzini 2006), and home sellers (Genesove and Mayer 2001); and in experimental settings (Weber and Camerer 1998). We focus on the disposition effect because it is a robust empirical finding and it is relatively easy to measure.

suggest that most of the learning by trading occurs as individuals learn about their own ability and low-ability investors stop trading. This implies that by ignoring investor attrition the existing literature substantially overstates how quickly investors become better at trading.

To clarify the model of learning we have in mind, consider the case of an individual who decides to begin trading. The investor must decide which of the myriad sources of market information and investment advice available to her to take seriously. She could consult standard news sources, internet sites, investment newsletters, and neighbors or friends. She might also consider the advice of brokers, news analysts, authors of books and magazines, and finance professors. To the extent that these sources fail to agree completely, individuals must determine how much decision weight to assign to each source. Moreover, the quality of these sources is likely to differ across individuals: some investors may know executives at a firm, while others will not. As investors begin trading, they can learn to which of the various sources they should pay more attention. This can be thought of as *improving* ability through learning. Investors who only have access to poor sources of information cannot improve by focusing more on particular sources. Instead, they will learn that they have no useful information and will stop trading actively, choosing instead to invest in a passive investment such as an index fund.³ This is also a type of learning, but rather than improving her ability, the investor learns *about* her ability.

Since investors who learn that their ability is poor should stop trading, we empirically examine whether investors learn about their ability by examining attrition in our data. Once we account for time-invariant individual heterogeneity and *endogenous* attrition, any improvement in investors' performance and any reduction in their disposition effect that comes with experience should reflect a direct improvement in their ability. By tracking the survival rates, performance, and the level of the disposition effect for each individual over time, we are able to differentiate between these

³This example is similar in spirit to well-known 'bandit' problems. For example, Bolton and Harris (1999) study the strategic interaction of agents in an experimentation game. In our setting, we can think of investors as being randomly assigned slot machines with different unobservable expected payoffs, and experimenting to learn what the payoffs are. The random assignment is analogous to the assignment of inherent ability.

two types of learning. In doing so, ours is the first paper to identify and measure both types of learning.⁴

We test our hypotheses with a remarkable dataset that includes the complete trading records of investors in Finland from 1995 to 2003, including more than 22 million observations of trades placed by households. We use these data to estimate disposition and calculate performance at the account level. Our disposition estimates indicate that a median individual in our sample is 2.8 times more likely to sell a stock when its price has risen since purchase than when its price has fallen. We exploit the panel structure of our data to examine whether individual investors learn to avoid the disposition effect and improve their performance as they trade. In particular, we estimate the mean return and the disposition effect for each account and year in our sample and relate these estimates to experience, past returns, and various demographic controls.

We identify learning about ability in the data by adjusting for survivorship and heterogeneity with a modified Heckman selection model that allows for individual fixed effects. The model is a two-stage instrumental variables model that adjusts for the possibility that the composition of the sample is endogenous while accounting for any cohort effects. We construct two variables for use as instruments: the first is an indicator of whether the individual inherited shares in the previous year due to the death of a relative, and the second is the proportion of accounts in an investor's zip code that are held by active traders. We argue that both of these variables are likely to satisfy the necessary exclusion restrictions—that is, they are likely to affect the probability of an investor remaining in the sample, but are unlikely to affect changes in the investor's performance or disposition effect except through their effect on survival.

Estimates of the selection model confirm that investors with poor performance are more likely to cease trading. An investor whose performance is one standard deviation worse than the mean

⁴Two recent papers examine the relation between investor sophistication and behavior. Calvet, Campbell, and Sodini (2007) find that more sophisticated households in Sweden are more likely to begin trading and less likely to stop than are less sophisticated households. Korniotis and Kumar (2008) show that investors in the U.S. have better performance when their demographic characteristics indicate they are likely to have high cognitive abilities. Neither of these studies directly examines whether investors learn.

is about 15 percent less likely to continue trading. Our instruments are also good predictors of whether individuals continue trading. The second stage of our selection model suggests that investors learn by trading: after accounting for survivorship, an extra 100 trades is associated with an improvement in average returns of approximately 3.6 basis points (bp) over a 30-day horizon (or about 30 bp per year), and a reduction in the disposition effect of about two percent. If we measure experience with the number of years an individual has been trading instead the number of trades she has placed, the improvement is negligible. This implies that individuals actually have to place trades to learn, or that learning truly occurs by trading. Perhaps more important, the magnitude of the learning estimates presented above are about *two to four times higher* when not adjusted for investor attrition. This implies that the learning we document occurs primarily as investors learn about their ability.

Differentiating between types of learning allows us to determine in a meaningful way how quickly investors learn. Estimating learning without accounting for heterogeneity and attrition results in inflated estimates of improvement that do not correspond to the experience of any particular type of investor. Correctly measuring the speed of learning is important for a number of reasons. If investors learn quickly and there is low turnover in the population of investors, behavioral biases are unlikely to affect asset prices significantly. Moreover, if they learn relatively quickly then the ‘excessive’ trading documented by Odean (1999) and Barber and Odean (2001) may be justified, because investors may optimally choose to trade more actively if they know they will improve with experience. Finally, thinking about both the speed and the type of learning has implications for market efficiency. For example, if many inexperienced investors begin trading around the same time, and they learn slowly, their trading could lead to time-varying market efficiency. In fact, we find evidence that suggests that investors learn less about their ability in years with positive market returns. In this respect, our paper contributes to the theoretical literature concerned with

the survival and price impact of noise traders (see, for example, DeLong, Shleifer, Summers, and Waldmann (1991) and Kogan, Ross, Wang, and Westerfield (2006)).

Some evidence about learning by trading exists. Feng and Seasholes (2005) give evidence that investors, in aggregate, display significantly less disposition over time, estimating that for sophisticated investors the disposition effect is essentially attenuated after about 16 trades. These estimates do not appear to be consistent with the findings of Frazzini (2006), which shows that mutual fund managers, who trade substantially more than most individuals, display a significant disposition effect. Furthermore, since Feng and Seasholes (2005) do not adjust for heterogeneity and attrition, it is impossible to tell whether their estimates imply that a particular investor who trades 16 times will no longer exhibit the disposition effect or whether they imply that most investors who exhibit the disposition effect will cease trading before they place 16 trades. Nicolosi, Peng, and Zhu (2008) show that the trading performance of individuals appears to improve with trading experience, estimating that individuals can improve their risk-adjusted portfolio return by about two percent per year (or about 0.8 bp per day) over a 3-year period. This seems quite large. Again, it is not possible to tell whether this estimate is driven by the most successful investors surviving or by the least successful investors improving their ability.⁵

While our tests have some features in common with existing papers, they differ from the literature in a number of important respects. First, unlike other papers, our tests use measures of performance and estimates of the disposition effect that are specific to *individuals*, allowing us to track particular individuals over time. This permits us to control for investor heterogeneity and survivorship effects, which allows us to separate the two types of learning. It also ensures that each observation in almost all of our tests is an average or regression coefficient for one individual in one particular year. Finally, using only one observation per person per year reduces the likelihood that our standard error estimates are incorrect because of correlation among our regression residuals.

⁵Other papers that find some learning in various settings include Pastor, Taylor, and Veronesi (2007), Choi, Laibson, Madrian, and Metrick (2007), Linnainmaa (2006), Barber, Odean, and Strahilevitz (2004), Pastor and Veronesi (2003), and List (2003).

Given the unique features of our data and our test methods, the results of our hypothesis tests add significantly to the literature on financial learning.

The rest of the paper is organized as follows. Section 1 provides detail on our data, while Section 2 describes the hypotheses we test and our statistical methods. Section 3 discusses our results and Section 4 concludes. We give some details about our statistical methods in an appendix.

1 Data and Methods

The data used in this study come from the central register of shareholdings in Finnish stocks maintained by Nordic Central Securities Depository (NCSD), which is responsible for the clearing and settlement of trades in Finland. Finland has a direct holding system, in which individual investors' shares are held directly with the CSD. Since our data come from the CSD, they reflect the official record of holdings and are therefore of extremely high quality. The data cover all trading in all Finnish stocks over a nine-year period. Grinblatt and Keloharju (2000, 2001a, 2001b) use a subset of the same data, comprising the first two years of our sample period.⁶ The data include the transactions of nearly 1.3 million individuals and firms, beginning in January, 1995 and ending in December, 2003. In all, more than 22 million trades by individual investors are included. On average, individuals hold 12.6 percent of all equity in Finland while financial institutions hold 9.6 percent, 34.7% is held by the government, and nonfinancial firms hold 33.4% during our sample period. (Additional statistics on investor trading are provided in Table 1, below.)

Our data allow us to reconstruct the portfolio of stocks held by each account on a daily basis. Using these holdings, we construct a proxy for wealth by calculating the average daily marked-to-market portfolio value for each investor. We also calculate the average value of trades placed by an investor each year. To measure sophistication, we note that investors who trade options are likely to be more familiar with financial markets. This is particularly true in our setting because many

⁶These references provide a detailed discussion of the data.

of the options in our data are granted to corporate executives as part of compensation. Therefore, while we do not include options trades in our estimates of disposition, we use whether an investor ever trades options as a proxy for sophistication. We also count the number of distinct securities traded by an investor over the sample period, and use this as a measure of portfolio diversification.

Table 1 provides summary statistics for the new accounts in our dataset. New accounts are those that place their first trade in 1995 or a subsequent year; they have no recorded initial positions. Panel A includes all new accounts that place at least one trade during our sample period (1995–2003), while Panel B gives results only for those new accounts for which we are able to estimate the disposition coefficient at least once. We only attempt to estimate the disposition coefficient if an individual has placed at least seven round-trip trades in a given year, although even with this restriction the procedure to maximize the likelihood function does not always converge. The last two rows of each panel are indicator variables, taking a value of one if the investor: (a) trades options; or (b) is female; and zero otherwise.

Comparing Panels A and B, it is apparent that the subset of investors for whom disposition coefficients are available is somewhat different than the larger population. By construction, the accounts in Panel B place more trades, but they also have larger portfolios, trade larger amounts of money, trade in a wider selection of securities, and are somewhat younger. As well, investors for whom we can estimate disposition are more likely to trade options (17%) than the overall sample (3%). Since we are only able to estimate disposition for investors who trade with some frequency, this likely results from the fact that investors who trade options are simply more likely to trade in general.

Figure 1 shows the number of accounts (including both new and existing accounts) that place one or more trades in each year. There is considerable variation in the number of accounts placing trades over time, from a low of 54,196 accounts in 1995 to a high of 311,013 accounts in 2000.

Additions of new accounts follows a similar pattern. We discuss entry and exit from the sample in more detail in the next section.

