



City Research Online

City, University of London Institutional Repository

Citation: Hayley, S. and Marsh, I. W. (2016). What do retail FX traders learn?. *Journal of International Money and Finance*, 64, pp. 16-38. doi: 10.1016/j.jimonfin.2016.02.001

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/15691/>

Link to published version: <http://dx.doi.org/10.1016/j.jimonfin.2016.02.001>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

What do retail FX traders learn?[☆]

Abstract

What is the benefit of experience? Using data from a leading trading platform we find no evidence that retail FX traders learn to trade better, but they do appear to learn about their innate abilities as traders and respond appropriately. In particular, following an unsuccessful trading day they are more likely to cease trading, to trade smaller amounts and to trade less frequently. These effects are stronger for younger and less experienced traders who might be expected to have more to learn than older, more experienced traders. As regards learning through experience, surprisingly we find that more seasoned traders demonstrate a slight decline in performance once we account for the endogenous decision to cease trading, and even very experienced traders consistently lose money.

Keywords: Market microstructure, foreign exchange, retail trading, learning

1. Introduction

What is the benefit of experience? To address this we use a new database from one of the fastest growing sectors of financial markets: retail foreign exchange trading. Historically, foreign exchange (FX) trading has been the preserve of large institutions largely due to the enormous costs of trading FX at the retail level. The Bank for International Settlements viewed retail FX trading as negligible in 2001. However, advances in technology, combined with the consensus that FX markets are both liquid and have low correlations with other markets, has led to the establishment of numerous platforms offering even very small retail investors cheap access to FX markets. By 2010, trading in the retail segment of the FX market was estimated to be \$125-150 billion per day, equivalent to 8-10 percent of global spot turnover (King and Rime (2010)). The retail trading sector of the foreign exchange market is in the spotlight after the French financial markets regulator, the Autorité des Marchés Financiers (AMF), released results of a study¹ concluding that 89% of participants examined lost money, and that even active and regular retail investors saw their losses mount. “[I]ndividual investors learn little over time” concluded Natalie Lemaire of the AMF’s Retail Investor Relations Directorate, adding that “Foreign exchange trading is a market that individual investors should avoid.”

We use data from one such retail trading platform which details daily activity levels of over ninety-five thousand individual investors over a thirty month period to investigate how individual traders learn about FX trading and how this affects their decisions to trade. The academic literature has identified two specific ways in which financial market participants might learn. The first, classical, approach is through “learning by doing” whereby traders improve their ability to trade as they gain experience (Arrow (1962); Grossman, Kihlstrom, and Mirman (1977)). This appears to be the type of learning AMF had in mind. An alternative is that traders learn about their inherent trading abilities (Mahani and Bernhardt (2007); Linnainmaa (2011)). If they infer from trading that they have skill, they will continue to trade. Conversely if they infer a low level of ability they cease trading. Both types of learning may coexist and, as discussed further below, any learning may or may not be rational. We investigate the importance of learning by doing and learning about ability for our large sample of retail FX traders.

¹AMF press release, 13 October 2014. See www.amf-france.org

We find evidence that traders learn from their trading experiences as they are more likely to cease trading, to trade smaller amounts and to trade less frequently following losses. These effects are stronger for traders that might be expected to have the most to learn. Trading decisions of novice traders (lacking trading experience) and young traders (lacking life experience) are more sensitive to performance signals than more experienced, older traders. Controlling for the strong endogenous decision to quit results in an economically small negative association between performance and experience. Traders do not learn how to trade better with experience and so even relatively long-lived traders consistently lose money. Non-rational learning, where traders interpret signals about their abilities in excessively positive ways, can explain this result. Our findings suggest that the concerns expressed by regulators may be justified.

In *learning by doing* models, traders learn how to trade by trading.² This could happen in a variety of ways. At a simple level, traders may become more skilled at using the trading platform, leading to better performance. Similarly, they may learn to avoid making trading mistakes, and this may include learning to avoid or reduce the effects of behavioural biases such as the disposition effect. At a more sophisticated level, they may learn better models of how the exchange rate is determined or how to read markets better such that they apply their models more appropriately depending on market conditions. This way of learning leads traders to alter their trading approach through time in response to performance signals but, crucially, suggests that performance should improve with experience, though potentially at a decreasing rate. Empirical evidence in support of a considerable learning by doing effect in equity markets is presented in Feng and Seasholes (2005) and Dhar and Zhu (2006). These papers suggest that traders learn through experience not to exhibit the behavioural bias towards realising gains early (the disposition effect) and that this drives their performance improvement over time. Kaustia, Alho, and Puttonen (2008) report that experienced investors seem less prone to anchoring effects than novice investors. Nicolosi, Peng, and Zhu (2009) show that individuals exhibit considerable improvement in performance as their experience grows (a risk-adjusted portfolio return increase of around 2% per year). Similarly, Barrot, Kaniel, and Sraer (2014) show that experienced equity retail traders trade faster than less expe-

²Dhar and Zhu (2006), Nicolosi, Peng, and Zhu (2009) and Seru, Shumway, and Stoffman (2010).

rienced traders and outperform as a result, suggesting they have learned better trading skills.

In *learning about ability* models, traders initially do not know how good they are at trading. For example, Linnainmaa (2011) and Mahani and Bernhardt (2007) present rational learning models whereby investors actively trade in order to learn about their innate abilities, knowing that the population is heterogeneous with some traders being skilled and others unskilled. Traders behave as Bayesian updaters, learning from their trading histories and changing behaviour accordingly. Once a trader receives a strong enough set of signals that she lacks trading skills she ceases trading. A trader receiving positive signals continues to trade and should increase trading activity. As traders gain experience their sensitivity to new performance signals declines. Long-lived traders should perform well as these are the traders who have learned that they have skills. In this approach to learning, the models and abilities of a trader can be thought of as fixed and the trader learns over time whether these are good enough to warrant continued trading.

One important implication of the Linnainmaa (2011) model is that if an individual places a high enough value on observing another signal about her ability she will trade even though she rationally expects to lose money. Since she expects to lose money she only trades small amounts. If she receives a positive signal through making a successful trade, she infers skill and subsequently trades more. If she makes a losing trade she infers a lower level of skill and trades less. Eventually, if an investor receives enough negative signals she may cease trading entirely. The impact of additional signals of ability is largest at the start of a trader's career and for traders with diffuse priors. The incremental effect of a positive signal on the decision to continue trading or on the scale of trades declines with experience.

The basic predictions of alternative non-rational learning models are very similar. In these models traders may again trade to learn. However, their learning is biased by one or more behavioural tendencies - including overconfidence or attribution bias - so even poorly performing traders continue to trade.³ Chiang, Hirshleifer, Qian, and Sherman (2011) consider irrational investors subject to naive reinforcement learning. Here, in-

³Odean (1998), Gervais and Odean (2001), Barber, Lee, Liu, and Odean (2011), Chiang, Hirshleifer, Qian, and Sherman (2011).

vestors expect the returns that they have personally experienced to recur even when this expectation is unjustified. Experiencing a trading success leads them to expect future trading success - the same tendency as a rational Bayesian learner - however naive reinforcement learners place too much weight on their personal experience and do not update their beliefs optimally. Gervais and Odean (2001) describe how another behavioural bias - overconfidence - can also lead investors to make mistakes as they gain experience. They model the process by which individuals learn to be overconfident about their trading abilities. A positive signal leads a trader to update her belief about her skills. However, attribution bias leads her to do so to an inappropriate extent, over-weighting the probability that success is due to superior ability and under-weighting alternative explanations such as luck. Conversely, unsuccessful trades are deemed due to external forces to too great an extent. In Gervais and Odean's model investors do not start out overconfident. However, overconfidence increases over initial trading periods before declining. As she gains experience she develops a more realistic assessment of her abilities. But while she is overconfident she behaves sub-optimally, trading too aggressively which lowers her expected profits. In these non-rational models even experienced traders may still perform badly and performance may deteriorate with experience. It is, however, anticipated that performance finally improves as rational updating prevails.⁴

This is the crucial difference between the rational and non-rational approaches to learning. In a rational model, expected returns increase with experience, either because traders have learned by doing how to trade better or because traders have learned about their abilities by trading and only the good remain active. Conversely, expected returns may actually decline with experience for non-rational models, at least until investors have learned to avoid behavioural biases. This distinction allows us to differentiate between alternatives by relating performance to experience.

Most of the empirical work on learning in financial markets has used data from equity markets, often considering the equity trades of Finnish or Taiwanese investors.⁵ We add to our understanding of learning in financial markets by considering data from the foreign exchange market. This comes with the added advantage, relative to secondary market

⁴List (2003), Agarwal, Driscoll, Gabaix, and Laibson (2008) and Kaustia, Alho, and Puttonen (2008).

⁵Finland: Seru, Shumway, and Stoffman (2010); Linnainmaa (2011). Taiwan: Chiang, Hirshleifer, Qian, and Sherman (2011); Barber, Lee, Liu, and Odean (2011).

equity data, of not having to separate active trades from long-term passive investments. Due to the nature of the trading architecture used, the overwhelming majority of the trades in our database are active.

We first test for cross sectional heterogeneity in trading ability. In a rational learning about ability model traders trade in order to learn whether they are skilled or not so at least some retail traders must demonstrate skill. If no individual trader is skilled and this is known by all, then each individual knows that they are not skilled and hence would not trade.⁶ The learning by doing explanation instead relies on traders being able to improve their performance over time through the act of trading. Again, this would imply that some traders are better than others. Anticipating our results, we find very strong evidence of cross-sectional heterogeneity in trading ability supporting the basic premises of the competing learning models.

