Learning By Trading

Amit Seru, Tyler Shumway, and Noah Stoffman*

Stephen M. Ross School of Business
University of Michigan

This version: January 3, 2007
First version: March 15, 2006

Abstract

We test whether investors learn from their trading experience. Using a large sample of individual investor records over a nine-year period, we analyze both the disposition effect and trading performance at the individual level. Disposition is costly: the median investor who suffers from the effect earns 3.2% to 5.7% lower annual returns on average than an investor with no disposition. Disposition falls, and performance improves, as investors become more experienced. An extra year of experience decreases the disposition effect of the median investor by about 4%, which accounts for about 5% of the increase in returns earned by these investors. By controlling for survival and unobserved individual heterogeneity, we show that investors in aggregate learn partly by attrition, but that learning at the individual level is also important. We also find that unsophisticated investors and investors who trade more learn faster, and we show that the trading style of investors changes with experience.

*701 Tappan Street, Ann Arbor, MI, 48109. Seru can be contacted at (734) 764-7197 or `aseru@umich.edu`, Shumway can be reached at (734) 763-4129 or `shumway@umich.edu`, and Stoffman can be contacted at (734) 764-7197 or `stoffman@umich.edu`. 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 partial funding. We thank Mark Seasholes, Ning Zhu, and seminar participants at Carnegie Mellon University, UC Irvine, the University of Manchester, the University of Toronto, and the University of Michigan finance brown bag for helpful comments. Any remaining errors are ours.

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. Considering these findings, it is natural to ask how skilled or informed investors acquire their advantage. In this paper, we test whether individual investors learn by trading. Studying both the investment performance and the strength of a behavioral bias of individual investors, we examine the hypotheses that biases decline and performance improves with investment experience.

Individual investors have been shown to trade too much (Odean 1999, Barber and Odean 2001), hold their employer's stock in their retirement funds (Benartzi 2001), and hold undiversified portfolios (Goetzmann and Kumar 2005). While investors may exhibit many different behavioral biases, our tests focus on the empirical regularity widely known as the disposition effect. The 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). Like several of these papers, we present evidence that the disposition effect is a behavioral bias. We focus on the disposition effect because it is a robust empirical finding and it is relatively easy to measure.

While the typical individual investor achieves relatively poor performance and exhibits behavioral biases, there is growing evidence of cross-sectional dispersion in the information or ability of individuals. Coval, Hirshleifer, and Shumway (2005) document significant performance persistence among individuals. Ivkovich and Weisbenner (2005) find that individuals place more informed trades in stocks of companies located close to their homes, and Ivkovich, Sialm, and Weisbenner (2005) show that individuals with more concentrated portfolios tend to outperform those who are more diversified. Linnainmaa (2005b) finds that individuals who trade with limit orders suffer particularly poor performance. This literature suggests that the market is not perfectly efficient, making it possible to ask whether some of the cross-section of ability we observe is due to learning.

Since there are various ways in which investors might learn in financial markets, we need to be clear about the type of learning we hope to measure. There is a large literature about market participants learning the values of the parameters that describe their investment opportunity sets (e.g. Lewellen and Shanken (2002)). The learning that we hope to measure is a much broader concept than this sort of learning. If investors change their behavior with experience, in any way that leads to improved investment performance or to reduced behavioral bias, we consider this change to be learning by trading. Thus, while the learning that we consider might incorporate making inferences about important parameters, it is not constrained to parameter estimation in the context of a particular model.

As an example of the type of learning we hope to measure, consider the problem faced by an investor trying to decide which of the myriad sources of market information and investment advice to take seriously. Investors are free to update their beliefs based on standard news sources, internet sites, investment newsletters, and neighbors or friends. They may also consider the advice of brokers, news analysts, authors of books and magazines, and finance professors. To the extent that these sources fail to completely agree, individuals must determine how much decision weight to assign to each source. If historical data on the performance of these sources are unavailable, individuals will have to learn about each source's value by observing their performance as they trade. We hypothesize that investors are more capable of identifying successful strategies as they gain experience. That is, they 'learn by doing' (Grossman, Kihlstrom, and Mirman 1977).

As another example of the type of learning we hope to measure, consider the problem faced by a new investor who may be subject to behavioral biases. Before observing her own reaction to profits and losses, news events, or volatility in market prices, she may not know the extent to which her responses to these stimuli will be biased. With the benefit of hindsight, she may be able to identify biases she has exhibited in the past and avoid those biases in the future. We conjecture that new investors learn to avoid their own behavioral biases as they become more experienced.

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. Investors may also learn by observing market quantities and considering the outcomes of hypothetical trades based on, for example, a particular information source. By estimating performance improvement and bias reduction as a function of both years of trading and transactions executed, we can estimate the relative magnitudes of both types of learning.

Our learning hypotheses, which we present in the next section, are important for at least four reasons. First, knowing whether investors learn by trading helps us understand the nature of the investment problem. In a standard neoclassical setting, we should not find evidence of learning by trading among individual investors. In such a setting investors have complete access to public information and unlimited cognitive ability, so they can back-test all possible investment strategies before they ever trade. Since these investors will use all available information to optimize their trading strategies, their strategies will not significantly improve with experience. Thus, if we find evidence of significant learning, this is also evidence for either costly information or cognitive constraints.

A second motivation for our empirical work is to better understand learning in an economic context. While there is a great deal of theoretical literature in both finance and economics about learning (as discussed by Sobel (2000)), direct empirical evidence about learning by economic agents is still relatively uncommon. Since we essentially estimate learning curves[1] for investors as a function of both time and cumulative trading activity, we can ask whether investors learn by actually trading or simply from the passage of time. And since we measure learning both as a reduction in the disposition effect and as any unspecified changes in behavior that lead to better performance, we can also test whether learning in these different dimensions is correlated. Finally, by examining the attrition of investors from our sample, we can differentiate between a population that learns because its unsophisticated members leave, and a population that learns because its members learn. We do this by examining the exit of individuals from our sample and implementing a Heckman selection model to control for survival.

A third reason to study learning concerns the possibility of important time-variation in the degree of market efficiency. If the population of traders changes significantly over time, and if newer traders are particularly subject to behavioral biases, periods in which many new investors are trading may correspond to periods in which prices do not reflect fundamental values. For example, Grinblatt and Han (2005) argue that trading by investors

[1] Learning curves are discussed in Argote (1999).

with the disposition effect causes stock price momentum. If the momentum effect varies significantly with time, prices might deviate from fundamentals substantially. This popular explanation of the 'technology bubble' in the late nineties has been argued by Shiller (2005) and Greenwood and Nagel (2006), and more generally by Chancellor (2000), among others.

Our final reason for considering learning in the context of the disposition effect is to more fully understand the nature of this effect. While there is substantial evidence that the disposition effect is a behavioral bias, it is possible that the effect could be explained by informed trading, rebalancing, or transactions costs (Strobl 2003). If the disposition effect is a behavioral bias and if investors learn with experience, we would expect investors to display less disposition with experience. Thus, confirming that more experienced investors display less disposition helps us differentiate the bias explanation for disposition from other explanations.

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 2.9 times as 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 by trading, and how this learning affects their returns. In particular, we estimate the disposition effect for each account and year in our sample and relate these estimates to experience, returns, and various demographic controls.

Our results provide robust evidence of learning. An additional year of trading experience decreases the disposition of a median investor by about 4%. Moreover, an additional year of experience improves performance by 40 basis points (bp) when returns are measured over 30-day horizons. We also find that disposition is costly, since investors earn higher returns when they do not suffer from the bias. The reduction in disposition that comes with one year of trading experience explains 4–6% of the increase in returns earned by these investors. Our results are unaffected if we control for unobserved investor heterogeneity (such as innate ability) and survival.

In addition to this main finding, we examine particular subgroups of investors whose disposition estimates decline. We expect that learning should be concentrated among unsophisticated investors and those who start out earning consistently poor returns. This is consistent with the notion that investors who perform poorly at the beginning of their investing careers and those who are unskilled might make larger improvements as they gain experience. We find support for this in our sample. Moreover, we find that learning primarily takes place when the market as a whole is down, which is consistent with investors learning particularly when their performance is not confounded with a market in which most stocks' prices are rising.

Our results contribute to a growing literature about learning by market participants. Barber, Odean, and Strahilevitz (2004) investigates the purchases of investors who have previously owned and subsequently sold a stock. They find that investors repurchase stocks that they previously sold for a gain, and stocks that have lost value subsequent to their prior sale. That is, investors repeat decisions that have been profitable in the past while avoiding those that have not, which the authors argue is a naïve form of learning.

Our tests are more closely related to those of Feng and Seasholes (2005), which documents that investors, in aggregate, display less disposition over time. Feng and Seasholes perform their tests with a hazard model, using Chinese data. They estimate the model with the trading records of all individuals together, rather than estimating the model for individual investors, as we do. Nicolosi, Peng, and Zhu (2004) shows that the trading performance of individuals improves with trading experience, estimating simple regressions with data from a large discount U.S. brokerage firm. Linnainmaa (2005a) examines the learning behavior of day traders using data from Finland that are similar to ours.

While our tests have some features in common with each of the papers above, they differ from the literature in a number of important respects. First, unlike any of these papers, our tests use estimates of the disposition effect that are specific to *individuals*, allowing us to track the disposition of particular individuals over time. This allows us to use a fixed effects specification, ensuring that our results are not driven by some unobserved individual characteristic such as intelligence or innate ability. Second, having access to individual-specific disposition estimates allows us to control for survivorship biases, as we describe below. This allows us to differentiate between two kinds of learning that are possible for the representative agent: investors can learn, or investors with poor performance can stop trading, which is a kind of aggregate learning. We find evidence for both types of learning, but our evidence suggests that learning at the individual level remains important after controlling for attrition. Third, we examine both investment performance and disposition together in a longer and larger dataset than any of the papers discussed above. Fourth, estimating the impact of learning on both performance and disposition allows us to test whether these two types of learning are correlated. Fifth, we show that the trading style of investors changes with experience, which provides further evidence of learning. 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 describes the hypotheses we test and some of our statistical methods, while Section 2 provides detail on our data. Section 3 explains our methodology and discusses our results. Section 4 concludes.

1 Hypotheses and Methods

To determine whether individual investors learn by trading, we test a number of related hypotheses. This section motivates and describes our hypotheses in more detail. It also describes some of the methods of our statistical tests.

1.1 Measuring disposition

The most direct way to test our hypotheses about the disposition effect require estimating the extent to which individuals in our data exhibit the effect. 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 each position on each day that an account sells any security as an observation. 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.