1.1 Measuring performance

Investor performance is the primary variable that we correlate with experience to test our hypotheses. Measuring the performance of individual investors is a significant challenge for a number of reasons. For example, it is not obvious how we should compare the performance of investors with different holding periods. Given the challenges associated with calculating performance, we take a straightforward approach that is nevertheless likely to capture much of the relevant information in the individual's returns. In short, we calculate the returns earned by each purchased stock in the 30 trading days following each of an investor's purchases. Importantly, we truncate this calculation window at the length of the actual holding if it is shorter than 30 days. We choose to focus on 30-day returns because the median holding period in our data is 39 trading days, but all of our findings remain unchanged if we use a 10-, 45-, or 60-day holding period. Our annual performance measure is the average of these 30-day returns. We provide details of our performance measures in the Appendix.⁷

1.2 Measuring the disposition effect

The other outcome variable that we track to evaluate whether investors learn with experience is the disposition effect. We estimate the disposition effect using a hazard model estimated with daily transaction and return data in each year for each investor in our sample. Estimating the effect for each year means that each individual/year coefficient is estimated with a unique data set which is completely unaffected by the individual's trades in other years. We estimate a hazard model of the

⁷Some returns in our observation period are quite volatile. Particularly in the few years at the end of the decade, some stocks would see large returns in short periods, and some investor returns are quite large or sensitive to holding periods. In our main analysis, we confirm that our results are not sensitive to removing investors with extreme portfolio returns, or winsorizing returns.

form,

$$h_{i,j}(t) = \phi_i(t) \exp\left\{\beta_i^d \text{PriceDummy}_t + \text{controls}\right\}, \quad (1)$$

where $h_{i,j}(t)$ is investor i 's probability of selling position j on date t , conditional on not having sold prior to date t . PriceDummy_t takes a value of one when the stock price on date t is above investor i 's purchase price, and zero otherwise. β_i^d therefore measures investor i 's susceptibility to the disposition effect, which we term the 'disposition coefficient'. Details of the estimation and additional discussion are provided in the Appendix.

1.3 Measuring experience

We measure investing experience with both the number of years that an investor has been trading and with the cumulative number of trades that an investor has placed. Of course, investors may gain experience by actively trading securities and observing the results of each trade. If this is the primary way in which investors learn, then cumulative trades will predict future investment performance and the disposition effect. However, investors may also learn by observing market quantities and considering the outcomes of hypothetical trades based on, for example, a particular information source. If investors mainly learn this way, then years of experience will be a better predictor of investment performance and the disposition effect than cumulative trades.⁸

2 Hypotheses

We test two hypotheses in this paper to examine if individual investors learn by trading. Both hypotheses are about how the characteristics of individuals change with time, so we test both with panel data. Our first hypothesis (H_1) is that investors learn to improve their ability over time—

⁸It is also possible, of course, that investors learn by considering returns to hypothetical trades before they ever start trading ('paper trading'). If this is the only way in which investors learn, then we should find no evidence of learning by trading (either H_1 or H_2). Learning from paper trading may be difficult for a number of reasons, including that it requires significant discipline to keep track of hypothetical trades, and that transaction prices for hypothetical trades are not observable.

that is, ‘learning by doing’ occurs. The second hypothesis (H_2) is that investors learn about their inherent ability by trading.

H_1 is in the spirit of classical learning-by-doing models (Arrow 1962, Grossman, Kihlstrom, and Mirman 1977), which argue that the productivity of agents increases with experience. The notion here is that investors improve over time since they should, for instance, be able to choose better combinations of various information sources that help them choose which stocks to hold. The main testable prediction of H_1 is that performance should improve over time for investors who remain active.

H_2 is best understood in the context of Mahani and Bernhardt’s (2007) model, where individuals do not initially know their ability. However, individuals obtain information about their skills through their trading experience. Only a small fraction of the population is adept at identifying profitable trading opportunities. In equilibrium, each period some novice speculators enter the financial markets. Recognizing that most traders lack financial acumen, the novices first experiment on a small scale. They then use the information contained in their trading profits to decide whether to continue. Those who earn sufficient profits conclude that they are likely to be skilled and expand their speculative activities. However, some speculators do less well, conclude that they are more likely to be inept traders, and stop actively trading. The main testable prediction of H_2 is that investors who experience losses stop actively trading, while those with sufficient profits remain and increase their trading intensity over time.

If either H_1 or H_2 are true, we should find an improvement in the average performance of the pool of investors over time. We therefore begin by testing a simple learning model to look for evidence of learning overall. If we assume a world with random attrition and no heterogeneity in ability among investors, a simple regression of performance on experience gives an estimate of the first type of learning (learning-by-doing). Alternatively, it gives an estimate of the second type of learning (learning-about-ability) if the pool of investors improves over time as the low-ability

investors stop trading and those with sufficient profits continue to trade. To disentangle H_1 and H_2 , we therefore need to examine the role played by individual heterogeneity and attrition in our data. Accounting for unobserved individual heterogeneity and attrition provides us with an assessment of H_2 and any learning that remains provides an estimate of H_1 .

3 Results

We present our empirical findings in this section. We begin by presenting the results of our tests relating to performance and disposition in Section 3.1. The simple learning model is presented in Section 3.2, and we deal with unobserved heterogeneity in Section 3.3. Section 3.4 examines the importance and magnitude of survivorship effects, and we present evidence on changes in trading intensity and risk taking in Section 3.5. Finally, a number of robustness tests are presented in Section 3.6.

3.1 Performance and disposition

We start by examining trader performance to be sure that our performance measure has similar properties to the measures used in the literature. Previous papers, in particular Odean (1998), have shown that average investor performance is worse than that of the market portfolio. Poor performance by average investors is also a prediction of the model in Mahani and Bernhardt (2007). We test whether individuals on average outperform a market index by calculating the average return to a stock purchased by an individual investor net of the market return. Calculating this average at a 30-day horizon (using our convention of using a shorter holding period if the individual sells the stock before 30 days) yields an average return net of the market of -48 basis points. At a 60-day horizon, the average net return is -50 basis points. At both of these horizons, returns net of the market return are quite statistically significantly negative.

While the average performance of individual investors is likely to be quite poor, Coval, Hirshleifer, and Shumway (2005) shows that some individuals persistently outperform others. Again, performance persistence among individuals is an implication of Mahani and Bernhardt (2007). We test whether there is any persistence in investor performance in three related ways. First, we regress each investor’s average 30-day return in year t on the investor’s average return in year $t-1$ and year fixed effects. Using year fixed effects adjusts for time series variation in average market returns. The estimated coefficient in this regression is 0.183 ($p < 0.0001$), very statistically and economically significant. Second, we calculate each investor’s average return in two disjoint time periods, 1995–1999 and 2000–2003. We then calculate the Spearman rank correlation between the return series from the first period with that from the second period. This correlation is 0.164 ($p < 0.0001$), again quite statistically and economically significant. Our third test involves sorting investors in each year into performance quartiles, and then plotting the average performance of each of those quartiles for the next several years. This plot, which appears in Figure 2, again gives evidence that the most successful investors in the past continue to outperform the least successful investors for at least a couple of years. Results calculated with alphas instead of raw returns are qualitatively the same. These results confirm that there is a degree of persistence in individual returns.⁹

Next, we examine the other outcome variable that will use in our analysis—the disposition effect. The disposition effect is quite large in our data. Figure 3 is a plot of the relation between the propensity to sell an existing position (the hazard ratio) and the position’s holding period return. To generate this plot, we group all investors and estimate one hazard model each year. We group the data for this procedure so we can estimate a model with many covariates, but almost

⁹We use a four-factor model to adjust returns for the known risk factors of Fama and French (1993) and Carhart (1997). To construct the HML and SMB factors, we take quarterly data on shares outstanding and book value of equity using definitions as similar as possible to those of Fama and French, and construct portfolios in the manner they describe. We measure a firm’s risk in a particular year by regressing daily returns in excess of the risk-free rate on the daily returns of the four factor-mimicking portfolios using the Scholes and Williams (1977) method to avoid well-known problems associated with nonsynchronous trading. As with the raw returns, we calculate holding period returns over a fixed interval of 30 days, but truncate the holding period at the length of the actual holding period.

all of the tests that follow are based on individual-level results.¹⁰ The conditional hazard ratio is remarkably similar across years. The plot shows an obvious kink in the hazard ratio near zero: investors are clearly more likely to sell a stock if it has increased in value since the purchase date. This provides strong support for the presence of a disposition effect in aggregate, consistent with the extensive literature cited above.

Turning to our main individual-level disposition regressions, we require that an investor place at least seven round-trip trades in a year to be included in the sample, and we run the regression for each investor-year to generate a separate disposition coefficient whenever possible. While this filter drastically reduces our sample size, it is necessary to ensure that our coefficients of interest are identified. Panel 3 of Table 1 summarizes these estimates. The median disposition coefficient in the cross-section of investors across years is 1.04, which is economically quite large. This coefficient implies that the median new investor in our data is $e^{1.04} = 2.8$ times more likely to sell a stock whose price is above its purchase price than a stock that has fallen in value since the time of purchase. None of the controls is statistically significant in the cross-section.

Before we can consider whether investors learn to avoid the disposition effect, we need to argue that the effect is in fact a behavioral bias. In particular, one necessary condition for disposition to be a costly behavioral bias is that investors with more disposition have inferior investment performance. If disposition is unrelated to investment performance, investors with the effect have little incentive to learn to avoid it. To get a sense of how returns vary with disposition, we examine average investor returns across quintiles of the disposition coefficient. In this sort, the disposition coefficients are always estimated one year before the average returns are calculated, so disposition coefficients and average returns are not mechanically correlated in any way. For each quintile, Figure 4 graphs the average return earned by investors over different horizons from the purchase date. Returns

¹⁰Rather than using only one indicator variable as in Equation (7), we use 20 dummy variables corresponding to different 1 percent return ‘bins’ in this model. In Figure 3 we plot the dummy variable coefficients by year. The sum of these coefficients times their corresponding dummy variables is multiplied by the baseline hazard rate to give the actual conditional hazard rate.

are substantially higher in the lowest disposition quintile than in the highest disposition quintile. For example, in the 30 days following a purchase, a stock's price increases 46 bp on average when bought by an investor in the lowest disposition quintile, compared to a decline of 54 bp if purchased by an investor in the highest disposition quintile. The differences between high- and low-quintile average returns range from 17 bp at the 10-day horizon to 131 bp at the 45-day horizon. These differences are both economically and statistically large, leading us to conclude that individuals with high disposition effect coefficients have relatively poor investment performance.