Second, we test whether the decision to cease trading is related to performance. If traders learn about their innate abilities by trading then their performance each day will give them a signal. If the learning about ability explanations are correct we would expect these signals to influence the decision of an individual whether to continue trading or not. We also test whether the sensitivity to performance signals is higher for the traders we would expect to have the most to learn (e.g. novice traders or younger traders with less life experience).

Third, and closely related to the decision to quit, we test whether traders receiving positive signals about their abilities alter their trading activities. As traders learn, either about their innate abilities or how to trade, a positive signal makes subsequent trading more attractive. Hence we expect traders to trade in greater volume or more frequently following a profitable trading day. Again, our results strongly support the hypotheses that traders learn from their performance. The decision to quit trading is strongly negatively correlated with both performance on the previous trading day and career trading performance. The sensitivity of the quit decision to performance declines with experience and age, consistent with learning effects being stronger for traders with the most to

⁶Irrational traders may have unrealistic priors about their own abilities or those of the population that are unsupported by the data. Even if no individual trader actually has skill an irrational trader may still trade because they irrationally believe that they (or some proportion of the population to which they may belong) has skill.

learn. Finally, traders also change their trading behaviour following success. Trade size increases following good performance and the gap between trades decreases.

However, there are some aspects of the relationship between trading signals and subsequent trading decisions that are harder to reconcile with purely rational learning about ability. For example, while there is a generally positive relationship between the dollar amount earned on day t and the probability of an increased trade volume on day $t + 1$, this is actually strongest in the immediate vicinity of breaking even. That is, both tiny gains or losses predict increased future trading volume much more strongly than do larger gains. We put forward an explanation based on traders learning about risk, and hence increasing their exposures if they make a daily profit or loss that they consider trivial.

The final test is specifically designed to discriminate between learning models. The test examines whether performance is related to experience. The learning by doing and rational learning about ability models both suggest performance should improve with experience. In its pure form, the rational learning about ability model assumes that each trader's ability is fixed. Average performance improves with experience in the population because unskilled traders cease trading and only skilled traders remain. That is, performance improves because of the endogenous decision to quit in a heterogeneous population. Once this is accounted for, performance should not be related to experience. The learning by doing model predicts a positive relationship between performance and experience even after the endogenous decision to quit is taken into account. Traders who remain in the sample get better at trading with experience (i.e. ability is not fixed). The irrational learning alternative suggests that performance may instead deteriorate with experience. Even correcting for the quit decision, the performance of traders who remain active may worsen with experience as their behavioural biases lead them to 'learn to fail'.

Analysing the relationship between performance and experience is complicated in the presence of heterogeneity and endogenous quit decisions. We use a Heckman-style approach that corrects for cross sectional heterogeneity and selection biases. Both of these corrections have significant impacts on our estimates of the experience-performance relationship. Once they are taken into account we find that performance deteriorates with experience, consistent with the non-rational learning model. The magnitude of these effects is not large and a trader with 100 days of active trading experience - large in terms

of our sample - is only 5.8% less likely to have a winning month than a complete novice. Overall, retail FX traders appear to learn in two main ways. First, they learn about their innate skills by trading and are more likely to trade less or even cease trading after bad performance signals. The sensitivity of trading decisions to negative signals is much larger for traders who are likely to be learning the most (novice traders and younger traders). These effects are statistically significant and economically large. While broadly consistent with rational learning, there are also elements of non-rational inference in these results. Second, after accounting for these learning about ability effects, a small amount of non-rational learning remains. Trader performance deteriorates slightly with experience and this explains why even relatively long-lived traders who might otherwise be expected to have learned that they have skill and/or learned how to trade still do not perform particularly well.

Our paper is related to another study of learning in foreign exchange markets. Ben-David, Birru, and Prokopenya (2014) show that retail FX traders attribute random success to their own skill as suggested by Gervais and Odean (2001). This attribution bias leads to overconfidence, which they demonstrate has consequences for risk taking by traders. These findings complement ours well. We build on their insights to show that such biases can explain why some traders' performance deteriorates with experience, though the magnitude of these effects is small.

The rest of the paper proceeds as follows. The next section describes our data set, section 3 details our statistical analysis and we conclude in section 4.

2. Data

2.1. Data source

The data used in this paper come from an on-line retail foreign exchange trading platform that wishes to remain anonymous. The data are in two files, a trading database and a trader characteristics database. The two databases can be linked through identification codes which are unique to each trader.

The trading database contains daily records of the complete trading history of a random sample of traders using the platform. All trading is for real money. The data provider

does offer paper trading facilities to customers but these are not included in our data. The platform is continually active Monday through Friday but trading is not possible on weekends. Mark-to-market reconciliation of open trading positions takes place at 9pm GMT each day. The platform allows for trading in all the major currencies primarily in the spot market although a small number of cash-for-difference trades are included in the data. The data provider informs us that the majority of trades are in euro-dollar spot. Traders may use the usual range of market orders, limit orders, stop loss and take profit orders.

For each trader the database gives a daily record of:

1. Number of trades made
2. Total value of trades (USD)
3. Total value of positions that remain open after the 9pm reconciliation (USD)
4. Profit/loss, realised and marked-to-market as appropriate, after trading costs⁷
5. Capital injections and extractions (USD)

The trader characteristics file contains the following information:

1. Age of trader
2. Location of trader by country
3. Location of trader by city (for some countries)

Details of individual trades are not available, only the daily aggregates. Only the US dollar equivalent of open positions is disclosed. The database does not record the currencies of any positions or the direction of any exposures. Daily profits/losses are calculated directly from trading positions open on that day and do not include any translation profit/losses resulting from the effect of exchange rate fluctuations on the value of cash held on account.

While the trader location fields indicate that traders are located in 98 different countries, the client base is extremely focused. Only four countries have more than one percent of

⁷Spreads are relatively narrow in this market. The data provider reports a typical spread of 1.9 pips in dollar-euro, 1.8 pips in yen-dollar and 2.7 pips in dollar-sterling.

the client base and two countries between them have 89% of traders. The city of residence is included for most clients.

The data begin on 4 January 2010 and end on 29 June 2012. The data provider gave us the complete trading history of a random sample of 95,617 unique traders, amounting to almost 4.8 million trader-day entries. Since this is a new database and information at such fine detail on retail trading is scarce we present detailed descriptive statistics in the following subsection. For confidentiality reasons, however, we cannot disclose certain statistics.

A new trader can be identified by the placing of an initial deposit into his or her trading account.⁸ Traders typically start trading very soon after this deposit is made, often on the same day. Some ten percent of our sample of traders commenced trading before our data begin and so our data are left censored. In the majority of the analysis below such traders are included in the sample. For the survival analysis, however, these traders are removed from the sample (and right censoring where traders have not stopped trading within the sample is accounted for in the analysis).

A trader is deemed to have ceased trading if his last observed active day is more than one month before the end of the sample. The mean interval between trading days is just four days in the sample and the median is one day. Just two percent of intervals exceed one month and so a one-month cut-off should eliminate most misidentified exits. Our results are very robust to alternative cut-offs.⁹

We define a trader to be “trading” on a particular day if the number of trades made on that day is non-zero. We define a trader to be “active” on a given day if he either trades on that day or has maintained an open position from the previous day.

2.2. Descriptive Statistics

Table 1 shows that the mean (median) trader in the sample is active on 45 (20) days. There is large variation in this number across traders, however, and the top quarter are

⁸We only have gender identification for a small proportion of traders but the overwhelming majority are male and so for ease we will identify traders as masculine in the rest of the paper.

⁹It is possible that traders join our data provider’s network after trading elsewhere, or that traders leave the network to trade elsewhere. In these cases we would mismeasure their experience and lifespans which would add noise to our analysis.

active on at least 51 days, the top ten percent are active on at least 115 days and the top 5% are active on over 182 days. Conversely, the bottom quartile trade on fewer than eight days. The large number of relatively inactive traders is in common with related databases. Hieronymus (1977) states that over one-third of futures trading accounts were traded only a few times and Linnainmaa (2011) shows that the least active quartile of Finnish short-term equity traders only ever make two trades.

< Table 1 about here >

The median trader's average daily trading volume is approximately \$41,000 and the cross-section median of average trades per day is seven (one round-trip counts as two trades).¹⁰ The distribution of trading volume is extremely skewed and widely dispersed. More than one-quarter of traders have an average daily volume in excess of \$100,000 and more than 5% trade in excess of \$0.5 million per active day on average. The distribution of the average number of trades per day is much tighter. Ninety-five percent of traders make fewer than 25 trades per day on average.

The median trader has a success ratio (percentage of profitable days) of 50%. The mean is slightly lower. Again, the dispersion of success rates is high. More than five percent of traders lost money every day they traded. At the other end of the distribution, ten percent of traders gained on at least three out of every four active days.

The median trader loses \$5.56 on average each active day (plus the opportunity cost of capital which we ignore). The distribution is left skewed and the average trader's average loss is much larger.¹¹ The overwhelming majority of our traders cannot live off the proceeds of trading since the average daily profit of even the trader at the 99th percentile is measured in just tens of dollars. Conversely, the lowest few percentiles are losing upwards of several hundred dollars per day, on average. Our traders are therefore likely to be similar to Mahani and Bernhardt (2007)'s prototypical novice speculator,

¹⁰We use medians rather than mean values in this section for confidentiality reasons.