We estimate the disposition effect with a Cox proportional hazard model with time-varying covariates. Our time-varying covariates include daily observations on some market-wide variables (5-day moving averages of market return, market return squared, and market volume) and daily observations of whether each position corresponds to a capital loss or gain. One advantage of our method is that the hazard model, which directly models the stock holding period, implicitly considers the selling versus holding decision each day. Another advantage is that we can estimate our model for each account with sufficient trades in our dataset.

Hazard models have been extensively applied in labor economics and elsewhere. Proportional hazard models make the assumption that the hazard rate, $\lambda(t)$, or the probability of liquidation at time t conditional on being held until time t is

$$\lambda(t) = \phi(t) \exp(x(t)'\beta), \tag{1}$$

where $\phi(t)$ is referred to as the 'baseline' hazard rate and the term $\exp(x(t)'\beta)$ allows the expected holding time to vary across accounts and positions according to their covariates, $x(t)$. The baseline hazard rate is common to all the trades in the sample. 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. In this model the covariates may vary with time, and as mentioned above, each of our covariates changes daily. The Cox proportional hazard model does not impose any structure on the baseline hazard, $\phi(t)$. Cox's (1972) partial likelihood estimator provides a way of estimating β without requiring estimates of $\phi(t)$. It can also handle censoring of observations, which is one of the features of our data. Details about estimating the proportional hazard model can be found in Cox and Oakes (1984).

1.2 Hypotheses about disposition

For investors to have an incentive to learn to avoid the disposition effect, it must be a behavioral bias that is costly to them. One necessary condition for disposition to be a behavioral bias is that disposition is a somewhat stable, predictable attribute of a particular investor. We examine this feature of disposition by testing our first hypothesis,

Hypothesis 1. *There is persistent cross-sectional variation in the degree of the disposition effect among individual investors.*

We test this hypothesis 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 neglected in the model estimated with second period data. We test Hypothesis 1 by estimating the rank correlation of account-level disposition coefficients over the two periods, and by testing whether the rank correlation is significantly different from zero.

A second necessary condition for disposition to be a behavioral bias is that investors with more disposition have inferior investment performance. If disposition is unrelated to inferior investment performance, investors with the effect would have little incentive to learn to avoid it. We test whether disposition is costly with Hypothesis 2,

Hypothesis 2. *Investors with high disposition effect coefficients have relatively poor investment performance.*

We test this hypothesis by sorting investors into disposition quintiles based on their coefficients estimated in one year and then examining stock returns by disposition quintile in the next year. We calculate post-purchase returns over a number of different holding periods, ranging from 10 to 45 trading days. We also regress returns on indicator variables that are only defined for statistically significant disposition coefficient estimates.

1.3 Hypotheses about learning

The focus of our paper is learning by individual investors, and we test several learning hypotheses. We first look for evidence of learning in investment performance. Specifically, we test our third hypothesis,

Hypothesis 3. *Investors with more experience have relatively good investment performance.*

We test this hypothesis by regressing investors' average returns on measures of investor experience. Returns are measured over the 30 trading days following each purchase. Our primary experience variables include the number of years that an investor has been in our data and the cumulative number of trades the investor has placed. We include a quadratic term for each experience variable to allow investors to learn more slowly over time. We also control for investor heterogeneity by including individual and year fixed effects. Finally, we carefully control for survivorship bias using a procedure introduced by Heckman (1976). It is important for us to control for survivorship bias, since it is clear that investors with weaker performance may be less likely to continue trading long enough for us to estimate their performance in future periods.

In addition to testing for learning in investment performance, we examine the extent to which disposition attenuates with experience. We exploit our estimated disposition coefficients for each investor in each year of our sample to test our fourth hypothesis,

Hypothesis 4. *Investors with more experience exhibit less of the disposition effect.*

We test this hypothesis by regressing investors' disposition coefficients on measures of investor experience. As in our performance results, we use the number of years that an investor has been in our data, and/or the cumulative number of trades to measure experience. We also cluster standard errors by investor and estimate fixed effects models, and we control for survivorship bias using the same Heckman (1976) procedure. Comparing the results of our tests of Hypotheses 3 and 4 allows us to estimate what fraction of any improvement in performance might be associated with reduction in the disposition effect.

Our next hypothesis concerns the rate at which investors learn. If some investors have more significant behavioral biases than other investors, or significantly worse investment performance than others, it is natural that they will learn to avoid biases and improve performance faster than other investors. Specifically, we test the conjecture that,

Hypothesis 5. *Relatively unsophisticated investors learn faster than relatively sophisticated investors.*

We test this hypothesis by sorting investors into subsamples based on various characteristics that are likely to be related to their financial sophistication. For example, we sort investors by their wealth (proxied by each investor's average daily portfolio value), by whether or not they trade options, by past returns, and by several other characteristics. We then use each subsample to regress disposition on experience and other variables in essentially the same regressions we performed to test Hypothesis 4. Finally, we look at the experience coefficients in these subsample regressions to test whether learning across groups occurs at the same rate.

We also test the hypothesis that learning about behavioral biases is correlated with learning about trading strategies or styles that affect performance. Following from Hypotheses 4 and 5, if investors with significant disposition also have poor performance, and if investors with poor performance or stronger disposition learn faster than others, we should be able to show that reducing the disposition effect is associated with improving performance. Specifically, we test our next hypothesis,

Hypothesis 6. *The change in an investor's disposition coefficient is correlated with the change in that investor's performance.*

To test this hypothesis, we again sort investors into subsamples based on various characteristics that are likely to be related to their financial sophistication, including their years of investing experience. We then calculate both a disposition coefficient and an average performance for each group, aggregating the trades for all members of the group as if they were one individual. Finally, we regress the change in each group's disposition coefficient over one year on the change in each group's performance.

Our final hypothesis explores whether investors change their behavior in a measurable way as they learn. If we cannot observe any changes in behavior over time, it is difficult to believe that investors are truly learning. Thus, our final hypothesis is,

Hypothesis 7. *The trading behavior of more experienced investors is measurably different from that of newer investors.*

We test this hypothesis by examining the characteristics of the stocks that are traded by investors with low or high levels of experience, and test for significant differences in the characteristics' means. The characteristics we consider include market size, past returns, past volatility, and volume.

2 Data

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 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.[2] The data include the transactions of nearly 1.3 million individuals and firms, beginning in January, 1995 and ending in December, 2003.[3] In all, more than 22 million trades by individual investors are included.

While our dataset includes exchange-traded options and certain irregular equity securities, we focus on trading in ordinary shares. Trading in Finland is conducted on the Helsinki Stock Exchange, which is owned by OMX, an operator of stock exchanges in Nordic and Baltic countries. Trading on the Helsinki exchange begins with an opening call from 9:45–10:00 a.m., and ends with a closing call from 6:20–6:30 p.m. Continuous trading during regular hours is conducted through a limit order book.

Our transaction data include the number of shares bought or sold, corresponding transaction prices, and the trade and settlement dates, although trades are not time-stamped. Additional demographic data, such as the account-holder's age, zip code, and language are also included. In addition, we create proxies for wealth and a measure of investor sophistication.

To construct a wealth proxy, we use opening balances and subsequent trades to reconstruct the total portfolio holdings of each account on a daily basis. Using these holdings, we measure wealth as 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 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.

Despite the impressive richness of these data, they are imperfect. Only the direct holdings and transactions of individuals are available. This means that for an individual who

[2] These references provide a detailed discussion of the data.

[3] The data include all transactions that settled on or after January 1, 1995. Since settlement in Finland is generally $T + 3$, transactions in the last few days of December, 1994 are included in the dataset, as well as some trades with longer settlement times that took place earlier that month.

directly trades shares of Nokia and holds a Finnish mutual fund that owns shares of Nokia, we will observe only trades in the former. The trades of the mutual fund are included in the dataset, but are identified as holdings of the mutual fund company, and cannot be tied to the individual. However, our wealth calculations allow us to compare the importance of the individual investors as a group to that of other market participants. On average, individuals hold 12.6% of all equity held by Finnish investors, including financial institutions,[4] government funds, nonprofit organizations and nonfinancial corporations. This is more than financial institutions, which hold an average of 9.6% during our sample period. The majority of equity is held by the government (34.7%) and nonfinancial firms (33.4%), although these investors trade relatively less and may do so for strategic reasons that are not directly linked to profit-maximization.

As discussed below, we use survival analysis to investigate disposition at the account level. The enormous size of our dataset gives us a great deal of power with which to investigate learning. However, the computational requirements to undertake this analysis are considerable. Allowing for time-varying covariates requires each observation to have an entire history of price changes, which means several thousand variables for each observation. For example, to estimate the disposition effect for just one investor who holds three stocks over a nine-year period, we need approximately 6,750 data points.

Table 1 provides summary statistics for our dataset. Panel A includes all observations, while Panel B gives results just for those observations for which we are able to estimate a disposition coefficient.[5] 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 likelihood function does not always converge. The means in this table are taken over all observations, so most individuals are counted more than once. The effect of this is to put more weight on values for individuals who appear in the sample repeatedly.

The last three rows of each panel are indicator variables, taking a value of one if the

[4] A complication of our data is that trading of shares held in an American Depository Receipt (ADR) or by certain foreigners who need not register directly with the CSD is hidden in the orders of certain institutions that serve as registrars. It is possible, however, to separate these 'nominee' accounts from other institutional holdings by carefully analyzing the trading history of each institutional account. We implement such a procedure, the details of which are available upon request. Therefore, throughout this paper when we write 'institutions,' we mean Finnish institutions and not nominee-registered accounts or trading in American- or Swedish-listed depository receipts.

[5] The observations in Panel A are used in the estimation of a selection model to control for survivorship, which is discussed in Section 3.4.

investor: (a) is present at the beginning of the sample; (b) trades options; or (c) is female; and zero otherwise. Approximately 51% of the entire sample, and 49% of the subset with disposition estimates, are observations from individuals whose accounts were opened prior to the beginning of our sample period. Because we cannot determine how long the accounts have been open, these observations are left-censored; we therefore allow them to have a different learning coefficient in our regressions below.

Comparing Panels A and B, it is apparent that the subset of investors for whom disposition coefficients are available is qualitatively similar to the entire sample, with the exception of the number of trades placed per year (which is higher in Panel B by construction), and the number of distinct securities traded. As well, investors for whom we can estimate disposition are more likely to trade options (23.5%) than the overall sample (18.7%). 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.