We can also verify that disposition is costly by mimicking the tests of Odean (1998), as presented in Table 2. The idea is to compare the returns of stocks sold at a gain to those that could have been sold at a loss but were not. In order to see the difference in cost of the disposition effect to investors in the low- and high-disposition groups, we implement a difference-in-difference specification. Using a regression framework instead of simple averages allows us to include year fixed effects to absorb year-specific factors such as average market performance. We define low- and high-disposition groups as the bottom and top quintile of the disposition coefficient, and exclude investors who do not fall into either of these two quintiles. The negative coefficient in the first row indicates that the effect documented by Odean is reversed for the low-disposition group: stocks sold for a gain by low-disposition investors subsequently underperform those that could have been sold at a loss. The difference-in-difference coefficient is significantly positive at the three horizons we consider, indicating that the effect documented by Odean (1998) is stronger for individuals in our high-disposition group than in the low-disposition group. That is, high-disposition investors sell stocks that subsequently outperform the stocks they could have sold, but low-disposition investors do not. This test reiterates that disposition effect is costly to investors.

Another necessary condition for disposition to be a behavioral bias is that disposition is a somewhat stable, predictable attribute of a particular investor. We test this conjecture by estimating the disposition effect at the investor level in adjacent time periods. Each set of estimates comes

from a completely disjoint dataset. Any trades that are not closed at the end of the first period are considered censored in the model estimated with first period data. Therefore, any trades that are not closed at the end of the first period are completely ignored in the model estimated with second period data. We explore the stability of disposition coefficients by estimating the rank correlation of account-level disposition coefficients over the two periods, testing whether the rank correlation is significantly different from zero. We estimate the rank correlation between an investor’s disposition coefficient in year t and their coefficient in year $t - 1$ to be 0.364, suggesting that there is a fair degree of persistence in the individual’s disposition coefficient. This correlation is extremely statistically significant. Taken together, these results provide strong evidence that the disposition effect is a widespread and economically important behavioral bias that is present in each year of our study.

3.2 Simple learning model

We start our analysis by examining if more experienced investors have better investment performance. We should find an improvement in the average performance of the pool of investors over time if either H_1 or H_2 is true. Therefore, before disentangling H_1 and H_2 , the simple learning model helps us assess whether either or both of these effects exist. We adjust for individual heterogeneity and survivorship in the next section. Our analysis begins by estimating a simple learning model of the form:

$$y_{i,t+1} = \alpha + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2 + \delta X_{i,t} + \epsilon_{i,t} \quad (2)$$

where the dependent variable is either the investor’s performance (average return following purchases) or the investor’s disposition coefficient. The primary coefficient of interest in this specification is on Experience. In the specification we include Experience² to allow investors to learn faster during earlier years. We proxy for investors’ trading experience by either years of experience

or cumulative number of trades placed, and also include a vector of controls ($X_{i,t}$) that might be expected to affect individual's performance as he or she trades, such as the individual's average total daily portfolio value (PortVal).

Columns 1 and 2 of Table 3 report the results of the performance learning regression. When experience is measured either in number of years or cumulative trades, it is positively and significantly related to average returns. An additional year of experience increases average 30-day post-purchase returns by $41 - 4 = 37$ bp, or approximately 3 percent at an annualized rate. An additional 100 trades increases returns at slightly over one-fourth of this rate. Again, results estimated with alpha instead of raw returns are quite similar (unreported).¹¹ While these estimates are encouraging, the speed of learning they imply seems almost implausibly large. For instance, taking the regression parameters at face value, an investor with 8 years of experience should outperform a new investor by about 22 percent per year. While we observe some heterogeneity in investor ability (or some performance persistence) it is not nearly large enough to justify these large coefficients.

Columns 3 and 4 of Table 3 present our results for the disposition learning regressions. To reduce the weight given to disposition coefficients that are not estimated very precisely, we estimate the regressions with weighted least squares, where the weights are proportional to $1/\widehat{\text{Var}}(\beta^d)$ from our hazard regression (equation (7) in the Appendix). Column 3 shows that disposition declines with experience ($\beta_1 < 0$). Moreover, investors tend to slow down in their learning as they gain experience since $\beta_2 > 0$. Wealthier traders, investors who trade more securities, and investors who earned higher returns in the previous year all have lower levels of disposition. Column 4 indicates that an additional 100 trades reduces the disposition coefficient by 0.041, which is similar to the coefficient on Experience in Column 3. In other words, a year of experience or 100 trades has approximately the same effect on disposition. In each of the specifications the estimated YearsTraded and CumulTrades coefficients are statistically significant at the 1 percent level. Economically, how-

¹¹Results estimated with both YearsTraded and CumulTrades terms included in the same regression are also very similar to those reported.

ever, our results suggest that investors learn relatively slowly. Specifically, the estimates in Column 3 suggest that an additional year of experience corresponds to a reduction in the disposition coefficient of approximately 0.05. To provide some context for this estimate, note that the unconditional median disposition coefficient in our sample is 1.04. An extra year of experience decreases this by about five percent.

As mentioned in Section 2, these estimates could reflect the amount of learning-by-doing in a world with exogenous attrition and no investor heterogeneity or they could reflect learning-about-ability if the pool of investors improves over time due to attrition of low-ability traders. We now move on to testing H_1 and H_2 by examining the role played by individual heterogeneity and attrition in our data.

3.3 Learning model with individual fixed effects

The simple learning model does not account for unobserved investor heterogeneity. As indicated in Figure 1, investor participation in our sample changes significantly over time. Consequently, cohort effects might make it appear as if there is learning even if there is none. For instance, if the number of trades placed by an investor is a noisy measure of ability, our experience variables could capture learning even if investors have constant ability and high-ability investors trade more than low-ability ones. With this in mind, our next model exploits the long time series of performance and disposition estimates available for each investor to assess the impact of time-invariant unobserved investor-level factors such as ability, education or wealth on investor learning:

$$y_{it} = \alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2 + \delta X_{i,t} + \epsilon_{i,t}, \quad (3)$$

where the fixed effects, α_i , control for unobserved individual heterogeneity.

The performance results on `YearsTraded`, reported in Column 5 of Table 3, suggest that an investor with one year of experience will earn 22 bp more than an inexperienced investor over a

30-day horizon. Column 6 indicates that a similar increase in returns comes from an additional 400 trades. This suggests that though investor performance improves with experience, accounting for individual heterogeneity reduces the estimates by about 50 percent. Results from the disposition regressions including individual fixed effects are displayed in Columns 7 and 8. While YearsTraded is no longer significant in these regressions, Column 8 suggests that 100 trades reduces the disposition coefficient by approximately 0.03. Comparing these estimates to those from the regression without fixed effects, we find that the learning estimates again are reduced by roughly one-half.

These regressions assume that individuals stop trading for purely exogenous reasons. The model shows that some learning by trading is occurring, but it gives little information about the nature of that learning. In the model of Mahani and Bernhardt (2007), as low-ability investors trade they realize that their inherent level of ability is low and decide to stop trading actively. Thus, examining attrition in the sample is critical to disentangle how much of the estimates of β_1 and β_2 are impacted by the endogenous nature of our sample, or how much of β_1 and β_2 is driven by H_1 and how much is driven by H_2 . To separate our inferences about H_1 and H_2 , we carefully control for investor heterogeneity and survivorship in our next set of tests.

3.4 Impact of attrition on learning estimates

Endogenous attrition can significantly affect our learning estimates, a problem best understood if we consider the decision of the investor to continue trading. To represent an investor who continues to trade when her performance is good and stops once she gets a few bad draws, define

$$s_{it} = I(\alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2 + \delta X_{i,t} + \nu_{i,t} > 0), \quad (4)$$

where $I(\cdot)$ is an indicator variable. Here, s_{it} equals zero if investor i exists in the sample after year t , and one otherwise. That is, s_{it} equals zero only in the last year of investor i 's trading.

The problem with both the simple learning model estimates and the fixed effect estimates is that Equation (3) assumes $\nu_{i,t}$ is uncorrelated with $\epsilon_{i,t}$, which may not be the case.

We first present some overall attrition evidence by examining the rate at which investors who are in our sample (having placed seven round-trip trades) in one year fail to place any trades during the rest of our sample period. Figure 5 shows that attrition is a significant feature of our data. Since the rest of the sample period changes from year to year, the earlier years of our sample period provide more reliable estimates of true exit rates than the later years. Approximately 25 percent of those traders who enter the sample in one year fail to ever trade again. Of traders who trade for two or three years, about five percent permanently exit the sample.

We conduct the Verbeek and Nijman (1992) test to assess whether selection is a problem by including s_{it-1} in the fixed effects model. If selection is not a problem—so investor attrition is random—the coefficient estimate on the selection dummy should be insignificant. The coefficient in the Verbeek and Nijman test estimate on the lagged selection dummy (s_{it-1}) in regression (3) is found to be significant for both performance (coefficient 2.19; t -stat 6.44) and disposition (coefficient 0.10; t -stat 3.30). This again indicates that selection is severe in the sample and attrition is not random.

To examine directly how much ceasing to trade (or learning about ability) affects our inferences about learning, we use a modified version of the selection model introduced by Heckman (1976). While the classic Heckman model involves a two-stage procedure—a selection model in the first stage to predict which observations will be observable in the second stage and the regression of interest with an adjustment for survivorship bias in the second stage—it does not account for individual heterogeneity. The evidence in Section 3.3 suggests that accounting for individual heterogeneity may be important to control for cohort or similar effects. Thus, we modify this procedure to account for both survivorship bias and individual heterogeneity, adopting the empirical strategy of Wooldridge (1995), which modifies the Heckman model to allow for fixed effects. Intuitively,

this approach accounts for survivorship by estimating the selection model every year and including the inverse Mills ratios (the conditional probability that an individual continues to trade) of each selection equation in the learning regression model. Individual time-invariant heterogeneity is accounted for in this method by running the regression in first-differences. More concretely, the learning regression model we estimate is

$$\Delta y_{i,t+1} = \beta \Delta x_{i,t} + \rho_{96} I(t = 96) \lambda_{96} + \dots + \rho_{02} I(t = 02) \lambda_{02} + \epsilon_{i,t}, \quad (5)$$

where $\lambda_{96}, \dots, \lambda_{02}$ are the inverse Mills ratios from cross-sectional probit model (the selection model) in each of the years 1996–2002, and $I(\cdot)$ is an indicator variable. (1995 is omitted to avoid perfect collinearity of the model.) Including these variables in the learning regression accounts for the impact of the selection equation. Note that a joint test of $\rho_t = 0$ for $t = 1996, \dots, 2002$ is a test of whether survivorship bias is a concern (Wooldridge 1995).

The first-stage uses cross-sectional probit regressions to predict whether or not the individual ceases to trade in a given period. The probit regressions include a constant, linear and quadratic experience terms, the number of different stocks the investor trades, the individual’s average return in the previous year, the cross-sectional standard deviation of the individual’s previous-year 30-day return, and a dummy variable that is one if the individual’s average daily marked-to-market total portfolio value is in the top quartile of all investors. As instruments, we use the following variables: (1) a dummy variable for whether an investor inherited shares in the previous calendar year (*Inheritance*); and (2) the proportion of active traders in the zip code of the individual (excluding the individual) in the previous calendar year (*Activeness*). As we will explain below, both these variables are likely to satisfy the necessary exclusion restrictions—that is, they are likely to affect the probability of remaining in the sample, but are unlikely to affect changes in an individual’s performance or disposition effect except through their effect on survival.