¹¹This left skew is consistent with a significant disposition effect (as has been found in other markets) which leads traders to close out profitable positions more rapidly than loss-making positions. This skew would also be consistent with traders having a daily profit target, but on average traders lose money, so only a minority are likely to hit any positive profit target. As Table 5 later shows, traders tend to quit after loss-making trades which is consistent with learning about a lack of ability, but not consistent with a lifetime profit target being met.

Kiyoshi Wakino, and pursue trading alongside a regular job rather than trading to make a living.¹²

Losing money on average is another typical finding in retail trading markets. Linnainmaa (2011), for example, reports that the median Finnish active equity trader makes a loss of 21 euros per trade on average. As such, it appears that the level of trading in the foreign exchange market by retail customers is excessive. However, if individuals are uncertain about their trading abilities they can learn by trading. If the value of learning through making a trade exceeds his expected loss on that trade then it makes sense for an individual to trade ‘too much’. The results of Dorn and Sengmueller (2009) suggest an alternative interpretation: That some equity investors derive nonpecuniary entertainment benefits which encourage them to trade even though active trading reduces the expected payoff from their investments. We consider this explanation in section 3.4 below.

We do not report percentage returns in Table 1 due to data limitations. The database reveals the number and value of trades made on a day and any profit or loss made on that day (realised or marked-to-market). However, we cannot map profits or losses to individual trades. Dividing profit by traded volume at a daily frequency results in extreme values, not least because of instances where marked to market profits may occur on days without trades. Similarly, dividing profits by the capital value of the account yields extreme values when capital values are low. In reality, the capital supporting these trades is in the bank accounts of traders rather than in their account with the trading platform. For these reasons we do not typically use returns in analysis performed at the daily frequency except as robustness checks, and focus instead on success ratios. In the final part of the paper we move to a monthly frequency. Here we do use a proxy for returns (*Ret*) calculated as the total cumulated profits over a month divided by the total volume traded during the month. This proxy has more reasonable properties, though it is still subject to a few huge outliers which we address through winsorisation. We also continue to use a monthly version of the success ratio in parallel with *Ret* and our results from these two proxies for performance are comparable.

¹²We cannot divulge exact values in this section of the paper for confidentiality reasons.

3. Analysis

3.1. Cross sectional heterogeneity in trading ability

Heterogeneity in trading ability is an important component of conventional learning models. If traders trade in order to learn whether they are skilled or not then at least some retail traders should demonstrate skill.¹³ Learning by doing implies that traders are able to improve their performance through the act of trading and this also implies that some traders are better than others. Abbey and Doukas (2014) select 428 foreign exchange accounts with relatively long trading histories (more than 30 active trading days) from a population of over 9,000 accounts. They show that one-quarter of these selected accounts display skill, significantly outperforming a four factor currency model.¹⁴

We test for cross sectional heterogeneity by regressing the i th-trading day success dummy on the career success rate over all previous trades. We augment the regression with monthly fixed effects and dummy variables capturing the number of days experience each trader has. If no differences in performance exist or if any differences are merely transitory then the coefficient on career success rate will be zero. Systematic outperformance by some traders will result in a positive coefficient.

The results presented in Table 2 indicate that some traders systematically outperform others, supporting the results of Abbey and Doukas (2014). The coefficient on career success rate is statistically and economically significant in each variant of the regression.

< Table 2 about here >

We also split the data into investors' first ($N \leq 10$), early ($10 < N \leq 25$), intermediate ($25 < N \leq 50$) and late ($N > 50$) trades. The coefficients are positive and statistically significant in each sub-sample, and grow larger as we consider traders' successively later

¹³It is not essential that the skilled traders are active on the particular platform we analyse. However, if our sample of traders is representative of the population of traders we might expect to see some skilled traders here. Even if no trader is in fact skilled, a participant may still trade if he (wrongly) believes some traders have skill and that he may be one of them.

¹⁴Barber, Lee, Liu, and Odean (2014) show that around 1,000 Taiwanese equity day-traders are able to earn predictable abnormal returns net of fees out of a population of 360,000 day-traders. Linnainmaa (2011) demonstrates considerable cross-sectional heterogeneity in the performance of frequent retail traders in the Finnish equity market, as do Barrot, Kaniel, and Sraer (2014) in their analysis of French retail traders.

trades. The estimated coefficient increases for later trades since the average career success rate becomes more precisely estimated as the number of observations per trader increases. The final row of the table gives the standard deviation of the career success rate and shows that it falls as we consider later career trades. Nevertheless, the net effect is that heterogeneity persists, and actually increases, as we consider successively later career trades. A one standard deviation increase to average career success rates increases today's probability of success by 5.5% in the first trades sample, 7.4% for early trades, 8.5% for intermediate trades and 10.1% for late trades. Without conditioning on experience, a one standard deviation shock to career success rates increases today's probability of success by 7.5%.

We might be concerned that these results are influenced by the disposition effect and that the career success rate is an imperfect measure of performance. We therefore replicate these regressions using a daily measure of returns (the total profits divided by the total volume traded). As noted above this is a problematic variable since it results in extreme values caused by marked-to-market returns of open positions on days with no or low trading volumes. We address this in two ways. First, we simply winsorise daily returns at the one percent level. The benchmark regression analogous to column (3) of Table 2 results in a coefficient on mean career return of 0.20 (t-stat = 19.5) suggesting that high average career returns predict higher future returns. Second, we only include frequent traders in the regressions. These are defined to be those traders in the top quintile of average number of trades per day and in the lowest quintile of overnight open positions (calculated as the average value of positions left overnight divided by the average daily trading volume). Our returns measure is likely to be more appropriate for these traders since trading within the day is dominant, and overnight positions small. The benchmark regression for this reduced sample of 12,044 traders also results in a significantly positive coefficient of 0.12 (t-stat is 9.7).

These results show that there is clear cross sectional heterogeneity in the performance of our traders.¹⁵ It is therefore reasonable that traders may participate in this market in order to learn about their own ex ante unobservable abilities. It is also possible that

¹⁵In section 3.4 we also demonstrate that a small proportion of traders in our sample has persistent ability to earn profits.

some traders are better than others because the former group has learned how to trade better, consistent with the learning by doing approach.

We might be concerned that it should not be possible to detect persistence in performance in this market given that exchange rates are frequently characterised as following a driftless random walk. Substantial robustness testing leads us to believe that our persistence findings are not driven by outliers or even small subsets of the data but is instead pervasive. Rather, we believe that the results are driven by the fact that exchange rates in our sample did not follow a random walk. As noted above, we do not know exactly which exchange rates are being traded but the data supplier told us that euro-dollar trading dominates and so we will illustrate with reference to this rate. Our data spans the period January 2010 through June 2012, some 640 trading days.

Over this period the euro-dollar exchange rate exhibited slight mean reversion over horizons of up to one week and conversely a tendency to follow sustained trends over multi-week horizons.¹⁶ Thus within our sample there will have been trading strategies which tended to be profitable, specifically, short-term range-trading and longer term trend-following strategies. Indeed, there were multi-week periods (significant compared to the relatively short active trading lives of many of the traders in our dataset) when simply being consistently long or short EUR/USD would have generated attractive profits. Thus within our database traders who followed such strategies will appear to exhibit persistent skill.

3.2. The decision to quit

The implications of the learning about ability class of models are that following a positive signal a trader will (a) continue to trade since he infers skill and (b) increase his level of activity (frequency of trading and/or size of positions). Activity levels increase since

¹⁶This is not strong evidence of weak form market efficiency, since the period is relatively short and more powerful tests suggest instead that developed market exchange rates can be characterised as fairly close to random walks (Pukthuanthong-Le and Thomas (2008)). Our database is a panel comprising a relatively short time period, but a very powerful cross section of individual traders. For this reason we do not aim to test market efficiency, and instead focus on how traders learn from their trading. If, based on analysis of much longer time series, we were to regard FX markets are close to efficient, then we would expect these traders to ultimately find that their strategies are not consistently profitable. However, this would not alter the fact that traders who are basing their decisions on the much smaller dataset of their own trading record will rationally regard trading profits as a signal that they have trading skill.

exploratory trades, where a trader expects to lose but is willing to pay to learn about his skill level, are likely to be very small to minimise the cost of learning. As he becomes more confident in his ability – and in the non-rational models he may become overconfident of his ability – he increases his activity level in the expectation of positive returns. Conversely, if a loss is made then the trader infers less skill and will trade less. Traders with a poor performance record will eventually infer that they are unskilled and, since trading is costly for them, will quit.

In this section we consider the decision to quit. We model this decision using a Cox proportional hazard rate model. Cox regressions model how different outcomes affect the hazard rate without the need to estimate or specify the baseline hazard. We include two main variables in the regression, the most recent trading day success dummy and the career success rate over all previous trades. We expect both variables to exert negative influence over the decision to quit since a success on the most recent trading day or a higher than benchmark career success rate should encourage a trader to continue as they are positive signals of his ability.

The basic model estimated is:

$$h(t|x) = h_0(t) \exp(a + b_1 \text{Day}(t) \text{Success} + b_2 \text{CareerSuccessRate} + \text{Controls}) \quad (1)$$

where $h_0(t)$ is the unspecified baseline hazard.