3 Results

We present our empirical findings in this section. Each finding is related to the hypotheses laid out above, so we deal with each hypothesis in turn. In Section 3.1 we show that the disposition effect is widespread and economically important in our data. Section 3.2 provides evidence that the disposition effect is costly in the sense that investors who suffer from it earn lower returns than those who do not. We then show that investors learn to avoid the disposition effect as they become more experienced in Section 3.3. We confirm that these results are not driven by a survivorship bias in Section 3.4. In Section 3.5 we investigate learning among subsamples of our data, which gives us instruments to use in constructing group-level disposition estimates for additional tests in Section 3.6. Finally, we examine how learning is manifested in the trading styles of investors in Section 3.7.

3.1 Disposition estimates

To measure the disposition effect, we wish to estimate the probability that an investor sells any stock that they hold at a given point in time. In particular, we want to know how this probability is affected by the stock price path since the initial purchase date. We measure holding periods as the time from the first purchase of a stock by an investor, i, to the time of the first sale. The next purchase of that same stock begins another holding period.[6] Many purchases are not followed by a sale within our sample period, so holding periods are right-censored. We use a Cox proportional hazards regression, and estimate

$$\lambda_i(t; x_i) = \phi_i(t) \times \exp\left(\beta_d I_{\{p_t > p_b\}} + \beta_r \bar{R}_{m,t} + \beta_s \sigma_{m,t} + \beta_V V_{m,t}\right) \tag{2}$$

for each investor. Here, $I_{\{p_t > p_b\}}$ is an indicator whose value is one if the price of a stock on date t is greater than its purchase price, and zero otherwise. Investors who suffer from the disposition effect will have positive values of β_d, which we therefore call the 'disposition coefficient.' We include 5-day moving averages of three controls: market volume ($V_{m,t}$), market returns ($\bar{R}_{m,t}$), and squared market returns ($\sigma_{m,t}$) to ensure that we are not capturing selling related to market-wide movements. The time-varying baseline hazard rate for each investor is denoted by $\phi_i(t)$. We repeat this estimation via maximum likelihood each year from 1995–2003.[7]

Hazard models are a natural framework with which to estimate the disposition effect, but most studies have used a simple logit setting.[8] 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, as has been done with logit models. This makes our disposition estimates more precise, and gives us more power with which to investigate learning.

Before turning to our individual disposition estimates, we present in Figure 1 a graph of the relation between the propensity to sell (hazard ratio) and holding period return. To generate this graph, we group all investors and run one regression each year. Rather than using only one indicator variable as in equation (2), we use 20 dummy variables corresponding to different 1% return 'bins.' (We group the data for this procedure so we can estimate a regression with many regressors. All of the tests that follow are based on individual-level

[6] Alternative definitions of a holding period, such as first purchase to last sale, or requiring a complete liquidation of a position, do not change our results.

[7] All of the results that follow remain qualitatively unchanged if we include a 'December dummy' in (2) or remove partial sales from our sample. This rules out tax-motivated selling or rebalancing as possible explanations for the disposition effect.

[8] Feng and Seasholes (2005) is an exception, although they pool their data and estimate the hazard regression only once. Since our focus is on estimating disposition at an individual level, we estimate the hazard regression for each investor and year.

results.) The graph 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 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. Even with this condition, some of our disposition coefficients are estimated with considerable noise; we therefore use weighted least squares (WLS) estimation of the models discussed below, where the weights are proportional to the reciprocal of the estimated variance of the disposition coefficients.

Table 2 summarizes the distribution of our disposition estimates, which we use to investigate our first hypothesis. Panel A provides information on all investors for whom we have estimates. There are 35,009 observations in our panel, comprised of 20,929 unique accounts. This distribution is the first hint that we will have relatively few data points when we include individual fixed effects in our regressions, as we discuss below. The effect of a high-flying market is apparent in the number of observations each year, which rises considerably and then declines somewhat in the latter part of our sample.

The median disposition coefficient is 1.07, and it increases over time, which could be a result of new disposition-prone investors entering the sample; we consider the effects of such selection in Section 3.4. The rank correlation between an investor's disposition coefficient in year t and their coefficient in year $t-1$ is 0.364, suggesting that there is indeed a fair degree of persistence in the disposition coefficients. This correlation is extremely statistically significant, which provides strong evidence in favor of Hypothesis 1. As a robustness check, we estimate the proportion of individuals who remain in the same quintile in years t and $t+1$. Averaging across years, we find a high proportion (73%) of individuals stay in the same quintile. This again provides strong support for Hypothesis 1.

Using the estimated standard errors for each investor, we can classify estimates as significant or not at any given confidence level. The last two columns of Panel A show the proportion of investors who have a significantly positive or negative disposition coefficient at the 10% level. Over our entire sample period, 43.5% of investors have a disposition coefficient that is statistically greater than zero. Panel B shows results for only those investors whose disposition coefficients are significant (either positive or negative). For this subset of investors, the median disposition coefficient is 1.82. Panel C gives summary statistics for the other coefficients in the hazard model. None of the controls is statistically significant in the cross-section.

These results provide strong evidence that the disposition effect is widespread and economically important in each year of our study. An investor with the median disposition coefficient is $e^{1.07} = 2.9$ 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.

Many of the results that follow focus on explaining these estimated disposition coefficients. Our goal is to understand whether disposition decreases, and performance improves, for individuals over time, and whether observable investor characteristics are associated with the decline.

3.2 Returns and disposition

We turn now to our second hypothesis, that the disposition effect is a costly behavioral bias. In particular, we expect investors who are prone to selling only their winners to experience lower returns than investors who do not behave this way. We analyze this by calculating the returns of stocks over several horizons, ranging from 10 days to 45 days. For each purchase in our dataset, we calculate the return to a stock held for 10, 20, 30, or 45 trading days after the purchase.[9] The h-day returns are calculated as the return earned either over the actual holding period or h days, whichever is shorter. For example, if an investor holds a position for 25 days, the 20-day return will be the return earned over the first 20 days, while the 30-day return is the return earned over the actual 25-day holding period. The median holding period in our sample is 39 days, which suggests that the 30- to 45-day returns most closely reflect the investors' realized returns.[10] We calculate returns using closing prices on both the day of purchase and the day of sale to ensure that our results are not affected by

[9] All of our results continue to hold if we use 60- or 90-day returns.

[10] Our data give us some ability to calculate total portfolio returns at the investor level. However, we are hampered by the fact that investors can deposit or withdraw funds from their equity accounts and we don't observe returns on bonds, real estate, or other investments. We therefore opt for using returns calculated this way, which has the added benefit that it is consistent with the returns used in the hazard regression in (2) to determine $I_{\{p_t > p_b\}}$.

the bid-ask spread.

To get a sense of how returns vary with disposition, we first examine average investor returns across quintiles of the disposition coefficient. The quintiles are calculated using all disposition estimates pooled together. For each quintile, Figure 3 graphs the average return earned by investors over different horizons from the purchase date. Returns are 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.

Hypothesis 2 is explored in more detail in Table 3. The table displays the results of regressions of returns on various measures of the disposition effect and year dummies. Each panel shows the results of five separate models with the explanatory variable being the variable identified in the first column. The regression is

$$R_{i,t}^{h} = \alpha + \beta_x X_{i,t-1} + \gamma_t + \epsilon_{i,t}, \tag{3}$$

where $h = \{10, 20, 30, 45\}$ denotes the horizon over which returns are calculated. Some of the regressions, labeled 'Significant at $x\%$', use as the regressor a dummy variable that takes the value of 1 (-1) if the disposition coefficient is statistically greater (less) than zero at the 10%, 5%, or 1% levels, respectively, and zero otherwise. Note that we regress returns from year t on disposition estimates from year $t-1$, so these are out-of-sample tests in an important sense.

Consistent with Hypothesis 2, the coefficients on disposition are generally significant, and always negative. Focusing on the 30-day returns (Panel C), we see that a one-standard deviation decrease in disposition (=1.6) leads to a 58 bp increase in returns, or roughly 4.8% per year. Put another way, Model 2 indicates that moving down one quintile increases 30-day returns by 18 bp, or approximately 1.5% per year.[11] Looking at the coefficients on the 'significant' dummy variables, we see that the stronger is our evidence that an investor suffers

[11] This calculation assumes 250 trading days per year. This number is likely overstated, as it is unlikely that returns could be scaled up over a year. We present this calculation simply for comparison purposes and for ease of interpretation.

from the disposition effect, the worse are their returns. Investors who have a coefficient that is large enough and measured with sufficient precision to be statistically greater than zero at the 1% level generate 68 bp less in a 30-day period than those who do not. This corresponds to an annual loss of roughly 5.7%.

Another way to understand the economic magnitude of these results is to consider the returns earned by an investor with median disposition. Recall from Table 2 that the median disposition coefficient is 1.07. Combining this with the disposition estimates, we have

$$\frac{1.07 \times (-0.36)}{100} = -0.0039, \tag{4}$$

so an investor with median disposition earns 39 bp less in a 30-day period than an investor with no disposition. This corresponds to a value of -3.2% per annum. Values obtained from the other panels in the table range from -1.9% to -2.4%.

3.3 Learning

Our primary objective in this paper is to determine whether more experienced investors are more likely to avoid the disposition effect and have better investment performance. To do this, we first test Hypothesis 3 by regressing average returns on our experience variables, the number of years since the investor first placed a trade in our sample (labeled 'YearsTraded') and the cumulative number of trades placed by the investor (labeled 'CumulTrades'). These experience variables are the focus of our analysis on learning. Having established that investor performance improves with experience, we next test our fourth hypothesis, that disposition declines with experience.

For both of these hypotheses, our regressions use panel data, where the dependent variable is (a) the investor's average return in a given year; or (b) the estimated disposition coefficients from equation (2). Specifically, we estimate the regression

$$y_{i,t} = \alpha_i + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2_{i,t} + \delta X_{i,t} + \gamma_t + \epsilon_{i,t}. \tag{5}$$

If investors' returns increase with experience, then we should find $\beta_1 > 0$ when we use an investor's average return as the dependent variable. Results examining the relationship between returns and investor experience are reported in Table 4. Similarly, if investors learn to avoid the disposition effect over time—so disposition falls with experience—then we would expect $\beta_1 < 0$ in our regressions where the dependent variable is the estimated disposition of the individual. Results examining relationship between disposition and investor experience are reported in Table 5. Moreover, to capture the fact that investors might learn faster during earlier years, we include a quadratic term in experience to allow for concavity in learning. The prediction here is that $\beta_2 < 0$ in the case of returns, and $\beta_2 > 0$ in the case of disposition. In all tables, we report results for 30-day returns, although the results are the same for the other horizons we considered.