For our first instrument, we conjecture that an individual who inherits shares is more likely to trade in the future, perhaps because their wealth has increased, or because the new shares cause them to pay more attention to the stock market. Another explanation, consistent with Jin and Scherbina (2008), is that inherited shares may prompt the recipient to trade since the shares may not fit with the investor's desired asset allocation. This satisfies the exogeneity condition since inheritance of shares from a relative is unlikely to directly affect changes in the performance or disposition effect of an individual. In the data, when an investor dies and shares are transferred to an heir, it appears as a transaction with the account of the deceased selling shares and the account of the heir purchasing shares. A special code identifies the transaction as an inheritance. In our sample death transfers are evenly spread over the sample (ranging from 62 in 1995 to as high as 443 in 2003).

Our second instrument is based on the proportion of accounts in an individual's zip code that place at least seven round trip trades in the previous year. The papers that argue that an individual investor is more likely to trade if his or her neighbors are trading (Hong, Kubik, and Stein 2004, Ivković and Weisbenner 2007) imply that this measure is correlated with investor activity. This instrument satisfies the exclusion restriction since having active traders in one's zip code is unlikely to directly affect changes in performance or disposition effect.¹² There are 1979 zip codes in the data and the measure varies between 0 and 88%.

We construct the sample to be used in the selection model as follows. An account observation is added to the selection sample if it places one or more trades in a given year. This differs from our main sample, where we require investors to have placed at least seven round-trip trades in order to estimate either the average performance or the disposition coefficient. Once an account is added, it remains in the selection sample until 2003, which is the end of our data. In some years, an account will have placed enough round-trip trades to be included in our hazard regressions, so the data will

¹²While these papers acknowledge that there may be other common factors, driving investor activity (such as access to similar financial advisors), these papers claim that there is a social component that is unrelated to performance that drives at least part of investor trading. Our identification relies only on this component.

include a performance average and disposition estimate for this account. However, each year we will also have data on many accounts for which we do not have performance and disposition estimates. If estimates are available, we treat the account as having been selected into our data. Results from the selection model, with two-step efficient estimates of the parameters and standard errors, are given in Table 4. The first-stage selection model uses 36,030 observations, while the second-stage regression (in first differences) use only 11959 observations in the performance regression and 16188 observations in the disposition regressions.

We estimate the first-stage regression for each year and construct inverse Mills ratios for each year. For brevity, we only report one set of pooled first-stage estimates with year fixed effects in Columns 1 to 3 of Table 4. Results for each of the years are qualitatively similar to those reported. The model in Column 1 shows the estimate on \bar{R}_{t-1} is positive and significant. This is consistent with H_1 —as low-ability investors trade, they learn about their inherent ability and cease trading when they get negative returns. More successful investors continue to trade actively. The estimate is also economically meaningful and suggests that, keeping other explanatory variables at their mean levels, a decrease in returns of one standard deviation increases the probability that the individual will cease to trade next period by around 15 percent. The other coefficient estimates reported in the first column of Table 4 also seem sensible: investors are more likely to remain in the sample and trade if they hold relatively diversified portfolios and have relatively more trading experience.

Following the model of Mahani and Bernhardt (2007), under H_2 , it is likely that individuals will cease trading when the variance of the signals they get from the performance of their trades is large, i.e., the signals are noisy. We proxy for the noise in the signals the investors receive from the performance by the variance in the returns across all positions taken by an account in the previous year and include it in the model in Column 2.¹³ We also include measures of ex-ante investor

¹³This measure is not the equivalent to the volatility of stocks purchased by an investor. Rather, it is the variance of the returns generated by each of the trades placed by the investor. Empirically, the correlation between this quantity and the average volatility of stocks held by an investor in our data is only 0.23.

sophistication proxied by whether they trade options or have significant wealth. The idea is that such investors are more likely to continue trading. The coefficient estimates reported on these additional variables in the second column also seem sensible: investors are more likely to remain in the sample if they receive precise signals about their performance and if they are relatively sophisticated.

Finally, in the third column we also add our instruments to the first-stage regression. As is reported, the coefficient estimates for both of our instruments are statistically significant at the five percent level or better, and they have the predicted sign. Specifically, both inheriting shares and having higher trading activity in the zip code of an investor increase the probability that the investor will continue trading in the next period. A joint χ^2_2 test for the significance of the instruments rejects the null hypothesis that the instruments are weak at the one percent level ($\chi^2_2 = 49.10$). In addition, we test the validity of exclusion restrictions related to the instruments using the Likelihood Ratio (LR) test proposed in Wooldridge (2003, Chapter 17).¹⁴ We find that the null hypothesis that exclusion restrictions related to the two instruments are violated is rejected at the 1% level, suggesting again that our instruments are doing a good job in explaining the selection equation. Both the instruments also have an economically significant impact. For instance, keeping other variables at their mean levels, inheriting shares increases the probability that the investor will continue trading in the next period by around eight percent. Similarly, a one standard deviation increase in investor activity (an increase of 0.20) in the zip code of the individual in the previous year increases the probability that the investor will continue trading next period by about five percent.

We find that accounting for selection has a significant impact on our learning estimates. Column 4 uses performance as the dependent variable in a regression of the form of Equation (5). Comparing the estimates in Table 4 to the simple model reported in Table 3, the coefficient on YearsTraded

¹⁴The idea is to examine the fall in the log-likelihood functions for the unrestricted (with instruments) and restricted (without instruments) models and compare with a chi-square distribution with appropriate degrees of freedom, which in our case is two.

is no longer statistically significant, and the coefficient on cumulative trades is reduced by about 90 percent.¹⁵ When we use disposition as the dependent variable in Column 5, the coefficient is reduced by slightly more than 50 percent.

The coefficients on the inverse Mills ratios in Columns 4 and 5 are also sensible. In particular, they suggest that factors that predict which investors stay in the sample are positively correlated with future performance and negatively related with disposition. This suggests that survivorship is indeed important. The endogenous attrition of low-ability investors significantly boosts the experience coefficient estimate in the simple performance regression and it diminishes the estimate in the simple disposition regression. The joint tests of statistical significance of the inverse Mills ratios also show that accounting for sample selection is important. In the disposition regression, the joint F-test of significance of the inverse Mills ratios yields an F-statistic of 12.18, much greater than the one percent critical value of 2.64. Similarly, in the performance regression, the F-statistic of 35.80 is much greater than the one percent critical value. Unreported results in which YearsTraded and CumulTrades are included in separate regression models yield almost the same coefficients, but regressions that include both variables are reported for brevity.

We also explore the properties of the coefficients on the inverse Mills ratios from the disposition and performance regressions in unreported tests. We find a strong negative relation between the coefficients on the inverse Mills ratios from the performance and disposition models. This is quite reasonable, since endogenous attrition of low-ability investors should significantly increase the estimates of experience in the performance regression and should reduce them in the disposition regression. Moreover, regressing the inverse Mills ratios coefficients on a dummy for whether the excess market return (over the risk-free rate) is positive, we find that there is less attrition in

¹⁵The coefficients on YearsTraded in these regressions must be viewed with caution. Because the change in YearsTraded is always equal to exactly one year, the coefficient on YearsTraded can only be identified if we leave the fixed effect for one year out of the regression. For robustness, we also estimate standard Heckman selection models without either individual or year fixed effects. The coefficients on YearsTraded in these models are 0.31 in the performance regression and -0.019 in the disposition regression. Both of these coefficients are marginally statistically significant (at 10 percent). The magnitude and significance of these coefficients is consistent with our other results.

high-return periods than in low-return periods. This is indicated by the coefficient on the return inverse Mills ratios being higher (by about 12.3; $p = 0.07$) and the coefficient on the disposition inverse Mills ratios being lower (by about -0.23; $p = 0.11$) during high-return periods. These results suggest that less learning-about-ability occurs in ‘good’ times.

Our results suggest that accounting for selection is important and significantly affects inferences about learning. Investor heterogeneity and survivorship effects account for something on the order of one-half to three-quarters of the learning estimates found in simple and aggregate models. This translates directly into slower learning than that inferred from simpler models. Taking the disposition-learning coefficient, for example, 100 trades corresponds to an improvement of about 0.04 in the simplest model, an improvement of about 0.03 in the model with individual fixed effects, and an improvement of about 0.02 in the survivorship/fixed effects model. Roughly speaking, if it takes about 100 trades to improve about 4 percent in the simple model, it takes about 200 trades to achieve the same improvement after adjusting for survivorship and individual heterogeneity.

Our estimates suggest that the fraction of learning that is driven by investors learning about their inherent ability is large. After adjusting for this type of learning, the portion of learning that is due to investors learning to improve their ability over time is statistically different from zero, but not excessively large. Overall, our findings are consistent with the two hypotheses we outline in Section 2, and the evidence supporting H_2 is much stronger than the evidence supporting H_1 .

3.5 Additional evidence of learning

3.5.1 Learning about ability: trading intensity

An important implication of H_2 is that investors increase their trading as they become more experienced and their beliefs about their ability become more precise, as in Mahani and Bernhardt (2007). We therefore examine whether investors who decide to continue trading increase their trading intensity. First, in Figure 6, we plot the average number of trades (dark blue) and the

value of shares traded in Euro (light blue) for each year of experience. We demean the results by calendar year to adjust for year fixed-effects, which ensures that market-wide changes in investor behavior do not contaminate our results. The plot clearly indicates that surviving investors place more trades as they gain more experience, both in quantity and in value of shares traded. This is confirmed using regressions, which are reported in Table 6. The regressions use the average number of trades and the value of trades as the dependent variables and they include year and individual fixed effects. In Columns 3 and 4 we also control for the endogenous attrition decision of investors. As is clear from the table, we find that investors trade more intensively as they become more experienced. This analysis does not include those investors for whom we are not able to estimate the disposition effect. These results are not qualitatively changed if we expand our sample to the entire population of investors.¹⁶

3.5.2 Learning by doing: risk and the speed of learning

As we outline in the hypothesis section, H_1 suggests that the productivity of agents increases with experience, for instance, by using better combinations of various information sources that help them choose the stocks to invest in. In order to examine this, we estimate our learning regressions with risk-adjusted returns (or 30-day alphas) instead of raw returns, and we regress the average factor betas of stocks purchased by investors on experience and our control variables. Our regressions also control for survivorship and individual and year fixed effects. The results of our regressions appear in Table 7. The table clearly shows that risk-adjusted returns improve with experience, with a coefficient that is actually larger than the coefficient we estimate for raw returns. Looking at the coefficients on average factor betas makes it clear why this is the case. With more experience,

¹⁶An alternative explanation that might also explain why some traders stop trading is related to the wealth of the investors: investors could stop trading because they run out of money, not because they realize that their ability is low. We do not believe that this alternative explains our results. First, our previous results directly control for initial wealth of the investor and for investor fixed effects. Doing so should partially account for any initial wealth related attrition. Second, in contrast to H_2 , this alternative makes no predictions on how trading intensity of an investor should change over time. Under H_2 , low ability traders should stop trading while high ability traders should scale up their trading intensity. The results in this section also provide support for H_2 rather than this alternative.

investors are actually both improving raw returns and taking less risk, or purchasing stocks with lower factor betas. This result is particularly strong for the market (RMRF) and size (SMB) factor betas. These tests provide additional support for H_2 .