< Table 3 about here >

The estimates reported in Table 3 suggest that the decision to quit is significantly influenced by performance of the most recent trading day and on career performance. A profit on the most recent trading day reduces the hazard rate by 16% relative to the baseline level of a loss-making day. That is, a trader is much less likely to quit trading tomorrow if he has made a profit from trading today. Similarly, the career success rate exerts a negative influence - the better the career success rate, the less likely it is that the trader will quit. Control variables such as the log of cumulative trading volume or number of trades, and month-year dummy variables are typically significant but their inclusion does

not affect the key coefficients.¹⁷

In learning about ability models, the sensitivity of the quit decision to success on the most recent trading day should be time-varying. Specifically, a novice trader learns a lot from the results of a day of trading and so the sensitivity of the decision to quit to success on that day is high. Conversely, an experienced trader with a longer history of trading learns less from an extra day of trading and so his quit decision is less sensitive to one day's performance. We follow Linnainmaa (2011) and perform a series of cross-sectional linear regressions. The dependent variable is an indicator dummy that takes the value one if the trader ceases trading after the current trade, and zero if he continues trading. The explanatory variable is the current trade success indicator dummy. Following the results of the Cox model above, this coefficient is expected to be negative as success should reduce the probability of quitting. The magnitude of the slope gives the sensitivity of the decision to quit to the information gleaned from the day's experience. We first run the regression for all traders active for their first trading day. The regression is then run for trader's second active day and another slope estimate is recorded. The process is continued up until the 50th active day.

< Figure 1 about here >

Figure 1 plots the evolution of this slope coefficient. The plot reveals the expected path. The sensitivity of the quit decision to a successful trade is always significantly negative but the sensitivity is much greater in the early days of a trader's career. A positive signal regarding the trader's ability in the form of a successful trading day reduces the likelihood that the trader will cease trading but this effect is reduced as the trader's track record extends. The learning about ability models interpret this as the speed of learning about the trader's skill level decreasing over time.

We can push this analysis a little further since demographic characteristics give us a cross sectional dimension to the test. A novice trader typically learns a lot about his likely level

¹⁷We include all traders in this analysis but since some traders were already active when our data begin we cannot calculate their true career success rate, only their success rate over the period they are in our database. At the suggestion of the referee we therefore re-ran the regressions using only new entrants for whom we do observe a true career success rate. The results are essentially unchanged for both Table 2 and Table 3.

of trading skills by trading. This speed of learning is likely to be greater the less informed about his ability the trader is initially. In particular, young novice traders are likely to be less informed about their abilities than an older novice trader since the older trader may also have received signals about his skill level as a trader from his life experiences. We therefore expect younger (older) traders' quit decisions to have greater (lower) sensitivity to trading performance. Figure 1 is consistent with this intuition. The evolving sensitivity for older traders is always above that of younger traders, implying that performance has a lower impact on older traders' decision to quit than younger traders' decisions even though they both have the same trading experience. However, other factors related to age (such as trader wealth) may also be at work here.¹⁸

3.3. The reaction to a signal

A second implication of traders learning about their abilities is that trading becomes more attractive after successes and less attractive after failures. This could manifest itself in several ways. Traders receiving a positive signal about their ability could increase their average trade size, trade more times per trading day, or trade on a higher proportion of days. We test the first two responses jointly by regressing the log of the US dollar value of trades on trading day t on trading day $t-1$'s success indicator and a lagged dependent variable. Traders who exit on day t are excluded from the regression since we already know that a loss increases the likelihood of quitting. The regression therefore measures whether outcomes affect trade volume even after excluding exits.

< Table 4 about here >

The estimates in Table 4 indicate that profitable trading on day t leads to a 19% ($=\exp(0.177)$) increase in the volume traded on the next active day relative to a trader making a loss on day t . We can again condition the analysis on the career status of the trader. Success during the first trades ($N \leq 10$) sees trade volumes increase by more than 20%, but this slightly falls to 18.8% for late career trades ($N > 50$). More distinctly, older traders increase their trade volumes by less than younger traders (21.4% versus 15%

¹⁸We thank an anonymous referee for this insight. Unfortunately, we do not have access to data on trader wealth and so we cannot explicitly test this explanation.

respectively). These findings are again supportive of a decreasing impact of a signal as traders become more experienced, either through having traded for longer or, especially, through being older.

Trading volumes are very volatile and we repeat the analysis with the dependent variable being an indicator variable taking the value of one if there is an increase in the volume traded on the next active day and zero otherwise. This is a simple linear probability model testing the effect of profitable trading on day t on the probability of increasing trading volumes on the next active day. Again we run this regression for the full sample and conditioned on the career status and age quintile of the traders. Results are reported in panel B of Table 4. The coefficients on the success indicator are all positive and highly significant, and suggest that volumes traded on day t are around six percent more likely to increase relative to volumes traded on day $t - 1$ if day $t - 1$ was a successful trading day. The magnitude of the effect shows the same slightly declining trend with experience but sensitivity to the trading signal is most clearly inversely related to the age of the trader.

Traders also trade more frequently following a positive signal about their abilities. In panel C of Table 4 we replace the volume of trade with the gap in days between active trading days as dependent variable. The coefficient on the lagged success signal is, as expected, significantly negative indicating that a positive signal of ability results in a shortening of the gap between trading activity. This effect is smaller for more experienced traders ($N > 50$) than novice traders. Again, the impact of success is much larger for younger traders than for older traders.

In our analysis to date, a successful trading day has been defined by a positive profit irrespective of the size of the profit.¹⁹ Results based on this simple indicator suggest that, compared to a non-profit day, a profitable trading day increases the probability of trading in higher volume on the next active day by almost six percent. In Figure 2 we plot the average probability of increasing trade volumes in period $t + 1$ for different values of day t profit or loss.²⁰

¹⁹Recall that we cannot reliably calculate returns due to data limitations.

²⁰The plot excludes traders who quit trading immediately after t and we exclude the profit or loss from the first trading day of each trader. We also exclude observations with exactly zero profit or loss from this plot.

< Figure 2 about here >

Three features of this plot are apparent. First, and consistent with the regression results, the probability of increasing subsequent trading volumes is higher following profits than losses. Second, the probability of increasing trade volume on the next trading day increases as we approach zero from either side. A gain or loss of just a few cents has the largest impact on subsequent trading volumes. Note that this effect is not driven by abnormally low trading volume on day t generating small profits or losses which is then followed by more normal volumes on day $t+1$. We obtain a very similar pattern if we only consider the effect of day t profits and losses generated by volumes that are above average for that trader. Third, there is a significant discontinuity at zero. While the probability of increasing subsequent trade volume is high for the very smallest losses (close to 50%), the probability jumps by around four percent for the very smallest profits. That is, a loss of one cent has a significantly smaller effect on the probability of increasing trade volume than a gain of one cent.²¹

It is hard to reconcile the last two features of the plot with rational learning. In a rational learning by doing framework, the most informative events are likely to be large gains (or losses). Yet the probability of trading more in the next period peaks following the very smallest gains or losses.²² One explanation may be that traders are learning about risk, rather than about their own ability. A trader may establish a position but be uncertain about the likely variance of profit and loss that will result, perhaps because of uncertainty about the prevailing volatility level of the market. Putting on a position that results in a near zero pay-off may signal that a larger position is needed, and hence very small pay-offs lead to increased subsequent position sizes. However, a rational trader seeking to learn about his ability (or about risk) ought to make very similar inference from making a few cents as he does from losing a few cents, yet the discontinuity at zero suggests this is not the case. These findings are perhaps more supportive of some form of behavioural biases in learning. We explore this further in subsequent sections.

²¹A formal regression discontinuity design analysis using the approaches detailed in Calónico, Cattaneo, and Titiunik (2015) confirms the magnitude and high statistical significance of the switch from loss to profit.

²²We performed a similar analysis on the decision to quit however unreported results reveal no particularly anomalous behaviour and the probability of quitting is monotonically declining in profit, even around the breakeven point.

The results of sections 3.2 and 3.3 suggest that decisions on when to quit and how much to trade are consistent with rational learning. The decision to start trading is more debatable. Formal learning-about-ability models require that some traders must be profitable so that it is worthwhile for new traders to enter the market in order to discover whether they have similar innate skill. Consistent with this, our data suggests some persistence in profitability, suggesting that some traders do have such skill.

Such persistent trading skill requires FX markets to be inefficient. The substantial body of academic research on the subject may lead us instead to conclude that developed FX markets are generally efficient. On this view, the apparent trading skill that we detect in our sample would reflect luck rather than skill and would not be expected to persist out of sample. The decision by traders to enter such an efficient market would thus have to be regarded as irrational.

The rationality of the decision to enter the FX market thus depends on what information we can reasonably expect traders to have access to beforehand. Judged against the sum total of publicly available information it could be judged irrational. With more limited prior knowledge, and influenced by marketing material suggesting that such trading is profitable, the decision to enter the market is more understandable. Failure to make full use of the available data could be regarded as evidence of behavioural bias: base rate neglect (a trader's failure to realise that his own direct experience represents a very weak dataset compared to the whole body of available data) or availability bias (giving excessive weight to the emotionally vivid personal experience). The question of whether traders' entry into this market is rational ultimately depends on the interpretation of factors which are outside the dataset that we investigate in this paper. By contrast, the learning behaviour that we can observe within our dataset is largely consistent with rational learning, perhaps with the exception of the behaviour described in Figure 2. These findings are, however, largely consistent with all classes of learning models. In the remainder of the paper we attempt to distinguish between the alternative learning models.