Other controls, $X_{i,t}$, in various specifications include 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 average daily portfolio value (PortVal), as well as year dummies, γ_t. Importantly, we also include a dummy variable that equals one only when an account was opened prior to the beginning of our sample to mitigate the problems of left-truncation bias. This term ('bgn') is interacted with the experience and experience-squared term to allow accounts that have been active for an indeterminate length of time to have an experience-returns or experience-disposition relationship that differs from other accounts.[12] Although it may be instructive to include lagged returns in the performance regression, we omit this variable to avoid the well-known bias in coefficient estimates of lagged dependent variables in a dynamic panel.

The base returns regressions are displayed in Columns 1 and 3 in Table 4. The results indicate that experience is associated with higher returns, both when measured by number of years or cumulative number of trades. While we discuss the economic interpretation of these coefficients shortly, it is worth noting that a possible explanation for the results in this section is that unobserved investor heterogeneity could explain investor learning. For instance, it is possible that an omitted variable such as investor intelligence, or access to insiders, could be related to both returns and experience, or disposition and experience. To rule this out, we exploit our time series of returns and disposition estimates, and adopt a fixed effects specification to control for unobserved heterogeneity at the account level. Our results are unchanged if we use a random effects specification instead, but we opt for the

[12] We also estimated the regressions with bgn interacted with all other variables and find that the coefficient estimates on the variables of interest (YearsTraded, CumulTrades) are similar to those reported. Moreover, the other variables remain insignificant when interacted with bgn. Alternatively, we re-estimated the regressions only for individuals with bgn=0. Though the number of observations drops, we still find that the estimates on all the variables (including YearsTraded and CumulTrades) are qualitatively similar.

fixed effects specification to allow for arbitrary correlation between α_i and $x_{i,t}$.

Overall, the returns regressions displayed in Table 4 provide uniform evidence that experience is associated with higher returns, which confirms the prediction of Hypothesis 3. For example, Column 1 in Table 4 indicates that an investor with one year of experience will earn 40 bp more than an inexperienced investor over a 30-day horizon. Column 3 indicates that a similar increase in returns comes from an additional 15 trades. For comparison, the median number of trades per year for this sample is 23 (Panel B of Table 1). Economically, this suggests that investors learn from both years of experience and actually placing trades; both are important for learning, as shown by the results in Column 5. Moreover, the coefficients on Experience^2 are negative (although generally insignificant), suggesting some slowing of learning over time. The estimates on the two interaction terms with bgn indicate that an additional year of experience has a smaller effect on accounts that have been open for some indeterminate period of time, which is what we would expect if learning slows with time.

Next, we examine the relationship between investor experience and disposition. Table 5 presents our results for the disposition regressions. To reduce the weight given to disposition coefficients that are not estimated very precisely, we estimate the regressions using weighted least squares (WLS), where the weights are proportional to $1/\widehat{\text{Var}(\beta_d)}$ from our hazard regression in equation (2). The base case (Column 1) shows that disposition declines with experience ($\beta_1 < 0$). Moreover, investors tend to slow down in their learning as they gain experience since $\beta_2 > 0$. In Column 2, we add the same controls used in Table 4 as well as the individual's average return from the previous year, $\bar{R}_{t-1}$. There is again evidence, although somewhat weak in this case, that accounts that were opened before our sample begins learn slower, since the signs on the two interaction terms go in the opposite direction from the main effect. Frequent traders, investors who trade more securities, and investors who earned higher returns in the previous year all have lower levels of disposition, but even with these controls our base results are qualitatively unchanged.

The results in Columns 3 and 4 mirror those in 1 and 2, although the measure of experience here is the cumulative number of trades investor i has placed up to time t. Column 3 indicates that an additional 100 trades reduces the disposition coefficient by 0.108, which is about twice the coefficient on Experience in Column 1. In other words, a year of experience or 50 trades have approximately the same effect on disposition. Interestingly, when we include both measures of experience in Column 5, we find that both cumulative trades and years of experience contribute to the change in disposition; neither type of experience drives out the other.

In each of the specifications the estimated YearsTraded and CumulTrades coefficients are statistically significant at 1% level. Economically, however, our results suggest that investors learn relatively slowly. Specifically, the estimates in Column 2 suggest that an additional year of experience corresponds to a reduction in disposition coefficient of approximately 0.04. To provide some context for this estimate, note that the unconditional median disposition coefficient in our sample is 1.07. An individual with this coefficient will be $e^{1.07} = 2.9$ times as likely to sell a stock whose price has risen since purchase than one whose price has fallen. An extra year of experience decreases this by about 4%.

It is also worth noting at this point that the effective number of observations used in the fixed effects estimation is considerably smaller than our initial sample. Since fixed effects estimation requires at least two data points for each investor, we are unable to include in this analysis individuals for whom we can only calculate the disposition effect in one year. This leaves us with 8,370 unique accounts and 24,955 data points, for an average of just under three years of data per account. This is of particular importance in considering the learning curve implied by the coefficients in Column 1, which is graphed in Figure 2. Disposition declines with the first several years of experience, reaching the maximum impact at between five and six years of experience. But as noted, we have very little data with more than six years of experience (from Table 1, the median years of experience is 2 and the 75th percentile is 4), so the right half of the graph should be interpreted cautiously. The curvature in this function is identified mainly from convexity in the impact of learning in the first couple years of experience, where we have the most data. Nevertheless, we view it as a strength of our results that they remain economically and statistically significant even when only the relatively few observations available to be used in the fixed effects specification are examined.

We end this section with an interesting calculation that illuminates the economic importance of the learning found in our data, along the lines of the calculation in (4). We have shown in this section that an additional year of experience reduces the disposition coefficient by 0.04. When combined with the results from Section 3.2, we can approximate the effect that this reduction in disposition has on returns:

$$-0.04 \times -0.0036 \times \frac{250}{30} = 0.0012, \tag{6}$$

where we continue to use 30-day returns as a benchmark. Moreover, the experience regression in Panel C of Table 2 indicates that an additional year of experience is associated with an increase of $0.0032 \times 250/30 = 0.027$ in returns. Therefore, approximately $0.0012/0.027 = 4.5\%$ of the increase in returns attributable to additional trading experience is due to the decrease in disposition. While investors are clearly learning to avoid their behavioral biases, a substantial proportion of their performance improvement must be due to other kinds of learning.

In summary, the results in this section confirm Hypotheses 3 and 4: the disposition effect declines with experience and performance improves with experience. In the next section, we confirm that these results hold even after controlling for a potential survivorship bias.

3.4 Controlling for survivorship bias

A major concern with our results is that survivorship bias might give an appearance of learning by investors when there is in fact no learning at the individual level. To understand how this might be the case, consider the following example under the null hypothesis of no individual learning. There are two types of investors: high-disposition and low-disposition. The high-disposition investors have poor performance, and decide to stop trading, thus exiting from our sample soon after entering. The low-disposition investors do well, and continue trading, thus remaining in our sample. If the data were generated in this setting, we would find evidence of learning, since the disposition effect would decline with experience—even though there is no learning by any investor. This is, however, a different type of learning; the representative agent in this setting is 'learning' by attrition. Our analysis therefore allows us to differentiate between these two types of learning and to document that both occur. We do this by augmenting our sample with investors who do not have disposition coefficient estimates, and implementing a Heckman (1976) selection model to control for these survivorship issues.

We construct the sample to be used in the selection model as follows. An account observation is added to the sample if it places one or more trades in a given year. This differs from our sample above, where we required investors to have placed at least seven round-trip trades in order to estimate the disposition coefficient. Once an account is added, it remains in our sample until the end, in 2003. In some years, an account will have placed enough round-trip trades to be included in our hazard regressions, so the data will include a disposition estimate for this account. However, each year we will also have data on many accounts for which we do not have disposition estimates. If a disposition estimate is available, we treat the account as having been selected into our data. The model of data we observe is

$$E(y_{i,t} \mid y_{i,t} \text{ is observed}) = x'_{i,t}\beta + \beta_\lambda \lambda_{i,t}, \tag{7}$$

where λ denotes the Inverse Mills Ratio. Including the Inverse Mills Ratio adjusts the coefficient estimates to account for the underlying selection model. Our first-stage regression includes a constant, linear and quadratic experience terms (using YearsTraded and CumulTrades), the number of securities traded (NumSec), a dummy for whether the account was open prior to our sample period (bgn), the individual's average daily marked-to-market portfolio value (PortVal), and, as instruments, the individual's average return in the previous year ($\bar{R}_{t-1}$), the standard deviation of the individual's previous-year return ($\sigma_{R_{t-1}}$). (Recall that we calculate returns for every trade an investor places, so in each year we have a number of return observations. $\bar{R}_t$ is the mean return, and σ_{R_t} is the standard deviation of returns.)

Results from the selection model—with two-step efficient estimates of the parameters and standard errors—are given in Table 6. The selection model uses 100,139 observations, while the second-stage regressions use only the 35,009 or 27,160 observations that were used in the regressions for returns and disposition, respectively, above.

The first-stage estimates seem sensible: investors with higher previous returns are more likely to remain in the sample, as are investors who hold relatively diversified portfolios. Higher variability in past performance also increases the probability of survival, suggesting that investors keep trading if at least some of their trades generate large returns, or perhaps that better average performance is also associated with less consistency. It is perhaps surprising that individuals with high portfolio values are more likely to leave the sample, although this could be partly explained by the death of older investors who have accumulated more wealth. The negative relation between survival and years of experience in the first column is caused mechanically by the right truncation of our data, which means that investors can't have more than eight years of experience. (In the fourth column, the coefficient on the cumulative number of trades is positive, but this becomes negative after only nine trades due to the quadratic term.) The selection model fits reasonably well, with a first-stage pseudo-R^2 of 0.265.

It is important to account for selection in this setting, as demonstrated by the highly significant estimate of ρ. (A test of $\rho > 0$ is equivalent to the test of $\beta_\lambda > 0$.) Even after accounting for selection, however, our earlier results with respect to experience continue to hold. The magnitude and significance of the coefficients for both disposition and returns in the second-stage are in fact very close to what we found in our earlier analysis. Thus even though survivorship bias is a legitimate concern, it does not affect the nature of our results. Put another way, while the representative agent learns from the attrition of some investors with poor performance, an important part of learning comes at the individual level.