As a final confirmation that agents improve their ability with experience, we examine learning for a group of less sophisticated (low-ability) investors. We conjecture that these ex-ante low-ability investors will learn at a faster rate than their higher ability peers. The model of Mahani and Bernhardt (2007) assumes that investors do not know their ability before they start trading. As a result it has no clear prediction on differential speeds of learning of ex-ante sophisticated and unsophisticated investors. However, if we assume that investors who are ex-ante sophisticated have some notion of their ability, our first hypothesis suggests that since these investors have little to improve, they will either learn slowly or they will not learn at all. To test this we need to sort investors by their level of sophistication. We rely on ex-ante observable characteristics of investors that we believe are related to their financial sophistication. In particular, we consider investors ex-ante likely to be relatively sophisticated if they trade options or have significant wealth.

Each row of Table 5 displays the mean of the disposition coefficient and the average returns (or performance) of each group, the simple regression coefficient of these variables on cumulative trades, estimates corrected for survivorship and unobserved heterogeneity, and the number of observations used in the calculations. Results for disposition are shown in Columns 1, 2 and 3, and for returns in Columns 4, 5 and 6. Looking at the table, it is clear that the means for each of our sophistication subgroups is significantly different, and each change in the mean across subgroups is of the sign we expect. In each pair of rows of the table there is a clear difference between the unsophisticated investors, who learn to avoid the disposition effect at a rate of about 10 percent per year, and sophisticated investors, for whom the learning coefficient is often insignificant. Moreover, the survivorship adjusted estimates show that even after correcting for attrition and investor fixed

effect, these investors do improve their ability to trade over time—that is, we confirm that H_1 is true. Similar results are obtained if we use `YearsTraded` as the experience variable instead.

3.6 Robustness tests

In this section we report the results of a few additional tests that relate to our main predictions. First, in unreported results, we substitute the market return for each stock’s return to see if individuals learn to time the market. If investors are learning to identify good times to buy then the market as a whole will tend to increase after their purchases; if instead they are learning to select stocks, we will not find evidence of learning when we look only at market returns. In fact, we find that the coefficient estimates on experience variables are insignificant, which suggests that performance improves because investors become better at stock selection. Second, we also conduct our tests using the Wooldridge (1995) method, taking data on investors who resume trading after ceasing to trade for a few years (the tests reported in the last section had dropped such investors, using only observations in two consecutive years). Including these investors increases the sample by around 250 observations in the second stage but does not affect the nature of the results reported. Third, all of the results on disposition remain qualitatively unchanged if we include a ‘December dummy’ in (7) or remove partial sales from our sample. This rules out tax-motivated selling or rebalancing as possible explanations for the disposition effect. Finally, in all the fixed effect regressions that control for individual heterogeneity, we cluster the standard errors at the individual level and find that our results are unaffected.

We also perform all of our tests with 30-day returns, regardless of when the investor actually sells. We find that our results are unchanged. This implies that our results are not driven by the selling behavior of investors. It is not the case, for example, that correlated liquidity shocks cause investors to sell simultaneously, drive down prices, and experience poor returns. In additional

tests, we find that an investor's sales are not concentrated in time, which is not consistent with their selling being driven by correlated liquidity shocks.

We also confirm that the learning we find is not driven by increasing precision in our disposition estimates. Although we estimate the disposition effect with disjoint data for each investor-year, since we expect surviving investors to increase their trading over time, it is likely that the precision of our disposition estimates will increase with time. It is, however, important to note that the precision of our estimates is also driven by the volatility of the stocks that investors hold, not just the number of positions they open and close. Moreover, we have no reason to believe that increasing precision will cause a reduction in our disposition estimates and not an increase. Nevertheless, to rule out this possibility, we conduct a bootstrap experiment that keeps constant the number of observations we use to estimate each investor's disposition effect. We assume that an investor places ten trades each year, regardless of the number of trades they actually place. We take a random sample of ten trades from their actual trades (with replacement) and estimate the disposition coefficient as before. We then estimate our survivorship and heterogeneity-adjusted model as before with these new disposition coefficients. Our results are robust to this procedure, and show clearly that our finding of learning is not driven by increasing precision in our estimates of the disposition effect.

4 Conclusion

We examine learning using the complete trading records of individual investors in Finland during the period 1995–2003. We correlate performance and disposition with investor experience and investor survival rates to determine whether and how investors learn by trading. We find that performance improves and the disposition effect declines as investors become more experienced, suggesting that investors learn by trading. Importantly, cumulative trades is a better measure of trading experience than the number of years that an investor has traded; our evidence that years of experience matters is relatively weak. We differentiate between investors learning about their

inherent ability by trading and investors learning to improve their ability over time by accounting for investor attrition. We find that a substantial part of this learning occurs when investors stop trading after learning about their inherent ability rather than continuing to trade and improving their ability over time. That is, the primary way that low-ability investors learn is by learning to stop trading. By not accounting for investor attrition and heterogeneity, the previous literature significantly overestimates how quickly investors become better at trading.

Our results suggest a number of interesting implications. First, since investors who continue trading learn slowly and there is great deal of turnover in the investor population, it is likely that behavioral biases are an important feature of financial markets. Agents do not learn fast enough to make it impossible for biases to affect asset prices. Second, while it would be nice to know how quickly investors who cease to trade would learn if they chose to continue trading, we have no way to estimate this speed. If we assume that those who continue trading learn more quickly than those who cease to trade, policy makers might enhance welfare by devising screening mechanisms, or tests that measure and reveal inherent investing ability. Allowing unskilled investors to learn of their poor ability without incurring significant costs might be more valuable than encouraging people to become active investors. Third, an open question in the literature is why there is such high trading volume, particularly among seemingly uninformed individual investors. Our results indicate that such trading may be rational; investors may be aware that they will learn from experience, and choose to trade in order to learn. Our results also suggest that differences in the expected performance of investors may arise from different experience levels. Finally, if many inexperienced investors begin trading around the same time, their trades could lead to time-varying market efficiency. Our evidence is therefore consistent with the recent results of Greenwood and Nagel (2008), and the more general discussion found in Chancellor (2000) and Shiller (2005).

5 Appendix

This appendix provides details of our approach to calculating individual returns and the disposition effect. It also describes the results of an alternative estimation procedure for the disposition effect that follows Feng and Seasholes (2005).

5.1 Measuring performance

As mentioned in the main text, measuring the performance of individual investors is a significant challenge. Our data do not include all non-equity securities that may be held by an investor, so it is impossible to measure the return for the investor’s entire portfolio. This is made more difficult by the fact that the amount of money an individual has invested in equities often fluctuates significantly over time. Since we cannot accurately measure portfolio returns, we measure performance by examining the average return of stocks purchased. However, this generates a new problem—comparing the returns on holding periods of different lengths. For example, it is particularly difficult to compare the performance of one investor who holds a stock for one week and earns a holding period return of 3 percent to that of another investor who holds a stock for one year and earns a holding period return of 15 percent.

We therefore calculate the returns earned by the purchased stock in the 30 trading days following each of an investor’s purchases. Importantly, we truncate this calculation window at the length of the actual holding if it is shorter than 30 days. That is, the 30-day return for investor i holding stock j is

$$R_{i,j}(t) = \frac{P_j(t + \min(s, 30))}{P_j(t)} - 1, \quad (6)$$

where $P_j(\cdot)$ denotes the stock’s closing price adjusted for splits and dividends, t denotes the purchase date, and s denotes the actual holding period.

Our approach is an attempt to deal with the problem of comparing returns over similar holding periods while ensuring that the actual selling decisions of investors affect their performance. By measuring returns this way, we hope to capture the value of short-term signals that the investor may have received. As a robustness check, we perform all of our tests with 30-day returns, regardless of when the investor actually sells, and find that our results are unchanged. Looking over longer horizons would introduce considerable noise into our return estimates.

5.2 Measuring disposition

Previous researchers have measured the disposition effect in a number of ways. Odean (1998) compares the proportion of losses realized to the proportion of gains realized by a large sample of investors at a discount brokerage firm. Grinblatt and Keloharju (2001) model the decision to sell or hold each stock in an investor’s portfolio by estimating a logit model that includes one observation for each position on each day that an account sells any security. Days in which an account does not trade are dropped from their analysis.

As Feng and Seasholes (2005) point out, a potential problem with these and similar approaches is that they may give incorrect inferences in cases in which capital gains or losses vary over time. Hazard models, which have been extensively applied in a number of fields including labor economics and epidemiology, are ideally suited to our setting. Since our focus is on estimating disposition at an individual level, we estimate the hazard regression for each investor and year. Implementation of the hazard model uses *all* data about the investor’s trading and the stock price path, rather than just data on days when a purchase or sale is made. That is, it implicitly considers the hold-or-sell decision each day. This improves our disposition estimates, and gives us more power with which to investigate learning. More important, individual level estimates enable us to test H_1 and H_2 by explicitly accounting for investor attrition and heterogeneity.

To measure the disposition effect for each investor-year, we use a Cox proportional hazard model with time-varying covariates to model the probability that an investor will sell shares that he currently holds. We count every purchase of a stock as the beginning of a new position, and a position ends on the date the investor first sells part or all of his holdings. Alternative definitions of a holding period, such as first purchase to last sale, or requiring a complete liquidation of a position, do not substantively alter our results. Our time-varying covariates include daily observations of some market-wide variables and daily observations of whether each position corresponds to a capital gain or loss.

Specifically, we estimate

$$h_{i,j}(t|x_{i,j}(t)) = \phi_i(t) \exp\left\{\beta_i^d I(R_{i,j}(t) > 0) + \beta_i^r \bar{R}_{M,t} + \beta_i^s \sigma_{M,t} + \beta_i^V V_{M,t}\right\} \quad (7)$$

where the hazard rate, $h_{i,j}(t)$, is investor i ’s probability of selling position j at time t conditional on not selling the position until time t , and $\phi_i(t)$ is the investor’s baseline hazard. Since we estimate the hazard model for each investor-year, the baseline hazard rate describes the typical holding period of just one investor in one particular year. The Cox proportional hazard model does not impose any structure on the baseline hazard, and Cox’s (1972) partial likelihood approach allows us to estimate the β coefficients without estimating $\phi(t)$.