3.4. *Persistent losers*

The literature has noted a problem with similar analyses of retail investors in different asset classes. If traders learn by doing, then long-lived traders should be profitable since they have learned to be competent. Similarly, within a rational learning about ability context, novice traders will be willing to incur small losses on average on their trades as long as there is a large enough positive expected payoff to learning that they have trading skills. While there is evidence that retail investors in equity markets learn by trading, the puzzle remains that even very experienced market participants are not, on average, successful. The average equity retail trader does not learn to be profitable, either by learning how to trade or by learning that he has innate ability and should persist. Non-rational learning about ability models are more consistent with the evidence in this regard. Unduly optimistic traders trade too frequently and too aggressively, both of which reduce expected returns. In these models, investors may learn to fail in that their performance worsens as they become more experienced.

< Table 5 about here >

We investigate the performance of our traders in Table 5. This table reports average performance statistics for traders with differing lifespans. We report three performance statistics: the average daily profit or loss measured in dollars, the success ratio, and an indicator variable ‘Winner’ that takes the value unity if the trader had a positive cumulated profit at the end of his trading career.²³ The average value of this final indicator tells us the proportion of traders who finished with a profit. We consider traders with differing lifespans (N^*) given in the column heading. The first row of each panel gives lifespan averages. So, for example, the average daily loss of traders active for between 10 and 25 days is \$20.22 over their trading lifetime. Subsequent rows give statistics for first ($N \leq 10$), early ($10 < N \leq 25$), intermediate ($25 < N \leq 50$) and late ($N > 50$) trades. So, those same traders active for between 10 and 25 days make an average daily loss of \$11.43 over their first ten trading days but this rises to an average loss of \$32.81 for the trades made between day 11 and their final trading day.

²³Or when the sample ended if the trader is right censored.

The first column of figures suggests that the most short-lived traders ($N^* \leq 10$) perform noticeably badly using all three performance statistics. They make losses of \$30.24 per trading day on average, make money on just 40% of trading days, and less than ten percent make money over their (short) trading lifespan. This is what we should expect from all models of learning. Traders who cease trading after relatively few periods have either concluded that they have no skill, or have not spent much time honing their trading techniques and so have not learned how to trade well. They would therefore be expected to perform badly.

All experience categories on average lose money and the proportion of traders who win over their career remains low irrespective of experience. These full-sample statistics are averages over the entire life of traders. Our earlier results show that traders alter their behaviour after signals of their abilities. In a learning by doing model, poor initial performance improves as traders learn how to trade. In learning about ability models, traders receiving positive signals increase their trading activities. While traders in our sample do not pay any fixed trading costs, if effort is also correlated with activity performance should also evolve, positively in a rational model but perhaps negatively in a non-rational one. The early-career performance of long-lived traders is therefore potentially different to their late-career performance.

The results in Table 5 support an evolution of performance, though not in the way the rational learning models predict. Traders who will go on to have long trading careers ($N^* > 50$) tend to have made comparatively small losses and have good success ratios in their first few trades ($N \leq 10$). A relatively large proportion (39%) make money over their first trades. However, performance deteriorates as their experience lengthens and by their late career trades ($N > 50$) they are making large losses per trade. Just 13% of these long-lived traders are making money over their late-career trades. The signals they received early in their careers appear to have encouraged them to continue trading but performance turns increasingly negative later in their career. Such a pattern is inconsistent with a learning by doing model and not easy to reconcile with a conventional rational learning about ability framework. Dorn and Sengmueller (2009) argue that the entertainment value of trading helps explain the high levels of turnover among retail equity traders, but such effects do not make it any easier to reconcile our findings of

deteriorating profitability with rational learning. Some of the forms of entertainment value that have been suggested would tend to be reduced by unprofitable trades (e.g. feelings of accomplishment or the ability to boast about successes), but the entertainment value of trading would actually need to increase over time to explain why rational agents should continue trading as profitability declines.

A deterioration in ability is, however consistent with an irrational learning model where, for whatever reason, traders interpret signals as positive and adapt their trading strategy but in doing so take larger but inappropriate trades, resulting in increasing losses. We explore this more formally in the following section.

3.5. Learning

To get a better understanding of the learning processes we estimate a series of regressions relating performance to experience. We begin with a simple model that we estimate at a monthly frequency:

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

where the dependent variable is a measure of monthly performance. We use two different measures of performance: A success indicator that takes the value one if positive total profits are made over the month, and zero otherwise (“*Win*”); and a proxy for return equal to total profits over the month divided by the volume traded during the month multiplied by 100 (“*Ret*”). We use monthly performance measures in this section for two reasons. First, the *Ret* measure is extremely volatile when measured at higher frequencies since it frequently apportions marked-to-market profits on carried-over positions to days with low or no trading activity. Aggregating over a month, while still imperfect, better balances trading profits and traded volumes. Second, later in this section we use a Heckman correction for selection bias that entails estimating a separate selection equation for each period in the sample. Working at a daily frequency this would mean estimating almost 600 selection regressions and including almost 600 additional terms in the second-stage regression. While arguably this would be possible due to the large size of our dataset, the use of daily selection equations would imply that we capture an individual’s decision to trade on day t . This is likely driven by many factors outside our dataset such

work commitments, vacations, or illness. Experimentation suggests that daily selection equations work very poorly. Conversely, moving to a monthly frequency reduces the number of selection equations needed and improves our ability to capture factors driving the decision to trade in any given month.

The key explanatory variable *Experience* is a proxy for investors' trading experience. We proxy for trader experience by either the number of days on which the trader has been active ("*LifeSpan*") or by the cumulative number of trades made ("*TradeCount*"). Both variables are calculated using the history of the traders until the start of the month in which performance is calculated. We allow traders to learn faster during their earlier years by including *Experience*². We also include a set of control variables, $X_{i,t}$, namely lagged log trade volume and lagged log number of deals, plus month-year time effect dummies. Some of the traders were present in the database at the start of our sample. We therefore under-estimate their experience using both of our measures of trading experience. Such traders are removed from the sample and we only use traders who are identified in the database as new entrants (around 62,500 individuals).²⁴

A positive coefficient on *Experience* in this regression would reflect learning by doing effects if attrition is exogenous and traders are not heterogeneous. Since both of these assumptions are strong and, given our earlier results, unlikely to hold for our data we relax them in subsequent specifications.

< Table 6 about here >

The results of estimating Equation 2 are reported in the first two columns of Table 6.²⁵ Three of the four performance-experience combinations suggest rational learning. However, these OLS results are merely reported to help benchmark our later ones since our previous results suggest that traders are heterogeneous and that attrition is endogenous. Equation 2 would then be misspecified and learning effects are potentially incorrectly estimated. We first deal with the heterogeneity issue.

²⁴The database flags the initial deposit made by a trader opening a new account allowing us to clearly identify new entrants. We may still underestimate the experience of traders if they have traded elsewhere before opening an account with the company supplying our data. We thank the referee for helping to clarify our thoughts in this area.

²⁵The results reported use OLS but we obtain similar findings if we use logit methods when the *Win* measure, a binary indicator variable, is the dependent variable.

Since trader participation in our sample changes over time, trader heterogeneity may induce cohort effects. For example, if innate heterogeneous ability positively correlates with the number of trades placed by an investor, Equation 2 would confuse heterogeneous but fixed-ability with learning. We can account for unobserved trader heterogeneity by including fixed effects in the regression:

$$y_{i,t+1} = \alpha_i + \beta_1 Experience_{i,t} + \beta_2 Experience_{i,t}^2 + \delta X_{i,t} + \epsilon_{i,t} \quad (3)$$

Accounting for heterogeneity has a substantial impact on our results, as documented in columns 3-4 of Table 6. Most obviously, any evidence of rational learning has disappeared and the performance of traders worsens as their experience grows. The magnitudes of the estimated “learning-to-fail” effects are quite large. A trader with a *LifeSpan* of 100 days is 14% less likely to have a winning month than a complete novice trader. Alternatively, the *Ret* of a trader with 100 days’ experience is 5.6bp lower than that of a novice.

Since we know the decision to quit is related to performance, our fixed effects estimates still potentially suffer from endogenous attrition. To assess whether this is a serious issue for our estimates we use a variant of the Verbeek and Nijman (1992) variable addition test. This entails including a lagged selection term $s_{i,t-1}$ into the fixed effects model, where $s_{i,t} = 1$ if trader i trades after period t and is zero otherwise. The Verbeek-Nijman test statistic is extremely significant in all specifications of the model, implying that selection is important in the sample and that attrition is not exogenous.

We account for both endogenous attrition and cross-sectional heterogeneity using a variant of the Heckman selection model due to Wooldridge (1995). This approach entails estimating a first stage Heckman selection model that predicts which observations will be observable in the second stage learning model for each monthly cross-section. The conditional probability that an individual continues to trade (the inverse Mills’ ratio) of each selection equation is then included in the learning regression model which corrects for the selection issue. The second stage learning model is estimated with fixed effects to account for individual time-invariant performance heterogeneity. Specifically, the second

stage regression we estimate is:

$$y_{i,t+1} = \alpha_i + \beta_1 Experience_{i,t} + \beta_2 Experience_{i,t}^2 + \delta X_{i,t} + \rho_1 I(t=1)\lambda_1 + \dots + \rho_j I(t=j)\lambda_j + \epsilon_{i,t} \quad (4)$$

where $\lambda_1, \dots, \lambda_j$ are the inverse Mills' ratios from the selection models in month 1 to j , and $I(t=j)$ is an indicator variable equal to unity in month j and zero otherwise.