3.5 Who learns?

Our next set of tests examine whether the patterns we find are consistent with how a Bayesian investor is expected to learn by participating in the market. Specifically, we investigate our fifth hypothesis that unsophisticated investors learn faster than sophisticated investors. Our empirical strategy is to re-estimate the experience regression for disposition and returns conditioning on a number of different variables. We hypothesize below that these classification variables are related to learning, and our results confirm this intuition. Importantly, we do not classify investors on the basis of the estimated disposition coefficient, β_d, because of concerns about measurement error. That is, the most extreme disposition estimates are likely those with the most error, and we would therefore expect these accounts to see a decrease in disposition in future years, even if these investors are not really learning. To avoid sorting on measurement error, we focus instead on observable variables that are related to learning. This approach is similar to an instrumental variables approach.

Each panel of Table 7 displays regression coefficients on Experience and Experience^2, as well as the mean of the dependent variable and number of observations (N). Results for disposition are shown in Columns 1–3, and for returns in Columns 4–6. All the results reported in this section are robust to controlling for survivorship bias and the other alternative specifications described above. However, in the interest of brevity, we do not report the estimates on all the coefficients and alternative specifications. The results of these specifications are available on request.

First, if what we find is in fact investors learning to avoid a behavioral bias, then, *ceteris paribus*, we would expect this effect to be confined to unsophisticated investors. Following this observation we classify investors ex-ante as sophisticated if they trade options (*Options*=1). We subsequently re-estimate equation (5) for both of these groups. As mentioned above, we expect $\beta_1 < 0$ and $\beta_2 > 0$ for investors who do not trade in options while for investors who trade in options we do not expect this pattern. Our results reported in Panel A of Table 7 indicate that this is indeed the case. There is a clear difference between the unsophisticated investors, who learn to avoid the disposition effect at a rate of about 10% per year, and sophisticated investors, for whom the learning coefficient is insignificant.

An alternative measure of sophistication could be an investor's wealth level. As discussed in Section 2, we measure wealth by calculating an end-of-day portfolio value for each investor and day using closing prices. The mean of this value serves as our wealth proxy. We classify an investor as 'wealthy' if they are in the top 25th percentile, and 'not wealthy' otherwise. This cutoff is obviously somewhat arbitrary, but our results are unaffected by using alternative thresholds. The correlation between this variable and the options variable is 0.349, suggesting that each variable captures different information. Nevertheless, the results in Panel B confirm the results in Panel A. In fact, the coefficients in the first column of Panel B are almost identical to those in Panel A. Therefore, our two measures of sophistication provide remarkably consistent evidence that learning takes place among *unsophisticated* investors.

Second, the work of Barber and Odean (2001) suggests that, *ceteris paribus*, males are more overconfident than females in their trading decisions. Consequently, we should find that males are more averse to learn from their mistakes than females, and therefore should learn slower than females. In Panel C we find that this is indeed the case. Remarkably, females learn to shed their disposition bias almost twice as fast as their male counterparts.

Third, it is plausible that investors learn more when the market in general is not doing well. During periods of high market returns, investors' incentives to learn about their biases could be reduced if they attribute their success to their ability, similar to the behavior modeled in Zingales and Dyck (2002) in the context of media and bubbles. Thus, we should find that investors are more likely to learn when the markets are not doing well rather than when they are. To test this we define the state of the market as an 'up-market' if the excess return on a broad Finnish index is positive in a given year, and as a 'down-market' if the excess return is negative. In Panel D we re-estimate our base regressions for each of the two states of the market and find that, as we hypothesized, individuals learn to avoid the disposition effect primarily when the markets are not doing well. Interestingly this result might also shed some light on strategic synchronized trading by informed investors as in Abreu and Brunnermeier (2002). If individual investors are reluctant to learn during boom times, it might explain why informed investors appear to take a long time before they decide to arbitrage away any price divergence during boom markets.

Our last test is related to the issue that motivated us to examine individual learning. Given the evidence in Coval, Hirshleifer, and Shumway (2005), who find that some individuals earn consistently high returns, we posit that learning should be confined to investors who start have poor performance when they first enter our sample. In other words, investors who start out consistently making profits should not be primarily responsible for explaining the disposition-experience relationship, since these investors may attribute their success to their own ability and therefore not learn, as in the theoretical model of Gervais and Odean (2001). To examine this issue, we classify investors who make excess profits that are in the top 75th percentile of the entire market in the first two years of their trading as 'winners.'[13] In Panel E we re-estimate our base regressions for the two groups and find that it is, in fact, primarily the individuals who are *not* winners that learn to avoid the disposition effect.

Interestingly, in general, the group that shows learning in these tests is the group that has a higher average disposition coefficient (e.g., females have a higher mean disposition than males (1.24 vs. 1.11), and females learn faster). This is consistent with high-disposition investors 'needing' to learn. However, these means do not account for survivorship, and must therefore be interpreted cautiously.

Our results in this section suggest that it is unsophisticated investors and investors who start out with poor returns who learn most. Males learn slower than females, again suggesting that males might be overconfident. Finally our results also indicate that most of the learning takes place when the market as a whole is not doing well, which we argue is when investors are more likely to believe that they are making mistakes, and take corrective actions. Taken together, these results provide strong evidence in favor of Hypothesis 5.

[13] Our results are not sensitive to alternative definitions of winners, such as using a one- or three-year classification period, or above-median excess returns.

3.6 Correlated learning

Thus far, we have estimated the disposition effect only for those investors with at least seven round-trip trades in a year. In this section, we implement an alternative procedure that allows us to include even those investors who trade infrequently. Importantly, the procedure also gives us estimates of returns that are less noisy, and we can therefore estimate in a straightforward way how the reduction in the disposition effect is related to investor performance, which allows us to directly test Hypothesis 6. We call this joint improvement of returns and a reduction in the disposition effect 'correlated learning.'

The results in the last section indicate that the disposition effect differs across groups of investors. We take advantage of those differences and use them as 'instruments' to create investor groups that are likely to have cross-sectional dispersion in both disposition and returns.[14] We then use the group-level disposition estimates and returns to estimate directly the impact that a reduction in disposition has on improved investor performance. For example, Panel A of Table 7 shows that the average disposition coefficient for investors who trade options is 0.99, whereas it is 1.17 for those who do not. Grouping investors by whether they trade options will give us the cross-sectional variation needed for a more powerful test.

We group investors along five dimensions: years of experience (0–8), whether the investor ever trades options (binary), whether the account was open prior to 1995 (binary), time-weighted average portfolio value (quintiles), and total number of trades placed (quintiles). Therefore there are $9 \times 2 \times 2 \times 5 \times 5 = 900$ groupings, although many of these have too few observations, so we end up with 411 groups.

We aggregate all of the trades for all investors in each group and estimate the disposition effect as above. We also calculate the average returns earned in the 20-, 30-, and 40-day periods subsequent to each purchase. Because we are aggregating the trades of many investors, this procedure allows us to include trades of individuals who have too few trades for a disposition effect to be estimated on their own.

Kernel density estimates for the distributions of disposition coefficients for these groups as well as for individuals (reported in Table 2) are graphed in Figure 4. The distributions are clearly centered at about the same location, but the density for the group estimates is

[14]This strategy is analogous to forming portfolios based on size or book-to-market to make asset pricing tests more powerful.

considerably less disperse.

We take advantage of the more precise estimates by examining the relationship between changes in disposition and changes in returns. The high level of noise in returns makes this test difficult to implement with the individual-level disposition estimates. For each group, we calculate the year-to-year change in disposition as $\Delta\beta_d = \beta_{d,t} - \beta_{d,t-1}$ and the year-to-year change in average returns as $\Delta\bar{r}_t^h - \bar{r}_{t-1}^h$ for horizons $h = 20, 30, 45$. Summary statistics are report in Panel A of Table 8. Changes in disposition and 20-day returns are insignificant, but returns at 30- and 45-day horizons are significantly positive in the cross-section.

To examine how a reduction in disposition affects returns, we estimate the regression

$$\Delta\bar{r}_t^h = a + b\Delta\hat{\beta}_d + \epsilon_t,$$

where $\bar{r}_t^h$ denotes the average return at a horizon of $h = 20, 30, 45$ days for each group in year t. For each return horizon, there is a significantly negative relationship between changes in returns and changes in β: decreases in disposition are associated with increases in performance, providing support for Hypothesis 6. Moreover, the decrease in disposition explains about 6% of the increase in returns, which is pretty close to our back-of-the-envelope calculation at the end of Section 3.3.

3.7 Changes in trading style

At this point, we have shown that while the disposition effect is costly, investors learn to avoid the bias with experience, and earn higher returns. The improved performance can come directly from a reduction of the disposition effect, or indirectly from concurrent changes in the investor's trading strategy.[15] Therefore, in this section we examine Hypothesis 7 by investigating what sort of changes in investors portfolios and trading occur as they become more experienced. We examine the size, liquidity, past returns, and past volatility of stocks traded by investors over time. We also examine the level of diversification of each investor's portfolio. We find that each of these characteristics changes as investors become more experienced. This provides additional support to the claim that investors learn through

[15] We attempted to differentiate between these direct and indirect effects, but unfortunately the considerable noise in returns prevented us from finding robust evidence of different effects. This is an interesting open question that we plan to explore in future research.

time; if we could find no measurable ways in which behavior was changing, it would be difficult to argue that investors were learning.

The variables of interest in this section are:

Diversification. We track the trading of investors over time to construct their portfolio holdings at the end of each day. We count the number of stocks held, and take a time-weighted average as a measure of that investor's level of diversification. We also count how many stocks from different 2-digit SIC codes are held.

Size. For each trade placed by an investor, we record the market capitalization of the stock using the closing price and shares outstanding from the trade date. We calculate an average across all stocks traded within a year for each investor to get an estimate of the size of stocks they trade.

Past returns. For each trade placed by an investor, we record the past return of the stock, measured over the period $t-25$ to $t-5$ trading days prior to the trade.

Past volatility. We calculate the volatility of the stock over the same period as the past returns.

These variables are calculated each year for each investor account. We then run regressions of the form

$$y_{i,t} = \alpha_i + \beta_1 \text{ Experience} + \beta_2 \text{ Experience}^2_{i,t} + \gamma_t + \epsilon_{i,t}, \tag{8}$$

where the α_i denote individual fixed effects and the γ_t are year dummies. Experience is measured by the number of years since the investor first placed a trade, but the results remain qualitatively unchanged if we use the cumulative trades variable instead. The quadratic experience term is insignificant and is therefore unreported for brevity. The inclusion of year dummies ensures that market-wide changes in stock characteristics will not contaminate our results, and the individual fixed effects control for any time-invariant heterogeneity.