The time-varying covariates, $x_{i,j}(t)$, are allowed to change each day. These include $I(R_{i,j}(t) > 0)$, an indicator for whether the total return on position j from the time of purchase up until time t is positive. Investors who suffer from the disposition effect are more likely to sell when this condition is true, so they will have positive values of β_i^d . We therefore refer to β_i^d as investor i ’s ‘disposition coefficient.’ The total return variable includes any dividends or other distributions, and is calculated using closing prices on all days, including the date of purchase. We use the closing price on the purchase date instead of the actual purchase price to ensure that our results are not contaminated by microstructure effects. As well, our data do not include transaction prices during the first three months of 1995, but we do have closing prices during this period. (Using ex-dividend returns leaves our results unchanged.) In addition, we include as controls three 5-day moving averages of market-level variables to ensure that we are not capturing selling related to market-wide movements: market returns ($\bar{R}_{M,t}$), squared market returns ($\sigma_{M,t}$), and market

volume ($V_{M,t}$). We repeat this estimation via maximum likelihood each year from 1995–2003 for each individual i who places at least seven round-trip trades in a year.

5.3 Estimates using the Feng and Seasholes method

We estimate the yearly hazard models presented in Table 8, which are comparable to the model of Feng and Seasholes (2005). We use a proportional hazard regression rather than the parametric Weibull model used by Feng and Seasholes. Our indicator variable is the same as their ‘Trading Gain Indicator’ (TGI), and the experience variable we use for this model is the total number of trades placed before the current trade, rather than `CumulTrades`, which is the total number of trades placed in previous calendar years. The estimates are from pooled hazard models estimated each year, in which all individuals are grouped together. Experience is interacted with an indicator for whether the price is above the purchase price, and the coefficient on this interaction term is interpreted as a learning coefficient. These models again give evidence of learning. However, the learning coefficient estimates are quite variable over time (0.149 to 0.063), and they are statistically insignificant in three of the nine years. Furthermore, the average disposition coefficient of an investor is estimated in aggregate to be around 0.65 in this model. This suggests that, depending on the period chosen, an additional year of experience corresponds to a reduction in the disposition coefficient of approximately 10 percent to 25 percent, significantly higher than the estimate from our heterogeneity- and survivorship-adjusted learning model. Table 8 also lists the number of observations available each year, and the fraction of the observations that are censored, or the purchased stocks that are not sold by the end of each year.

The low level of the average disposition effect, the high variability in the annual learning estimates, and the high level of learning found both by Feng and Seasholes (2005) and in our implementation of their model suggest that there are significant differences between our approach and theirs. One important difference is that we estimate a different disposition coefficient for each individual. This allows us to control for each individual’s baseline hazard function (or their average holding period). Another difference is that we estimate learning over a much longer period of time, since Feng and Seasholes (2005) only have about two years of transactions data. Thus, the year-to-year variation in the annual estimates in Table 8 is much less of a concern for our analysis.

References

- Arrow, K., 1962, "The Economic Implications of Learning By Doing," *Review of Economic Studies*, 29, 155–173.
- Barber, B. M., and T. Odean, 2001, "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116, 261–292.
- 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, UC Berkeley.
- Bolton, P., and C. Harris, 1999, "Strategic Experimentation," *Econometrica*, 67, 349–374.
- Calvet, L. E., J. Y. Campbell, and P. Sodini, 2007, "Fight or Flight? Portfolio Rebalancing by Individual Investors," NBER Working Paper.
- Campbell, J. Y., 2006, "Household Finance," *Journal of Finance*, 61, 1553–1604.
- Carhart, M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.
- Chancellor, E., 2000, *Devil Take the Hindmost: A History of Financial Speculation*, Plume, New York, N.Y.
- Choi, J., D. Laibson, B. Madrian, and A. Metrick, 2007, "Reinforcement Learning and Investor Behavior," Working Paper, Yale University.
- Coval, J. D., D. A. Hirshleifer, and T. Shumway, 2005, "Can Individual Investors Beat the Market?," Working Paper, University of Michigan.
- Coval, J. D., and T. Shumway, 2005, "Do Behavioral Biases Affect Prices?," *Journal of Finance*, 60, 1–34.
- Cox, D. R., 1972, "Regression Models and Life Tables," *Journal of the Royal Statistical Society, Series B*, 34, 187–220.
- DeLong, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann, 1991, "The Survival of Noise Traders in Financial Markets," *Journal of Business*, 64, 1–19.
- Fama, E., and K. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.
- Feng, L., and M. S. Seasholes, 2005, "Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?," *Review of Finance*, 9, 305–351.
- Frazzini, A., 2006, "The Disposition Effect and Underreaction to News," *Journal of Finance*, 61, 2017–2046.
- Genesove, D., and C. Mayer, 2001, "Loss-Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics*, 116, 1233–1260.
- Greenwood, R., and S. Nagel, 2008, "Inexperienced Investors and Bubbles," Forthcoming, *Journal of Finance*.

- Grinblatt, M., and M. Keloharju, 2001, "What Makes Investors Trade?," *Journal of Finance*, 56, 589–616.
- Grossman, S. J., R. E. Kihlstrom, and L. J. Mirman, 1977, "A Bayesian Approach to the Production of Information and Learning by Doing," *Review of Economic Studies*, 44, 533–547.
- Heckman, J., 1976, "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models," *Annals of Economic and Social Measurement*, 5, 475–492.
- Hong, H., J. D. Kubik, and J. C. Stein, 2004, "Social Interaction and Stock-Market Participation," *Journal of Finance*, 59, 137–163.
- Ivković, Z., C. Sialm, and S. J. Weisbenner, 2008, "Portfolio Concentration and the Performance of Individual Investors," *Journal of Financial and Quantitative Analysis*, 43.
- Ivković, Z., and S. J. Weisbenner, 2005, "Local Does as Local Is: Information Content of the Geography of Individual Investors Common Stock Investments," *Journal of Finance*, 60, 267–306.
- Ivković, Z. Z., and S. Weisbenner, 2007, "Information Diffusion Effects in Individual Investors' Common Stock Purchases: Covet Thy Neighbors' Investment Choices," *Review of Financial Studies*, 20, 1327–1357.
- Jin, L., and A. Scherbina, 2008, "Inheriting Losers," Working Paper, University of California, Davis.
- Kogan, L., S. A. Ross, J. Wang, and M. M. Westerfield, 2006, "The Price Impact and Survival of Irrational Traders," *Journal of Finance*, 61, 195–229.
- Korniotis, G. M., and A. Kumar, 2008, "Superior Information or a Psychological Bias? A Unified Framework with Cognitive Abilities Resolves Three Puzzles," Working Paper, University of Texas at Austin.
- Linnainmaa, J., 2006, "Learning from Experience," Working Paper, University of Chicago.
- Linnainmaa, J., 2007, "Do Limit Orders Alter Inferences about Investor Performance and Behavior?," Working Paper, University of Chicago.
- List, J. A., 2003, "Does Market Experience Eliminate Market Anomalies?," *Quarterly Journal of Economics*, 118, 41–71.
- Mahani, R., and D. Bernhardt, 2007, "Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity," *Journal of Finance*, 62, 1313–1340.
- Nicolosi, G., L. Peng, and N. Zhu, 2008, "Do Individual Investors Learn from Their Trading Experience?," Forthcoming, *Journal of Financial Markets*.
- Odean, T., 1998, "Are Investors Reluctant to Realize Their Losses?," *Journal of Finance*, 53, 1775–1798.

- Odean, T., 1999, "Do Investors Trade Too Much?," *American Economic Review*, 89, 1279–1298.
- Pastor, L., L. Taylor, and P. Veronesi, 2007, "Entrepreneurial Learning, The IPO Decision, And The Post-Ipo Drop In Firm Profitability," Forthcoming, *Review of Financial Studies*.
- Pastor, L., and P. Veronesi, 2003, "Stock Valuation And Learning About Profitability," *Journal of Finance*, 58, 1749–1789.
- Scholes, M., and J. Williams, 1977, "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5, 309–327.
- Shapira, Z., and I. Venezia, 2001, "Patterns of Behavior of Professionally Managed and Independent Investors," *Journal of Banking and Finance*, 25, 1573–1587.
- Shefrin, H., and M. Statman, 1985, "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 40, 777–790.
- Shiller, R. J., 2005, *Irrational Exuberance (Second Edition)*, Princeton University Press, Princeton, N.J.
- Shumway, T., and G. Wu, 2006, "Does Disposition Drive Momentum?," Working Paper, University of Michigan.
- Verbeek, M., and T. Nijman, 1992, "Testing for Selectivity Bias in Panel Data Models," *International Economic Review*, 33, 681–703.
- Weber, M., and C. Camerer, 1998, "The Disposition Effect in Securities Trading: An Experimental Analysis," *Journal of Economic Behavior and Organization*, 33, 167–184.
- Wooldridge, J. M., 1995, "Selection Corrections for Panel Data Models Under Conditional Mean Independence Assumptions," *Journal of Econometrics*, 68, 115–132.
- Wooldridge, J. M., 2003, *Introductory Econometrics: A Modern Approach*, South-Western College Publishing, Cincinnati, Ohio.

Table 1: Summary Statistics

This table presents summary statistics for our data. Panel A includes all individual accounts in our data that started trading during the sample period. Panel B gives results just for those accounts for which we are able to estimate at least one disposition coefficient. We only estimate the disposition coefficient if an individual has placed at least seven round-trip trades in a given year. Number of trades is the total number of trades placed by an investor during the sample period. Average portfolio value is the average marked-to-market value of an investor's portfolio using daily closing prices. Panel C reports summary statistics for the estimates of the hazard model in equation (7) in the text.

Panel A: Entire Sample (322,454 accounts)

	Mean	25th Pctl	Median	75th Pctl
Number of years with trades	1.9	1.0	1.0	2.0
Number of securities traded	3.5	1.0	1.0	3.0
Number of Trades	15.4	1.0	3.0	8.0
Average value of shares traded, EUR	3,447	808	1,653	3,310
Average portfolio value, EUR	11,588	1,470	2,794	5,856
Age in 1995	39.3	27.0	39.0	51.0
Gender (1=female)	0.39			
Trades options (1=yes)	0.03			

Panel B: Accounts with Disposition Estimates (11,979 accounts)

	Mean	25th Pctl	Median	75th Pctl
Number of years with trades	4.4	3.0	4.0	6.0
Number of securities traded	22.3	12.0	18.0	28.0
Number of trades	222.3	68.0	117.0	224.0
Average value of shares traded, EUR	5,356	1,855	3,235	5,759
Average portfolio value, EUR	58,828	5,102	11,483	26,147
Age in 1995	35.3	27.0	34.0	44.0
Gender (1=female)	0.15			
Trades options (1=yes)	0.17			

Panel C: Hazard Function Estimates

Variable	Mean	<i>t</i> -stat	10th Pctl	Median	90th Pctl
β^d	1.32	6.81	-0.36	1.04	2.57
β^r	-0.04	-0.27	-0.65	0.03	0.93
β^s	0.01	0.46	-0.05	0.00	0.04
β^V	0.25	0.91	-2.62	0.24	3.08

Table 2: Ex Post Return Test

This table shows the results of a regression that effectively implements the same test reported in Odean’s (1998) Table VI. The idea is to compare the returns of stocks sold at a gain to those that could have been sold at a loss but were not. In order to see the difference in cost of the disposition effect to investors in the low- and high-disposition groups, we implement a difference-in-difference specification. Using a regression framework instead of simple averages allows us to include year fixed effects to absorb year-specific factors such as average market performance. We define low- and high-disposition groups as the bottom and top quintile of the disposition coefficient. The negative coefficient in the first row indicates that for the low-disposition group the effect documented by Odean is reversed: stocks sold for a gain subsequently underperform. (Robust standard errors, clustered by investor, are presented in parentheses.) The difference-in-difference coefficient is significantly positive at the three horizons we consider, indicating that the effect documented by Odean (1998) is stronger for individuals in our high-disposition group than in the low-disposition group.