The first stage selection model uses a set of cross-sectional probit regressions to predict whether or not a trader trades in a given month. It is desirable to have at least one instrument in the selection equation to ensure identification.²⁶ The probit regressions include a constant, the relevant performance measure (Win_{t-1} or Ret_{t-1}), linear and quadratic experience terms (with experience proxied by either *LifeSpan* or *TradeCount*), and an instrument similar to that use by Seru, Shumway, and Stoffman (2010) that takes advantage of the demographic characteristics we have available. The instrument is the proportion of all traders from the trader in question's city that choose to trade in that month (*PropActive*), and the argument is that an individual is more likely to trade if his neighbours are trading. This is motivated by the work of Hong, Kubik, and Stein (2004) and Ivkovic and Weisbenner (2007) who claim that there is a social component unrelated to performance that drives at least part of investor trading (while acknowledging that there may be other common factors driving trader activity). Identification relies on this social component.

< Table 7 about here >

The variables included in the first stage selection equations, including the instrument term, are each significant in most of the cross sections suggesting that they work well in explaining the traders' decisions to trade each month. For brevity, we only report pooled first-stage estimates in Table 7, using either *LifeSpan* or *TradeCount* as the experience proxy and *Ret* as the performance proxy. Results for each of the months are qualitatively similar to those reported. As expected, there is strong evidence that as investors perform poorly they cease trading. In particular the estimates on Ret_{t-1} are

²⁶The non-linearity of the inverse Mills' ratio may be sufficient to allow identification. However, the exclusion restriction that follows from including a variable in the first stage regressions but not in the second stage model aids identification.

both positive and very strongly significant. Consistent with earlier findings, this shows that low ability traders learn about their inherent ability by trading and eventually cease. More successful traders continue to be active. The other coefficient estimates reported in Table 7 are also sensible: traders are more likely to trade if they have more trading experience though at a decreasing rate, and, importantly, the coefficient estimate for our instrument is statistically significant and of the predicted sign. Specifically, a higher proportion of traders in the same city who are active in a given month increases the probability that the individual will trade in that month.

The key coefficient estimates from Equation 4 are given in the final columns of Table 6. We note first that the joint test of $\rho_j = 0$ for all months is strongly rejected in the second stage regression, indicating the presence of important selection bias effects consistent with the Verbeek-Nijman results (unreported p-values are all < 0.001). Second, all the linear experience coefficients are negative and three of the four are statistically significant. As with the fixed effects estimation, traders appear to perform worse with experience. The quadratic terms are all positive, but are only statistically significant in Panel A where the experience proxy is given by *LifeSpan*.

While the signs of the coefficients are similar to those obtained using simple fixed effects, coefficient magnitudes are smaller once selection issues are taken into account. Based on the Heckman regression estimates, a trader with a *LifeSpan* of 100 days is 6.4% less likely to have a winning month than a complete novice trader. Alternatively, the *Ret* of a trader with 100 days' experience is 3.7bp lower than that of a novice. Proxying experience with *LifeSpan* suggests that performance deteriorates until a trader has been active on around 325 days. By this time his *Ret* performance is 7.1bp worse than a novice (and his probability of a winning month is 8.2% worse). Very few of our traders are so experienced, however, with fewer than 1% having 325 days of trading experience. When we proxy experience with *TradeCount* the linear term bears a negative coefficient that is significant when performance is measured by returns, but not when we use win ratios. The quadratic terms, while positive, are not significant in either regression.

Accounting for both heterogeneity and endogenous attrition leads us to conclude that learning by doing effects are (a) relatively small and (b) negative for the typical lifespans of our traders. Put another way, the performance of traders deteriorates slightly with

experience. We suggest that this small “learning-to-fail” effect is consistent with some form of irrational learning by doing such as naive reinforcement learning or overconfidence. Alternative explanations exist, including cognitive deterioration with age (which we consider below). At this stage we prefer to emphasise that the decline we have identified is quite small, whatever the cause. While learning by doing appears to have a small negative effect, learning about ability effects are much larger.

We have accounted for cross-sectional heterogeneity in performance by including fixed effects in the regressions. However the Heckman selection correction approach pools all traders and so assumes that traders’ decisions whether to trade each month are homogeneous with respect to the driving variables. This is potentially invalid since we have shown in section 3.2 that traders respond to signals differently. In particular, young and old traders respond very differently to experiencing a successful trading day when deciding whether or not to quit. We account for this specific cross-sectional heterogeneity in selection by re-estimating the entire Heckman model (both first and second stages) on sub-samples of data based on the age of traders. Specifically, we rank traders by age and define young traders to be those in the lower two quintiles, and old traders to be those in the upper two quintiles.²⁷ The key second stage coefficients are reported in Table 8. For both sub-samples we observe the familiar pattern of negative coefficients on all the experience terms and positive coefficients on the quadratic terms. For both young and old traders, once we account for performance heterogeneity and heterogeneous attrition, performance initially declines with experience before recovering. Again, however, the recovery in performance only begins after unusually long trading lives. The magnitudes of these effects differ substantially between young and old traders when we proxy performance with the success indicator (*Win*). The probability of having a winning month declines relatively rapidly for young traders before recovering quickly. The effects are more muted for older traders (and are not significant at all if we measure experience using the number of trades made). Conversely, the effect of experience on returns is essentially equivalent across both groups of trader.

The negative relationship between performance and experience is at least as strong for

²⁷As defined, young traders are no more than 27 years old and have an average age of 23. Old traders are at least 32 and are on average 42 years old.

young traders as it is for older traders, and in some cases is much stronger. Without being able to directly test all possible alternatives due to limitations with the data, the results are, we feel, consistent with behavioural explanations. It does not seem unreasonable to think that younger traders with less life experience are more likely to be overconfident or to overemphasise their personal trading experiences relative to more mature traders. Nevertheless, both groups of trader see deteriorating performance with ability.

One alternative that we can address is the concern that the small negative relationship between performance and experience detailed in Table 6 may simply be capturing a decline in cognitive ability with age. There are several reasons we doubt this explanation. First, our sample is relatively short and cognitive decline would need to be rather rapid if we were able to detect it over just thirty months. Second, the average age of traders in our sample is around thirty. The age at which cognitive decline begins is a matter of controversy (see Salthouse (2009) and Schaie (2009)), but this illustrates the fact that any decline in relative young adults is likely to be modest. Third, the results in Table 8 where we disaggregate by trader age are not consistent with such an explanation unless we were to think that cognitive decline is stronger in younger people where the negative relationship is strongest.

The evidence in this section helps differentiate between competing learning hypotheses. Once we account for heterogeneity and selection effects, there is no evidence of positive learning with experience. There is some evidence that retail FX traders with more experience tend to underperform those with less experience. There is weak evidence that performance eventually improves but very few traders in our sample are sufficiently experienced to reap any benefits (and hence our estimates of long-term learning effects are very imprecise). Nevertheless, the magnitude of the irrational learning effects are relatively small. The learning about ability effects we have documented earlier are much more important. Traders mainly learn by realising that they have no skill and hence ceasing trading.

4. Conclusion

We analyse the performance of almost 100,000 retail foreign exchange traders over two and a half years. Our focus is on whether and how traders learn in this previously under-

explored marketplace. In particular, we attempt to distinguish between the three main categories of learning seen in the literature:

1. Rational learning about ability, where traders who are initially uncertain about their innate trading skills update their estimates of their ability based on their performance
2. Learning by doing, whereby traders get better at trading through experience
3. Irrational learning, where as traders become more experienced they update their estimates of their own skill levels or revise their trading techniques in irrational ways, perhaps due to behavioural tendencies such as overconfidence or attribution bias.

We find four main pieces of evidence. First, traders are significantly more likely to cease trading after a day on which they lose money (i.e. after they receive a negative signal about their ability). This is entirely consistent with traders learning about ability. The sensitivity of this decision to quit following a negative signal is much larger for traders who are likely to be learning the most. Novice traders and young traders react most to performance signals, while more experienced traders and older traders, who might have already learned either through being in the market for longer or through their life experiences, react much less.

Second, traders trade both more frequently and in greater volume following a positive signal. This too is consistent with rational learning about ability. Traders may be willing to make expected losses on their initial trades while they either learn how to trade effectively or learn what their innate skill levels are. To minimise costs, however, they should choose to trade only relatively small amounts. As they learn how to trade or as they receive positive signals about their abilities trading becomes more attractive and so they trade more and in larger amounts.

Third, despite the encouraging results for learning models we show that traders - even very experienced traders - perform poorly. This is common to the literature on retail traders' performance. Most retail FX traders make losses from trading on average, and these losses are not confined to their early trades. This is harder to reconcile with learning by doing or with conventional rational learning about ability models. It is more consistent with irrational learning approaches where traders interpret signals about their abilities in overly positive ways, adapting their strategies but in doing so taking on inappropriate

trades and making increasing losses.

Our final set of results explicitly tests for the relationship between performance and experience. Our empirical approach takes into account two key characteristics observed in the data. First, some traders are clearly better than others and, second, attrition is endogenous and the decision to quit is related to performance. Having corrected for these important effects we demonstrate that the balance of evidence points to a relatively small but statistically significant deterioration in performance with experience. We interpret this as traders learning to fail due to irrational learning.

Overall then, we find evidence supporting the hypothesis that traders learn about their abilities by trading, though this may not be fully rational. We also find evidence of irrational learning with experience. While some retail traders learn that they should not be active in this market and quit following losses, others irrationally continue to trade and simply make more losses. This suggests that regulators' concerns about this sector of the foreign exchange market have considerable justification.