The results presented in Table 9 provide support for Hypothesis 7. Investors hold more diversified portfolios as they gain experience, both in terms of number of stocks held, and number of stocks held from different industries. These results are strongly significant. The size regressions in the first row of Panels B and C indicate that more experienced investors trade larger stocks, both in purchases and sales. Turning to volume, there is a pronounced shift from more liquid stocks (stocks with high trading volume) to less liquid stocks. Investors are more likely to be momentum traders with experience: they buy stocks with high past returns, but there is no significant change in their sales. Finally, stocks sold by experienced investors are significantly less volatile than those sold by inexperienced investors; again, there is no significant effect among purchases.

The results in this section provide support to our main results that investors learn. Notably, we have considered only a few observable ways in which the investors can change their trading behavior. Of course, even if we found no evidence of changes in the particular trading styles considered in this section, it could still be possible that investors are learning; they could, for example, learn to time the market (which would be consistent with a reduction in the disposition effect), or pick different stocks with similar characteristics that nevertheless earn higher returns.

4 Conclusion

We examine learning in a large sample of individual investors in Finland during the period 1995–2003. We focus on both trading performance and the disposition effect, a behavioral bias that can be readily estimated from trading data. We use a hazard regression framework to estimate the effect at the investor level for each year in our sample. Consistent with earlier studies, we find that the disposition effect is widespread and economically important in our data. It is also costly: investors who suffer from the effect earn lower average returns than those who don't. In particular, the median investor who suffers from the effect earns 3.2% lower annual returns on average than an investor who does not suffer from the disposition effect.

Our major finding is that the disposition effect declines, and performance improves, as investors become more experienced; that is, investors learn. An extra year of experience decreases the disposition of a median investor by about 4%. Moreover, an additional year of trading experience is associated with an improvement in average returns of approximately 40 bp over a 30-day horizon. This is likely due to investors learning in a number of ways, but we are able to attribute approximately 5% of this directly to the disposition effect.

Importantly, we differentiate between two types of learning at the aggregate level: learning by attrition, and learning at the individual level. By controlling for survivorship of accounts over time, we show that both types of learning are important. Moreover, our results continue to hold after controlling for unobserved investor heterogeneity, such as innate ability.

In addition, we find that learning is particularly strong among specific groups of investors, including unsophisticated investors, investors who start out earning consistently poor returns, and females. Investors also learn faster during general market downturns. Finally, we show that the trading style of individuals changes as they become more experienced, suggesting that they learn along many dimensions.

Our results suggest a number of interesting policy implications. For example, 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. Moreover, the occasional entrance of many inexperienced investors could lead to time-varying market efficiency. Our evidence is therefore consistent with the recent results of Greenwood and Nagel (2006), and the more general discussion found in Chancellor (2000) and Shiller (2005).

Our research contributes to the growing literature on the investment behavior and performance of individual investors, which was highlighted by Campbell (2006) in his recent presidential address. While the extant literature has provided much evidence that individual investors on average make poor choices, we show that individuals do learn by trading.

References

Abreu, D., and M. K. Brunnermeier, 2002, "Synchronization Risk and Delayed Arbitrage," *Journal of Financial Economics*, 66, 341–360.

Argote, L., 1999, *Organizational Learning: Creating, Retaining and Transferring Knowledge*, Kluwer, Norwell, MA.

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.

Benartzi, S., 2001, "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock," *Journal of Finance*, 56, 1747–1764.

Campbell, J. Y., 2006, "Household Finance," *Journal of Finance*, 61, 1553–1604.

Chancellor, E., 2000, *Devil Take the Hindmost: A History of Financial Speculation*, Plume, New York, N.Y.

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., and D. Oakes, 1984, *Analysis of Survival Data*, Chapman and Hall, London; New York.

Cox, D. R., 1972, "Regression Models and Life Tables," *Journal of the Royal Statistical Society, Series B*, 34, 187–220.

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.

Gervais, S., and T. Odean, 2001, "Learning to Be Overconfident," *Review of Financial Studies*, 14, 1–27.

Goetzmann, W. N., and A. Kumar, 2005, "Why Do Individual Investors Hold Under-Diversified Portfolios?," Working Paper, Yale University.

Greenwood, R., and S. Nagel, 2006, "Inexperienced Investors and Bubbles," Working Paper, Stanford University.

Grinblatt, M., and B. Han, 2005, "Prospect Theory, Mental Accounting, and Momentum," *Journal of Financial Economics*, 78, 311–339.

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.

Ivkovich, Z., C. Sialm, and S. J. Weisbenner, 2005, "Portfolio Concentration and the Performance of Individual Investors," NBER Working Paper No. W10675.

Ivkovich, 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.

Lewellen, J., and J. Shanken, 2002, "Learning, Asset-Pricing Tests, and Market Efficiency," *Journal of Finance*, 57, 1113–1145.

Linnainmaa, J., 2005a, “Learning from Experience,” Working Paper, UCLA.

Linnainmaa, J., 2005b, “The Limit Order Effect,” Working Paper, UCLA.

Nicolosi, G., L. Peng, and N. Zhu, 2004, “Do Individual Investors Learn from Their Trading Experience?,” Yale ICF Working Paper No. 03-32.

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.

Shapira, Z., and I. Venezia, 2001, “Patterns of Behavior of Professionally Managed and Independent Investors,” *Journal of Banking and Finance*, 25, 1573–1587.

Sheather, S. J., and M. C. Jones, 1991, “A Reliable Data-Based Bandwidth Selection Method for Kernel Density Estimation,” *Journal of the Royal Statistical Society, Series B*, 53, 683–690.

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.

Sobel, J., 2000, “Economists’ Models of Learning,” *Journal of Economic Theory*, 94, 241–261.

Strobl, G., 2003, “Information Asymmetry, Price Momentum, and the Disposition Effect,” Working Paper, University of North Carolina.

Weber, M., and C. Camerer, 1998, “The Disposition Effect in Securities Trading: An Experimental Analysis,” *Journal of Economic Behavior and Organization*, 33, 167–184.

Zingales, L., and A. Dyck, 2002, “The Bubble and the Media,” in Peter K. Cornelius, and Bruce Kogut (eds.), *Corporate Governance and Capital Flows in a Global Economy*, chap. 4, pp. 83–104, Oxford University Press, New York.

Table 1: Summary Statistics

This table presents summary statistics for our data. Panel A includes all observations, while Panel B gives results just for those observations for which we are able to estimate a disposition coefficient. We only estimate the disposition coefficient if an individual has placed at least seven round-trip trades in a given year. Even with this restriction, the likelihood function does not always converge. All of the observations in Panel A are included in the estimation of the Heckman selection model, discussed in Section 3.4. The means in this table are taken over all observations in the panel, so most individuals are counted more than once.

Panel A: Entire sample (N=121,477[a])

	Mean	25th Pctl	Median	75th Pctl
Years of experience	2.8	1	2	4
Cumulative number of trades	22.6	3	11	26
Age	44.2	34	44	54
Number of trades per year	13.1	1	7	17
Number of securities traded	9.1	3	7	13
Average portfolio value, EUR	107,631	7,761	19,111	50,351
Average value of shares traded, EUR	6,242.2	1,661	3,214	6,283
Present at beginning of sample[b]	0.509			
Trades options[b]	0.187			
Gender[c]	0.145			

Panel B: Observations with disposition estimates (N=35,009)

	Mean	25th Pctl	Median	75th Pctl
Years of experience	2.8	1	2	4
Cumulative number of trades	48.2	14	27	56
Age	44.6	35	44	54
Number of trades per year	32.1	16	23	37
Number of securities traded	15.7	9	14	20
Average Portfolio Value, EUR	112,773	8,914	21,727	56,645
Average value of shares traded, EUR	6,857.1	2,277	4,048	7,470
Present at beginning of sample[b]	0.485			
Trades options[b]	0.235			
Gender[c]	0.135			

[a] The number of observations available in the selection model in Section 3.4 is reduced to 100,139 because of missing values for returns.

[b] Dummy variable: 0=no, 1=yes

[c] Dummy variable: 0=male, 1=female

Table 2: Disposition Estimates

This table reports statistics for our estimates of the disposition effect. β_d is the coefficient in the hazard regression (equation (2)). The regression is estimated for each account-year with ten or more round-trip trades. The cross-sectional median and 1st/10th deciles are reported. The columns labeled 'positive' and 'negative' report the proportion of investors with a statistically significant coefficient, where significance is measured at the 10% level using standard errors obtained from the maximum likelihood estimation of the hazard model. Panel B reports statistics only for those accounts with a significant (positive or negative) coefficient estimate. Panel C provides the same information for the controls in the hazard model, where β_r, β_s, and β_V are the coefficients on 5-day moving averages of market returns, market returns squared, and market volume, respectively.

Panel A: Entire sample

		β_d estimate			Significant at 10%	
Year	N Obs	10th Pctl	Median	90th Pctl	Positive	Negative
1995	341	-0.770	0.756	2.516	30.8%	1.8%
1996	578	-0.420	0.994	2.321	37.0%	1.6%
1997	1,165	-0.490	0.925	2.416	35.8%	2.2%
1998	2,064	-0.485	1.072	2.651	37.4%	1.6%
1999	4,057	-0.445	1.115	2.582	43.4%	1.5%
2000	10,429	-0.341	1.033	2.566	44.1%	1.5%
2001	6,695	-0.271	1.037	2.564	44.5%	1.2%
2002	4,722	-0.248	1.143	2.680	47.5%	1.3%
2003	4,958	-0.254	1.138	2.527	43.4%	1.1%
All years	35,009	-0.327	1.070	2.574	43.5%	1.4%

Panel B: Only significant β_d estimates

		β_d estimate		
Year	N Obs	10th Pctl	Median	90th Pctl
1995	105	1.080	1.877	3.084
1996	214	1.000	1.765	2.697
1997	417	1.036	1.757	2.964
1998	771	1.053	1.896	3.248
1999	1,762	1.026	1.905	3.024
2000	4,594	0.944	1.781	3.022
2001	2,977	0.931	1.762	3.042
2002	2,243	0.983	1.864	3.062
2003	2,152	1.032	1.854	2.968
All years	15,235	0.978	1.820	3.031

Panel C: Hazard function estimates

Variable	Mean	t-stat	10th Pctl	Median	90th Pctl
β_r	0.312	1.00	-0.63	0.02	0.83
β_s	-0.014	-0.88	-0.05	0.00	0.04
β_V	0.263	0.63	-2.31	0.25	2.77
β_d	1.130	135.66	-0.33	1.07	2.57

Table 3: Disposition and Returns: Regression Analysis

This table reports results from univariate regressions of returns on various measures of disposition and experience, with year fixed effects:

$$R^h_{i,t} = \alpha + \beta_x X_{i,t-1} + \gamma_t + \epsilon_{i,t}.$$

Returns are measured beginning with the closing price on the day of purchase, and calculated over different horizons as described in Section 3.2. We use average returns, calculated annually for each account, in all regressions in this table. These tests are out-of-sample, in that we regress returns in year $t+1$ on disposition statistics from year t. Each panel presents results for different horizons, and N gives the number of observations used in the regressions. The regressor of interest is shown in the second column. Disposition estimates come from the hazard model in equation (2). Disposition quintiles are formed over all disposition estimates and years. In the rows labeled 'Significant at 10%', etc., we use a variable that takes a value of 1 (-1) when the estimated disposition coefficient is statistically greater (less) than zero at the 10%, 5%, and 1% levels, respectively, and zero otherwise. Coefficients are reported in percent.