	Ex Post Return Horizon		
	30 days	45 days	60 days
Sold for a gain	-0.0164 (0.0021) ^{***}	-0.0520 (0.0024) ^{***}	-0.0541 (0.0026) ^{***}
High disposition	-0.0120 (0.0052) ^{***}	-0.0191 (0.0069) ^{***}	-0.0145 (0.0084)
Sold for a gain \times High disposition	0.0605 (0.0026) ^{***}	0.1063 (0.0033) ^{***}	0.1146 (0.0040) ^{***}
Year fixed effects	Yes	Yes	Yes

Table 3: Learning Models

This table presents results for regressions of the form

$$y_{i,t+1} = \alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2 + \delta X_{i,t} + \epsilon_{i,t}$$

where the dependent variable is either the investor's average 30-day return following purchases ($\bar{R}_{i,t+1}$) or the investor's disposition coefficient ($\beta_{i,t+1}^d$). In the Simple Learning Model, the α_i are held constant across all investors, while in the model with heterogeneity they are allowed to vary. Experience is measured by either years of experience (YearsTraded) or cumulative number of trades placed (CumulTrades). $X_{i,t}$ is a vector of controls including the number of trades placed by the individual in a given year (NumTrades), the number of securities held by the individual in a given year (NumSec), and the individual's average total daily portfolio value (PortVal). Data are from the period 1995 to 2003. Standard errors are in parentheses, and ***, **, * and * denote significance at 1%, 5% and 10%, respectively.

	Simple Learning Model		Learning Model with Heterogeneity	
	$R_{i,t+1}$	$\beta_{i,t+1}^d$	$R_{i,t+1}$	$\beta_{i,t+1}^d$
YearsTraded _t	0.414 (0.160)***	-0.050 (0.014)***	0.249 (0.401)	-0.010 (0.038)
YearsTraded _t ²	-0.043 (0.027)	0.007 (0.002)***	-0.022 (0.038)	0.01 (0.003)***
CumulTrades _t ($\div 10^2$)	0.110 (0.043)***	-0.041 (0.003)***	0.058 (0.059)***	-0.030 (0.005)***
CumulTrades _t ² ($\div 10^4$)	-0.0005 (0.0003)*	0.0002 (0.00002)***	-0.0009 (0.0004)**	0.0002 (0.00003)***
NumSec _t	0.096 (0.012)***	0.104 (0.012)***	0.003 (0.024)	-0.008 (0.002)***
NumTrades _t	-0.031 (0.005)***	-0.043 (0.007)***	-0.021 (0.010)**	-0.0007 (0.001)
PortVal _t ($\div 10^6$)	0.104 (0.061)*	0.099 (0.058)*	0.618 (0.602)	-0.067 (0.601)
\bar{R}_t		-0.005 (0.001)***	-0.002 (0.001)**	-0.002 (0.001)*
Individual fixed effects	No	No	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	13404	17715	13404	17715
R ² (%)	1.6	1.1	45.7	38.7
		1.6	45.9	38.6

Table 4: Learning with Survival Controls

This table reports estimates of selection model regressions with the fixed effects modification developed by Wooldridge (1995). The regressions are of the form

$$\Delta y_{i,t+1} = \beta \Delta x_{i,t} + \rho_{96} I(t = 96) \lambda_{96} + \dots + \rho_{02} I(t = 02) \lambda_{02} + \epsilon_{i,t},$$

where $\lambda_{96}, \dots, \lambda_{02}$ are the inverse Mills ratios from cross-sectional probit model (the selection model) in each of the years 1996–2003, and $I(\cdot)$ is an indicator variable. Including these variables in the learning regression accounts for the impact of the selection equation. (1995 is omitted to avoid perfect collinearity of the model.) A joint test of $\rho_t = 0$ for $t = 1996, \dots, 2002$ tests of whether survivorship bias is a concern. The first stage probit model is estimated each year and the inverse Mills ratios for each year are computed separately from each of these models. For brevity, the table reports estimates of the first stage probit model with data from all of the years of the sample pooled together. The second-stage regressions are estimated with all the variables in first differences, except the inverse Mills ratios. These first differences add fixed effects to the model. The dependent variable in the second stage is either the individual's average return in the following year, $\bar{R}_{i,t+1}$, or the individual's disposition coefficient, $\beta_{i,t-1}^d$. $X_{i,t}$ is a vector of controls described in the main text. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10% respectively.

Dependent variable:	Selection Models (1st Stage)			2nd Stage	2nd Stage
	Insample $_{i,t+1} = 1$			$\Delta R_{i,t+1}$	$\Delta \beta_{i,t+1}^d$
YearsTraded $_t$			0.292 (0.199)	0.451 (0.521)	-0.038 (0.042)
YearsTraded $_t^2$			0.028 (0.008)***	-0.009 (0.043)	0.007 (0.003)**
CumulTrades $_t$ ($\div 10^2$)			0.341 (0.021)***	0.036 (0.003)***	-0.021 (0.006)***
CumulTrades $_t^2$ ($\div 10^4$)			-0.0011 (0.000)***	-0.001 (0.000)	0.002 (0.000)***
\bar{R}_{t-1}	0.859 (0.064)***	1.040 (0.082)***	1.050 (0.082)***		-0.211 (0.162)
$\sigma(\bar{R}_{t-1})$		-0.160 (0.062)***	-0.181 (0.061)***	-6.539 (0.829)***	0.047 (0.097)
Trades Options		0.257 (0.018)***	0.257 (0.018)***		
Wealthy		0.430 (0.021)***	0.427 (0.022)***		
$I(\text{Inherit} = 1)$			0.134 (0.060)**		
Activeness			0.289 (0.043)***		
<i>Inverse Mills Ratios:</i>					
ρ_{96}				4.868 (0.883)***	-0.128 (0.062)**
ρ_{97}				12.57 (1.130)***	-0.141 (0.074)*
ρ_{98}				24.03 (1.350)***	-0.388 (0.143)***
ρ_{99}				-7.6973 (4.910)	0.291 (0.203)
ρ_{00}				4.417 (0.822)***	-0.146 (0.064)**
ρ_{01}				7.16 (0.636)***	-0.092 (0.050)*
ρ_{02}				12.874 (0.671)***	-0.229 (0.058)***
Observations (N)	36030	36030	36030	11959	16188
Time Fixed Effects	Yes	Yes	Yes		
Other Controls	Yes	Yes	Yes	Yes	Yes
Log likelihood	-26077	-22982	-22958		
R^2 (%)				12.7	0.83
<i>Joint test of $\rho_t = 0$</i>					
$F(7, N)$				35.80	12.18
$\text{Pr} > F$				0.000	0.000

Table 5: Heterogeneity in Learning

This table reports both means and simple learning coefficient estimates from regressions of the form,

$$y_{i,t+1} = \alpha + \beta_1 \text{CumulTrades}_{i,t} + \beta_2 \text{CumulTrades}_{i,t}^2 + \delta X_{i,t} + \gamma_t + \epsilon_{i,t}.$$

The variable of interest is either the disposition coefficient ($\beta_{i,t+1}^d$) or returns ($\bar{R}_{i,t+1}$). Regressions using the simple learning model are labeled Mod.1, while regressions using the adjustment for attrition are labeled Mod.2. Attrition-corrected regressions include the inverse Mills ratios described in Table 3. For brevity we only report β_1 coefficients in the table. We classify investors as ‘Trades options’ if they trade in options at any point during our sample. Similarly, investors are classified as ‘wealthy’ if they are in the top 25th percentile of average portfolio value. We include year dummies in all the regressions. Data are from the period 1995 to 2003. Standard errors are in parentheses and ***, ** and * denote significance at 1%, 5% and 10% respectively. All group means are significantly different at 1% level.

Classification	Dependent Variable								
	Mean	$\beta_{i,t+1}^d$			Obs	Mean	$\bar{R}_{i,t+1}$		
		Mod.1	Mod.2	Obs			Mod.1	Mod.2	Obs
No options trades	1.17	-0.024 (0.003)***	-0.02 (0.001)***	14078	-0.81	0.14 (0.06)***	0.32 (0.13)***	10389	
Trades options	0.99	-0.013 (0.004)***	-0.01 (0.004)***	3846	0.1	0.008 (0.04)	0.01 (0.01)	3020	
Low wealth	1.14	-0.023 (0.003)***	-0.018 (0.001)***	15426	-1.18	0.05 (0.01)***	0.029 (0.01)***	11411	
High wealth	1.11	-0.04 (0.05)	-0.01 (0.012)	2498	1.08	0.01 (0.023)	0.02 (0.021)	1198	

Table 6: Trading Intensity and Experience

This table reports regression results for trading intensity and experience. In Columns one and three the dependent variable is the number of trades placed by an investor in a year. In Columns two and four the dependent variable is the value of shares traded, measured in Euros. Attrition- corrected regressions include the inverse Mills ratios described in Table 3. For brevity, we only report the coefficients on experience variables in this table. All regressions include individual and year fixed effects. Data are from the period 1995 to 2003. Standard errors are in parentheses and ***, ** and * denote significance at 1%, 5% and 10% respectively.

	No Attrition Correction		With Attrition Correction	
	Num Trades	Val Trades	Num Trades	Val Trades
YearsTraded _t	10.074 (0.618)***	135257 (14091)***	13.32 (0.749)***	22713 (5626)***
YearsTraded _t ²	-0.462 (0.097)***	-9515 (2123)***	-0.896 (0.085)***	-1216 (557)**
Observations	13266	13266	13266	13266
Individual fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

Table 7: Risk Taking and Experience

This table reports the results of fixed effect selection model estimates of regressions of various performance and risk measures on experience measures. The method of these regressions and their associated first-stage estimates are described in Table 4. The dependent variables include each investors average 30-day risk adjusted return (alpha), and each investors average beta coefficient on four factors—RMRF, SMB, HML, and UMD. These betas are estimated in the standard way, as described in the text. Each regression includes the control variables and inverse Mills ratios described in Table 4, but only the experience variables coefficients are reported in this table. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10% respectively.