- Abbey, B. S., and J. A. Doukas, 2014, “Do individual currency traders make money?,” *Journal of International Money and Finance*.
- Agarwal, S., J. C. Driscoll, X. Gabaix, and D. Laibson, 2008, “Learning in the Credit Card Market,” NBER Working Papers 13822, National Bureau of Economic Research, Inc.
- Arrow, K. J., 1962, “The Economic Implications of Learning by Doing,” *Review of Economic Studies*, 29(3), 155–73.
- Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean, 2011, “Do Day Traders Rationally Learn about their Ability?,” *University of California Davis, working paper series*.
- Barber, B. M., Y.-T. Lee, Y.-J. Liu, and T. Odean, 2014, “The cross-section of speculator skill: Evidence from day trading,” *Journal of Financial Markets*, 18(0), 1 – 24.
- Barrot, J.-N., R. Kaniel, and D. A. Sraer, 2014, “Are retail traders compensated for providing liquidity?,” *Simon Business School Working Paper No. FR 14-15*.
- Ben-David, I., J. Birru, and V. Prokopenya, 2014, “Uninformative Feedback and Risk Taking,” *Fisher College of Business Working Paper*, (2014-03), 17.
- Calonico, S., M. D. Cattaneo, and R. Titiunik, 2015, “Robust nonparametric confidence intervals for regression-discontinuity designs,” *Econometrica*.
- Chiang, Y.-M., D. Hirshleifer, Y. Qian, and A. E. Sherman, 2011, “Do Investors Learn from Experience? Evidence from Frequent IPO Investors,” *Review of Financial Studies*, 24(5), 1560–1589.
- Dhar, R., and N. Zhu, 2006, “Up Close and Personal: Investor Sophistication and the Disposition Effect,” *Management Science*, 52(5), 726–740.
- Dorn, D., and P. Sengmueller, 2009, “Trading as Entertainment?,” *Management Science*, 55(4), 591–603.
- Feng, L., and M. S. Seasholes, 2005, “Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets?,” *Review of Finance*, 9(3), 305–351.
- Gervais, S., and T. Odean, 2001, “Learning to be Overconfident,” *Review of Financial Studies*, 14(1), 1–27.
- 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(3), 533–47.
- Hieronimus, T. A., 1977, *Economics of futures trading for commercial and personal profit*. Commodity Research Bureau.
- Hong, H., J. D. Kubik, and J. C. Stein, 2004, “Social Interaction and Stock-Market Participation,” *Journal of Finance*, 59(1), 137–163.
- Ivkovic, Z., and S. Weisbenner, 2007, “Information Diffusion Effects in Individual Investors’ Common Stock Purchases: Covet Thy Neighbors’ Investment Choices,” *Review of Financial Studies*, 20(4), 1327–1357.

- Kaustia, M., E. Alho, and V. Puttonen, 2008, “How Much Does Expertise Reduce Behavioral Biases? The Case of Anchoring Effects in Stock Return Estimates,” *Financial Management*, 37, 391–412.
- King, M. R., and D. Rime, 2010, “The four trillion dollar question: what explains FX growth since the 2007 survey?,” *BIS Quarterly Review*.
- Linnainmaa, J. T., 2011, “Why Do (Some) Households Trade So Much?,” *Review of Financial Studies*, 24(5), 1630–1666.
- List, J. A., 2003, “Does market experience eliminate market anomalies?,” *The Quarterly Journal of Economics*, 118(1), 41–71.
- Mahani, R., and D. Bernhardt, 2007, “Financial Speculators’ Underperformance: Learning, Self-Selection, and Endogenous Liquidity,” *Journal of Finance*, 62(3), 1313–1340.
- Nicolosi, G., L. Peng, and N. Zhu, 2009, “Do individual investors learn from their trading experience?,” *Journal of Financial Markets*, 12(2), 317–336.
- Odean, T., 1998, “Volume, Volatility, Price, and Profit When All Traders Are Above Average,” *Journal of Finance*, 53(6), 1887–1934.
- Pukthuanthong-Le, K., and L. R. Thomas, 2008, “Weak Form Efficiency in Currency Markets,” *Financial Analysts Journal*, (2008-3), 31–52.
- Salthouse, T. A., 2009, “When does age-related cognitive decline begin?,” *Neurobiology of Aging*, 30(4), 507–514.
- Schaie, K. W., 2009, “When does age-related cognitive decline begin? Salthouse again reifies the cross-sectional fallacy,” *Neurobiology of Aging*, 30(4), 528.
- Seru, A., T. Shumway, and N. Stoffman, 2010, “Learning by Trading,” *Review of Financial Studies*, 23(2), 705–739.
- Verbeek, M., and T. Nijman, 1992, “Testing for Selectivity Bias in Panel Data Models,” *International Economic Review*, 33(3), 681–703.
- Wooldridge, J. M., 1995, “Selection corrections for panel data models under conditional mean independence assumptions,” *Journal of Econometrics*, 68(1), 115–132.

Table 1: Descriptive Statistics

	Mean	25 percentile	Median	75 percentile	Std. Dev.
Active days	45	8	20	51	68
Avg daily trading volume (\$)	n.d.	14,651	41,047	115,291	1,307,496
Avg daily number of trades	n.d.	4	7	12	9.6
Avg daily profit (\$)	n.d.	-19.36	-5.56	-1.48	141
Avg success rate	0.48	0.37	0.50	0.61	0.21

Notes: n.d. denotes not disclosed for confidentiality reasons.

Table 2: Trade Performance Heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sample	Full	Full	Full	$N \leq 10$	$10 < N \leq 25$	$25 < N \leq 50$	$N > 50$	Age Q1	Age Q5
Career success rate	0.415 (0.003)	0.411 (0.003)	0.417 (0.003)	0.157 (0.002)	0.474 (0.004)	0.674 (0.005)	0.895 (0.006)	0.306 (0.005)	0.537 0.008
Year/month dummies	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Experience dummies	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.022	0.024	0.024	0.011	0.029	0.039	0.044	0.016	0.034
Std Dev. of Career success rate	0.179	0.179	0.179	0.307	0.173	0.143	0.117	0.198	0.167

Notes: This table reports the results of regressing a success dummy of trading day t (1 if profit, 0 if loss) on the career success ratio up to but not including time t plus dummy variables, as noted. Standard errors are reported in parentheses under coefficient estimates. Columns (4)-(7) restrict the sample according to trader experience. Column (8) examines traders in the youngest age quintile and column (9) examines the oldest quintile of traders. All standard errors are clustered by trader and are robust to heteroscedasticity.

Table 3: The Decision to Quit

	(1)	(2)	(3)
Success ($t-1$)	0.842 (0.008)	0.840 (0.008)	0.836 (0.008)
Career success rate ($t-1$)	0.848 (0.017)	0.855 (0.0178)	0.875 (0.018)
Year/month dummies	No	Yes	No
Controls	No	No	Yes

Notes: This table reports the results of estimating a Cox proportional hazards model with a success dummy of trading day $t-1$ (1 if profit, 0 if loss) and career success ratio up to time $t-1$ plus dummy and control variables, as noted. The control variables are logged career-to-date $t-1$ cumulated trading volumes and number of trades. Standard errors are reported in parentheses under coefficient estimates.

Table 4: The Effect of Success on Trade Activity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample	Full	$N \leq 10$	$10 < N \leq 25$	$25 < N \leq 50$	$N > 50$	Age Q1	Age Q5
Panel A: Trade volume							
Success ($t-1$)	0.177 (0.002)	0.183 (0.003)	0.177 (0.003)	0.179 (0.003)	0.172 (0.003)	0.194 (0.004)	0.140 (0.005)
Volume ($t-1$)	0.753 (0.001)	0.653 (0.002)	0.736 (0.001)	0.763 (0.002)	0.788 (0.002)	0.716 (0.003)	0.763 (0.003)
Panel B: Trade volume increase indicator							
Success ($t-1$)	0.0576 (0.001)	0.0623 (0.001)	0.0597 (0.001)	0.0592 (0.001)	0.0537 (0.001)	0.0618 (0.001)	0.0454 (0.002)
Panel C: Gap between active days							
Success ($t-1$)	-3.470 (0.026)	-4.662 (0.067)	-4.914 (0.060)	-4.258 (0.0527)	-2.203 (0.027)	-4.687 (0.063)	-2.356 (0.046)
Intensity ($t-1$)	0.024 (0.001)	0.022 (0.002)	0.022 (0.002)	0.025 (0.002)	0.029 (0.002)	0.017 (0.017)	0.039 (0.003)

Notes: Panel A of this table reports the results of regressing the log of the volume of trades on day t on a success dummy of trading day $t-1$ (1 if profit, 0 if loss), a lagged dependent variable, and year-month and experience dummy variables. Standard errors are reported in parentheses under coefficient estimates. Columns (3)-(6) restrict the sample according to trader experience. Column (7) examines traders in the youngest age quintile and column (8) examines the oldest quintile of traders. In Panel B the dependent variable is an indicator variable that takes the value of one if the volume of trades on day t is greater than the volume traded on trading day $t-1$. No lagged dependent variable is included in this case. Panel C repeats this exercise with the difference between active days (measured in days) as dependent variable and with a lagged dependent variable included in the regression. All standard errors are clustered by trader and are robust to heteroscedasticity.