Model	Dependent variable	Coefficient	t-stat	Adj-R^2
	Panel A: Average 10-day returns (N=27,241)			
1	Disposition estimate	-0.07	-1.83	0.069
2	Disposition quintile	-0.03	-1.28	0.068
3	Significant at 10%	-0.18	-2.96	0.069
4	Significant at 5%	-0.24	-3.93	0.069
5	Significant at 1%	-0.29	-4.38	0.069
	Panel B: Average 20-day returns (N=27,232)			
1	Disposition estimate	-0.14	-2.39	0.134
2	Disposition quintile	-0.07	-1.71	0.134
3	Significant at 10%	-0.20	-2.13	0.134
4	Significant at 5%	-0.19	-2.01	0.134
5	Significant at 1%	-0.38	-3.59	0.135
	Panel C: Average 30-day returns (N=27,220)			
1	Disposition estimate	-0.36	-4.65	0.229
2	Disposition quintile	-0.18	-3.75	0.229
3	Significant at 10%	-0.43	-3.55	0.229
4	Significant at 5%	-0.38	-3.09	0.228
5	Significant at 1%	-0.68	-5.09	0.229
	Panel D: Average 45-day returns (N=27,184)			
1	Disposition estimate	-0.40	-4.22	0.320
2	Disposition quintile	-0.20	-3.36	0.320
3	Significant at 10%	-0.51	-3.47	0.320
4	Significant at 5%	-0.36	-2.41	0.320
5	Significant at 1%	-0.66	-4.02	0.320

Table 4: Returns and Learning

This table reports the estimates of regressions of the form

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

where Experience is measured by either years of experience (YearsTraded) or cumulative number of trades placed (CumulTrades). The dependent variable in this estimation is individual i's return next year, R_{t+1}, calculated as described in Section 3.3. $X_{i,t}$ 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), the individual's average total daily portfolio value (PortVal), and a dummy variable (bgn), which equals one if an account was opened prior to January 1995 and zero otherwise, interacted with the Experience and Experience^2. We also include year dummies and individual-specific intercepts in each regression. Data are from the period 1995 to 2003. ***, ** and * denote significance at 1%, 5% and 10% respectively.

	Dependent Variable: R_{t+1}				
	(1)	(2)	(3)	(4)	(5)
CumulTrades_t ($\div 10^2$)			2.68	2.64	2.63
			(.940)***	(1.010)***	(1.111)***
CumulTrades^2_t ($\div 10^4$)			-2.92	-2.95	-2.93
			(1.133)**	(1.270)**	(1.261)**
YearsTraded_t	0.423	0.421			0.330
	(.105)***	(.157)***			(.149)**
YearsTraded^2_t	-0.020	-0.016			-0.009
	(.012)*	(.009)*			(.023)
$\text{bgn}_i \times \text{YearsTraded}_t$		0.228		0.287	0.249
		(.182)		(.190)	(.184)
$\text{bgn}_i \times \text{YearsTraded}^2_t$		-0.060		-0.077	-0.065
		(.155)		(.155)	(.155)
NumTrades_t		-0.036		-0.033	-0.033
		(.009)***		(.010)***	(.010)***
NumSec_t		0.007		0.010	0.011
		(.018)		(.018)	(.018)
PortVal_t ($\div 10^6$)		0.120		0.480	0.479
		(.060)**		(.222)**	(.220)**
Observations	27,160	27,160	27,160	27,160	27,160
Adjusted R^2 (%)	10.5	12.8	11.0	13.1	15.3
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes	Yes

Table 5: Disposition and Learning

This table reports the estimates of regressions of the form

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

where Experience is measured by either years of experience (YearsTraded) or cumulative number of trades placed (CumulTrades). The dependent variable in this estimation is the disposition coefficient, β_d, of an individual i in year t calculated as described in equation (2). $X_{i,t}$ is a vector of controls including the average return earned by the individual in the previous year ($\bar{R}_{t-1}$), 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), the individual's average total daily portfolio value (PortVal), and a dummy variable (bgn), which equals one if an account was opened prior to January 1995 and zero otherwise, interacted with the Experience and Experience^2. We also include year dummies and individual-specific intercepts in each regression. Regressions are estimated using WLS, where the weights are proportional to the reciprocal of the variance of the disposition estimates. Data are from the period 1995 to 2003. ***, ** and * denote significance at 1%, 5% and 10% respectively.

	Dependent Variable: β_d				
	(1)	(2)	(3)	(4)	(5)
CumulTrades_t ($\div 10^2$)			-0.108	-0.101	-0.101
			(.025)***	(.034)***	(.033)***
CumulTrades^2_t ($\div 10^4$)			0.006	0.005	0.004
			(.003)**	(.002)**	(.002)**
YearsTraded_t	-0.049	-0.044			-0.027
	(.011)***	(.017)***			(.014)**
YearsTraded^2_t	0.005	0.004			0.004
	(.002)***	(.002)*			(.002)*
$\text{bgn}_i \times \text{YearsTraded}_t$		0.009		-0.045	0.008
		(.027)		(.042)	(.028)
$\text{bgn}_i \times \text{YearsTraded}^2_t$		-0.005		0.005	-0.005
		(.004)		(.004)	(.004)
$\bar{R}_{t-1}$		-0.342		-0.350	-0.356
		(.143)**		(.143)**	(.144)**
NumTrades_t		-0.002		-0.001	-0.001
		(.0007)***		(.0008)*	(.0008)*
NumSec_t		-0.006		-0.006	-0.006
		(.002)***		(.002)***	(.002)***
PortVal_t ($\div 10^6$)		-0.030		-0.031	-0.028
		(.025)		(.027)	(.025)
Observations	35,009	35,009	35,009	35,009	35,009
Adjusted R^2 (%)	7.3	9.1	8.2	10.0	11.8
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes	Yes

Table 6: Disposition and Learning: Controlling for Survivorship

This table reports the results from the Heckman selection model with two-step efficient estimates of the parameters and standard errors. The second stage regression is

$$y_{i,t} = \alpha + \beta_1 \text{Experience}_{i,t} + \beta_2 \text{Experience}^2_{i,t} + \beta_\lambda \lambda_{i,t} + \delta X_{i,t} + \gamma_t + \epsilon_{i,t},$$

where the dependent variable is the disposition coefficient of individual i in year t estimated as described in equation (2), Experience is measured either by the number of years since the account entered the sample (YearsTraded) or the cumulative number of trades placed (CumulTrades), $\lambda_{i,t} = \rho\sigma_\epsilon$ is the Inverse Mills ratio constructed from the first stage regression. $X_{i,t}$ is a vector of controls including the number of securities held by the individual (NumSec), the individual's return in the previous year ($\bar{R}_{t-1}$), and a dummy variable that equals one if an account was opened prior to January 1995 and zero otherwise (bgn). A test of $\lambda > 0$ is equivalent to a test of $\rho > 0$ (see text). We also include year dummies in all the regressions. See Section 3.4 for detail on this procedure. Data are for the period 1995 to 2003. ***, ** and * denote significance at 1%, 5% and 10% respectively.

	First Stage	Second Stage		First Stage	Second Stage	
Dependent variable	Insample=1	$\beta^d_{i,t}$	$R_{i,t+1}$	Insample=1	$\beta^d_{i,t}$	$R_{i,t+1}$
CumulTrades_t ($\div 10^2$)				3.53	-0.112	2.54
				(.081)***	(.013)***	(.241)***
CumulTrades^2_t ($\div 10^4$)				-0.406	0.007	-2.05
				(.004)***	(.002)***	(1.08)**
YearsTraded_t	-0.118	-0.041	0.471	-0.181	-0.028	0.427
	(0.008)***	(0.021)**	(0.149)***	(0.021)***	(0.014)**	(0.150)***
YearsTraded^2_t	0.011	0.002	-0.051	0.017	0.001	-0.053
	(0.001)***	(0.001)**	(0.019)***	(0.002)***	(0.001)*	(0.019)***
bgn_i	-0.441	0.090	0.176	-0.275	0.089	0.245
	(0.019)***	(0.045)**	(0.175)	(0.008)***	(0.028)***	(0.120)**
NumSec_t	0.187	-0.004	0.151	0.105	-0.005	0.110
	(0.001)***	(0.008)	(0.126)	(0.001)***	(0.005)	(0.007)**
PortVal_t ($\div 10^6$)	-0.100	-0.011	0.120	-0.310	-0.028	0.480
	(.018)***	(.008)	(.061)**	(.079)***	(.025)	(.220)***
$\bar{R}_{t-1}$	0.660			0.548		
	(0.099)***			(0.042)***		
$\sigma_{R_{t-1}}$	0.577			0.266		
	(0.075)***			(0.028)***		
ρ		0.114	0.117		0.101	0.138
		(0.012)***	(0.024)***		(0.016)***	(0.015)***
Observations	100,139	35,009	27,160	100,139	35,009	27,160
Pseudo-R^2 (%)	23.3			25.7		
χ^2-test of $\rho = 0$		18.31***	22.25***		19.72***	23.26***

Table 7: Disposition and Learning: Conditioning on Attributes

This table reports the estimates from regressions of the form

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

conditioned on variables described above each of the five panels. The dependent variable in this estimation is either the disposition coefficient (1) or returns (2) for an individual i in year t calculated as described in (2). Controls in each regression include the number of securities held by the individual in a given year, the individual's return in the previous year, and a dummy variable which equals one if an account was opened prior to January 1995 and zero otherwise. For brevity only coefficients of interest are reported. We classify investors, *ex ante*, as sophisticated 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 define the state of the market as 'up' if the excess return on a broad market index is positive, and 'down' if it is negative. Finally, we classify investors who make excess profits that are in the top 75th percentile of the entire market in the first two years of their trading as 'winners.' We include year dummies in all the regressions. We also report average disposition and returns for each group, as well as the number of observations (N). Data are from the period 1995 to 2003. ***, ** and * denote significance at 1%, 5% and 10% respectively. All of the results remain qualitatively the same if we estimate a selection model as in Table 6. We do not report these results in the interests of space. All group means are significantly different at 1% level.