Coefficient	Dependent Variables				
	30-day α	β on RMRF	β on SMB	β on HML	β on UMD
YearsTraded	0.58 (0.41)	-0.034 (.0128)**	-0.016 (.0059)**	0.0544 (.0435)	0.0073 (.0250)
YearsTraded ²	0.02 (0.02)	0.00246 (.0012)*	0.00163 (.0005)***	0.0037 (.0028)	-0.00076 (.0019)
CumulTrades ($\div 10^2$)	0.05 (0.02)***	-0.01 (0.002)***	-0.002 (0.0009)***	-0.002 (0.004)	-0.016 (.015)
CumulTrades ² ($\div 10^4$)	-0.0002 (0.0001)**	0.00005 (0.00001)***	0.00001 (0.000006)**	0.00002 (0.00002)	0.00008 (0.00005)
Dep Var Mean	0.07	0.663	0.028	-0.086	-0.054
Standard Dev	3.30	0.288	0.103	0.431	0.440

Table 8: Disposition Estimates at Aggregate Level Using Alternative Estimation

This table presents learning estimates from pooled proportional hazards models, using a method similar to that of Feng and Seasholes (2005). Each year we pool the trades of all investors, treating them as if they were just one individual and estimating one learning coefficient for the entire population. The model is

$$h(t) = \phi(t) \exp\{\beta_1 I(R_{i,j}(t) > 0) + \beta_2 \text{Exper} + \beta_3 \text{Exper}^2 + \beta_4 \text{Exper}[I(R_{i,j}(t) > 0)] + \beta_5 \text{Exper}^2[I(R_{i,j}(t) > 0)]\},$$

where Exper is measured in years since first placing a trade and $I(R_{i,j}(t) > 0)$ is an indicator variable that takes a value of one when a stock has increased in price since its purchase date. For brevity, only the β_4 coefficient estimates are reported. Estimated β_5 coefficients are all insignificant. We also report the number of trades, or observations, considered by the model (in thousands of trades), the percentage of observations censored (trades not closed by the end of the year), and the number of accounts contributing observations to the model. Standard errors are in parentheses, and ***, ** and * denote significance at 1%, 5% and 10%, respectively.

Year	1995	1996	1997	1998	1999	2000	2001	2002	2003
β_4	-0.074 (0.039)*	-0.093 (0.062)	-0.096 (0.031)***	-0.149 (0.022)***	-0.031 (0.024)	-0.013 (0.011)	-0.063 (0.011)***	-0.069 (0.012)***	-0.123 (0.012)***
Trades	6.6	14.5	26.6	44.6	99.6	251.9	161.9	115.6	131.8
Censored	16%	19%	19%	19%	17%	16%	17%	18%	20%
Accounts	384	713	1360	2412	4953	11585	7445	5416	6056




Figure 1: Participation By Year

This graph shows the number of accounts that place one or more trades in each year, including both accounts that exist at the beginning of the sample and new accounts. There is considerable variation in the number of accounts placing trades over time, from a low of around 54,196 in 1995 to a high of 311,013 in 2000.




Figure 2: Returns Persistence

This figure plots the average 30-day returns earned by investors in years following their first purchase. Investors are grouped into quartiles in their first year of trading, and we then calculate average returns for each quartile in subsequent years. Returns are calculated using the approach discussed in the text, and are demeaned by calendar year, which removes the impact of any year fixed effects. Raw returns are reported here, but the results for risk-adjusted returns are not qualitatively different.

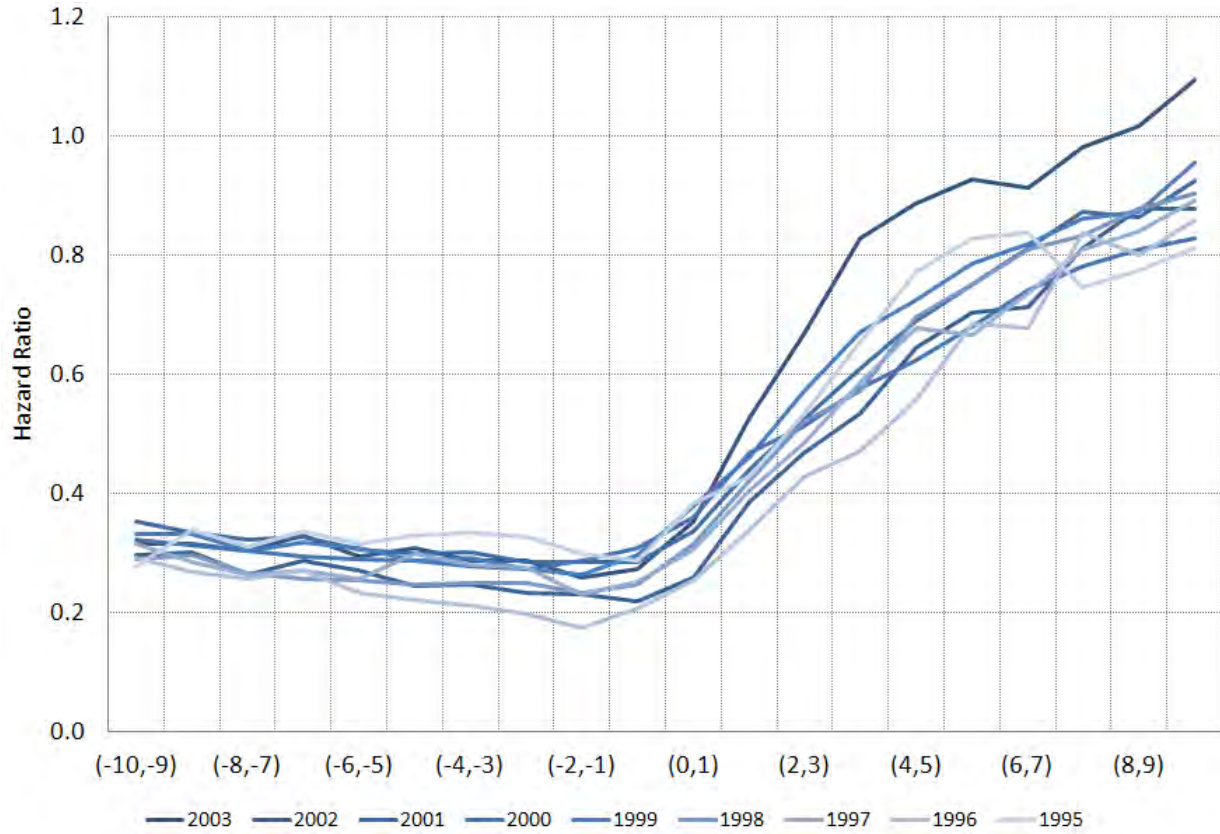


Figure 3: The Disposition Effect in Aggregate

This graph shows how the propensity to sell a stock depends on the stock's return since purchase. Each line plots the regression coefficients from one hazard regression modeling the conditional probability of selling a stock. The coefficients correspond to dummy variables for return 'bins' ranging from $[-10, -9)$ percent to $[9, 10)$ percent. In each year, there is a pronounced kink near zero, and the hazard increases rapidly for positive returns.

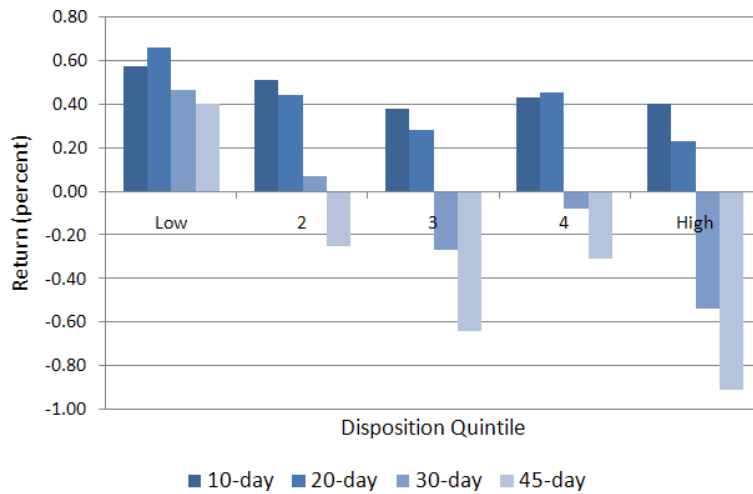


Figure 4: Returns by Disposition Quintile

This figure shows average 10-, 20-, 30-, and 45-day returns following a purchase for each disposition quintile. Returns are calculated using the approach discussed in the text. Raw returns are reported here, but the results for risk-adjusted returns are not qualitatively different. Returns earned by the lowest quintile (1) are higher than those earned by the highest quintile (5).




Figure 5: Proportion of Accounts Who Exit

This graph shows the percentage of investors who exit in 1 year, 2 years, ..., or 6 years following the first year in which they place at least seven trades. Exit is defined as placing no further trades during our sample period. Within each year group, the first bar indicates the percentage of investors who exited in the 1st year after placing a trade, the second bar indicates the percentage of investors who exited in the 2nd year, and so on. Data are missing for later years because we do not know how many investors stopped trading after 2003.

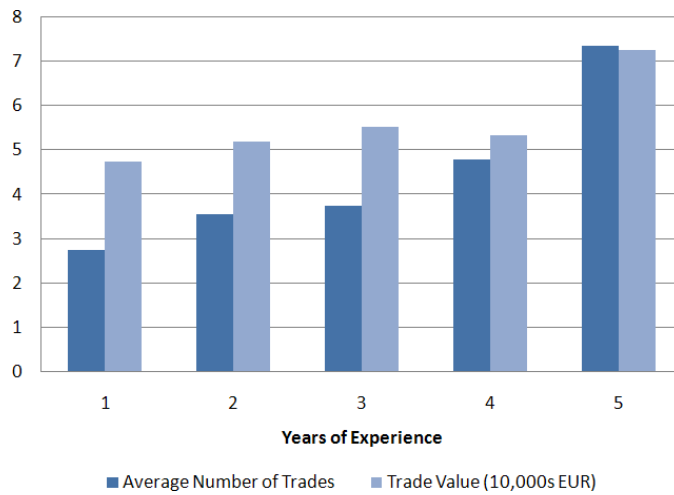


Figure 6: Trading Intensity and Experience

This figure shows how trading intensity changes with experience. Intensity is measured as the number of trades placed (dark blue) and the total value of trades placed (light blue; in 10,000's of EUR). Results are demeaned by calendar year to adjust for year fixed-effects. The results for the first partial year in which investors trade are omitted, so the positive values in the plot are all deviations from that first partial year period. We report results beginning in the investor's first full calendar year in the sample.