	(1) Disposition			(2) Returns			
Investor Classification	YearsTraded	YearsTraded^2	Mean	YearsTraded	YearsTraded^2	Mean	N
Panel A: Conditioning on Investor Sophistication							
Doesn't trade options	-0.096	0.012	1.17	0.501	-0.048	-0.81	26,698
	(0.024)***	(0.003)***		(0.184)***	(0.029)*		
Trades options	-0.031	0.003	0.99	0.233	-0.018	0.10	8,211
	(0.032)	(0.005)		(0.299)	(0.047)		
Panel B: Conditioning on Investor Wealth							
Not wealthy	-0.101	0.013	1.14	0.389	-0.029	-1.18	26,176
	(0.023)***	(0.003)***		(0.169)**	(0.015)*		
Wealthy	-0.015	0.001	1.11	0.191	-0.039	1.08	8,733
	(0.043)	(0.006)		(0.339)	(0.049)		
Panel C: Conditioning on Investor Gender							
Females	-0.140	0.021	1.24	0.482	-0.026	-0.38	4,703
	(0.047)***	(0.008)***		(0.129)***	(0.016)		
Males	-0.078	0.009	1.11	0.361	-0.044	-0.62	30,206
	(0.020)***	(0.003)***		(0.169)**	(0.029)		

	(1) Disposition			(2) Returns			
Investor Classification	Experience	Experience^2	Mean	Experience	Experience^2	Mean	N
			Panel D: Conditioning on Market Conditions				
Down market	-0.120	0.013	1.14	0.619	-0.064	-0.70	27,074
	(0.022)***	(0.003)***		(0.179)***	(0.029)*		
Up market	0.028	0.006	1.08	0.114	-0.017	1.77	7,835
	(0.077)	(0.021)		(0.471)	(0.122)		
			Panel E: Conditioning on Investor Profits				
Not winners	-0.840	0.010	1.15	0.732	-0.077	-0.62	32,456
	(0.021)***	(0.003)***		(0.161)***	(0.028)***		
Winners	-0.039	0.010	0.96	-0.641	-0.731	-0.03	2,453
	(0.074)	(0.011)		(0.703)	(0.672)		

Table 8: Disposition and Returns

This table presents the summary statistics and regression results using the group-level disposition estimates discussed in Section 3.6. Panel A provides summary statistics for the year-on-year change in disposition coefficients, β_d, and average return over a 20-, 30-, or 45-day horizon. Panel B presents the results of three separate regressions of the form

$$\Delta \bar{r}_t^h = a + b\Delta\hat{\beta}_d + \epsilon,$$

where $\bar{r}_t^h$ denotes the average return at a horizon of $h = 20, 30, 45$ days for each group in year t. The dependent variable is the yearly change in the average return earned over each horizon, as indicated in the first column. Each regression includes an intercept (not reported) and the disposition coefficient, β_d, as a regressor.

Panel A: Summary Statistics

Variable	Mean	Std. Dev.	t-statistic
$\Delta\beta_d$	0.0212	0.517	1.69
$\Delta\bar{r}_{20}$	0.0024	0.068	1.45
$\Delta\bar{r}_{30}$	0.0052	0.110	1.97
$\Delta\bar{r}_{45}$	0.0089	0.152	2.41

Panel B: Regression Estimates

Dependent variable	Coefficient	t-statistic	Adj-R^2
$\Delta\bar{r}_{20}$	-0.0113	-2.94	0.054
$\Delta\bar{r}_{30}$	-0.0182	-2.94	0.054
$\Delta\bar{r}_{45}$	-0.0215	-2.50	0.058

Table 9: Changes in Trading Behavior

This table presents the results of regressions of the form

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

where the α_i denote individual fixed effects and the γ_t are year dummies. Experience is measured by the number of years since the investor first placed a trade. For brevity, only β_1 coefficients are displayed. Each row corresponds to a separate regression. Panel A presents results for variables measuring the investor's level of diversification, using either the number of different securities in the investor's portfolio or the number of different industries held, where industries are defined using 2-digit SIC codes. Panel B presents results for regressions using only purchases, and Panel C uses only sales. See Section 3.7 of the text for an explanation of the other variables.

Dependent Variable	Coefficient	t-statistic
Panel A: Diversification		
Number of stocks held	0.3633	869.60
Number of industries held	0.1456	1088.53
Panel B: Buys		
Size (MM Euro)	1,962	42.34
Volume (MM Euro)	-13.6809	-23.41
Past return	0.0082	15.40
Past volatility	0.0000	-0.86
Panel C: Sells		
Size (MM Euro)	1,470	37.42
Volume (MM Euro)	18.1303	33.03
Past return	-0.0002	-0.72
Past volatility	-0.0004	-16.87

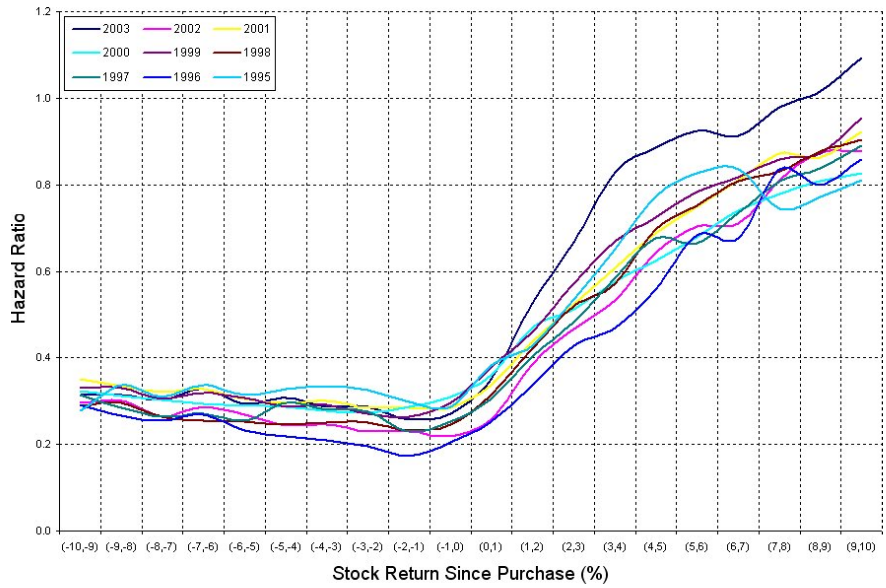

Figure 1: Propensity to Sell

This graph shows how the propensity to sell a stock depends on stock performance since purchase. There is a pronounced kink near zero, and the hazard increases rapidly for positive returns.

Figure 2: Experience Curve

This figure graphs the effect of years of experience on the disposition coefficient. The graph is generated from the estimates in regression (1) in Table 2.

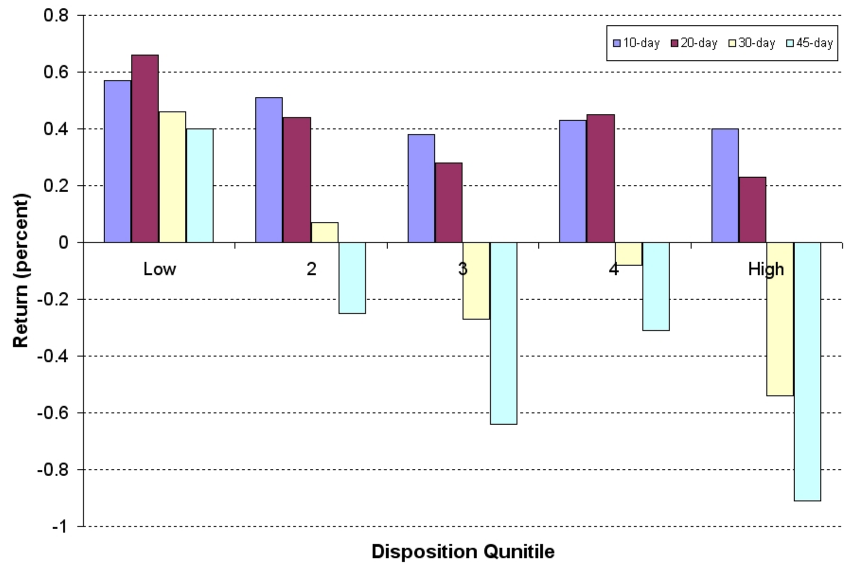

Figure 3: Returns by Disposition Quintile

This figure shows average 10-, 20-, 30-, and 45-day returns following a purchase for each disposition quintile. Returns earned by the lowest quintile (1) are higher than those earned by the highest quintile (5).

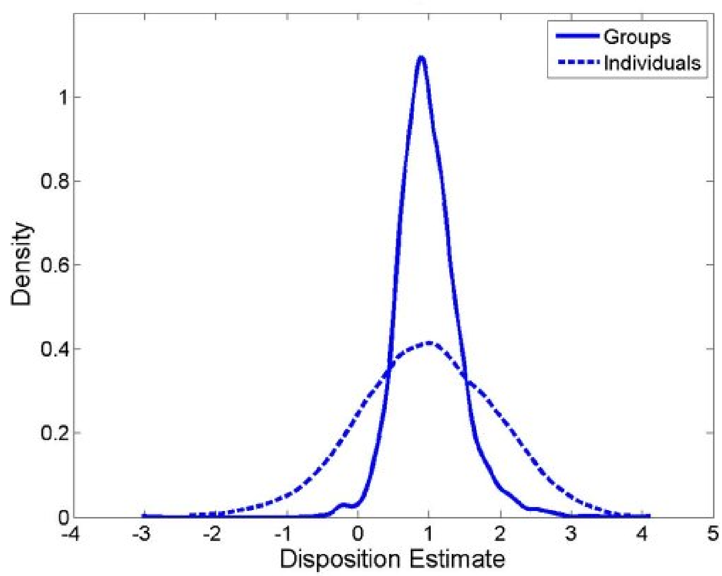

Figure 4: Kernel Density Estimate

Kernel density estimates for the disposition effect are graphed for the account-level estimates (dashed line) and the group-level estimates (solid line). The bandwidth for the density estimation is selected using the plug-in formula of Sheather and Jones (1991).