Table 5: Performance and Experience

Lifespan:	$N^* \leq 10$	$10 < N^* \leq 25$	$25 < N^* \leq 50$	$N^* > 50$	<i>Full</i>
Panel A: Average profit					
<i>Full</i>	-30.24	-20.22	-16.88	-19.21	<i>N.D.</i>
$N \leq 10$	-30.24	-11.43	-4.73	-5.18	-12.33
$10 < N \leq 25$		-32.81	-15.24	-13.45	-18.29
$25 < N \leq 50$			-29.91	-15.57	-19.18
$N > 50$				-23.52	-23.52
Panel B: Success ratio					
<i>Full</i>	40.49	47.99	49.91	51.29	50.27
$N \leq 10$	40.49	51.88	53.07	52.06	49.72
$10 < N \leq 25$		42.41	51.56	51.87	49.69
$25 < N \leq 50$			44.88	51.85	50.10
$N > 50$				50.87	50.87
Panel C: Winners					
<i>Full</i>	9.77	14.28	16.08	15.79	<i>N.D.</i>
$N \leq 10$	9.77	35.50	39.98	38.65	28.13
$10 < N \leq 25$		9.69	32.70	34.06	24.52
$25 < N \leq 50$			10.51	29.50	21.36
$N > 50$				13.39	13.39

Notes: The table reports the average daily profit measured in dollars, the mean success ratio (100.profitable days/total number of active days), and the proportion of traders who are profitable(“Winners”), broken down by trader lifespan in columns (denoted N^*) and experience in rows (denoted N). The final column gives the performance measure of the full sample, unconditional on the lifespan of traders. The first row in each panel gives the performance measure over the full lifespan of the trader. Subsequent rows give the performance measure over the traders’ first ($N \leq 10$), early ($10 < N \leq 25$), intermediate ($25 < N \leq 50$) and late ($N > 50$) trades. *N.D.* denotes not disclosed for confidentiality reasons.

Table 6: Learning Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	FE	FE	Heck	Heck
	<i>Win</i>	<i>Ret</i>	<i>Win</i>	<i>Ret</i>	<i>Win</i>	<i>Ret</i>
Panel A:						
<i>LifeSpan_t</i>	0.689 (0.044)	-0.013 (0.010)	-1.699 (0.075)	-0.668 (0.020)	-0.806 (0.082)	-0.436 (0.025)
<i>LifeSpan_t²</i>	-1.432 (0.158)	0.091 (0.030)	3.107 (0.171)	1.081 (0.043)	1.699 (0.183)	0.668 (0.054)
Panel B:						
<i>TradeCount_t</i>	0.190 (0.010)	0.032 (0.003)	-0.234 (0.040)	-0.094 (0.013)	-0.036 (0.032)	-0.045 (0.011)
<i>TradeCount_t²</i>	-0.056 (0.011)	-0.009 (0.002)	0.080 (0.028)	0.032 (0.010)	0.006 (0.019)	0.014 (0.008)

Notes: Panel A of this table reports the results of estimating equations 2 - 4. The dependent variable in each regression is a measure of monthly performance captured by one of two proxies: a success indicator that takes the value one if positive total profits are made over the month, zero otherwise (*Win*) and a proxy for return equal to total cumulated profits over the month divided by the volume traded during the month multiplied by 100 (*Ret*). The proxy used as dependent variable is given in the third row of the table. The key explanatory variables are *Experience* and *Experience*². In Panel A the proxy for *Experience* is the number of days on which the trader has been active/1000 (*LifeSpan*) while in panel B the proxy is the cumulative number of trades made/10,000 (*TradeCount*). The second row of the table gives the estimation method used (OLS, fixed effects or a Heckman-style model as detailed in the text). Standard errors are given in parentheses below coefficient values. All standard errors are clustered by trader and are robust to heteroscedasticity. All regressions also include month-year dummies, lagged monthly trading volume and lagged monthly number of trades. The Heckman-style model also includes separate inverse Mills ratios estimates for each month in the sample.

Table 7: First Stage Heckman Selection Regressions

	<i>LifeSpan</i>	<i>TradeCount</i>
Ret_{t-1}	3.359 (0.025)	3.133 (0.024)
$LifeSpan_t$	9.881 (0.106)	
$LifeSpan_t^2$	-16.871 (0.393)	
$TradeCount_t$		3.269 (0.042)
$TradeCount_t^2$		-0.960 (0.023)
$PropActive_t$	1.192 (0.070)	0.973 (0.070)
Pseudo- R^2	0.100	0.075


Notes: This table reports results of pooled selection regressions. The dependent variable is an indicator variable that takes the value one if the trader was active during month t , and zero otherwise. Traders enter the sample in their first active month. The explanatory variables are lagged performance (proxied by Ret_{t-1}), experience to the start of month t proxied by either $LifeSpan$ or $TradeCount$, quadratic experience terms, and the proportion of traders in the trader in question's city that trade in month t ($PropActive_t$). Regressions also include month-year dummies.

Table 8: Learning Regressions

	(1)	(2)	(3)	(4)
	<i>Win</i>	<i>Win</i>	<i>Ret</i>	<i>Ret</i>
	Young	Old	Young	Old
Panel A:				
<i>LifeSpan_t</i>	-1.417 (0.165)	-0.600 (0.105)	-0.448 (0.052)	-0.451 (0.032)
<i>LifeSpan_t²</i>	3.196 (0.411)	1.200 (0.215)	0.771 (0.119)	0.642 (0.065)
Panel B:				
<i>TradeCount_t</i>	-0.340 (0.096)	-0.039 (0.036)	-0.066 (0.025)	-0.053 (0.012)
<i>TradeCount_t²</i>	0.202 (0.090)	0.004 (0.016)	0.042 (0.023)	0.016 (0.008)

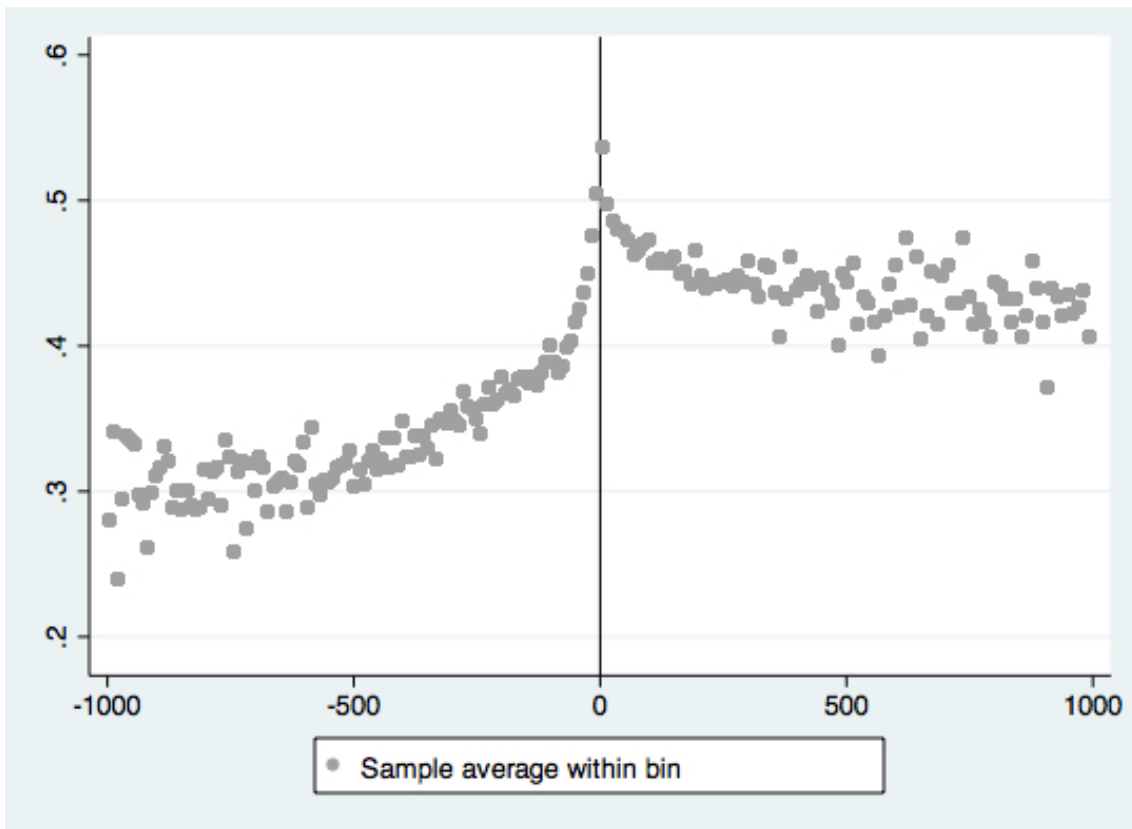
Notes: Panel A of this table reports the results of estimating equation 4 using samples of younger and older traders. The dependent variable in each regression is a measure of monthly performance captured by one of two proxies: a success indicator that takes the value one if positive total profits are made over the month, zero otherwise (*Win*) and a proxy for return equal to total cumulated profits over the month divided by the volume traded during the month multiplied by 100 (*Ret*). The proxy used as dependent variable is given in the second row of the table. The key explanatory variables are *Experience* and *Experience*². In Panel A the proxy for *Experience* is the number of days on which the trader has been active/1000 (*LifeSpan*) while in panel B the proxy is the cumulative number of trades made/10,000 (*TradeCount*). The third row of the table gives the sample used for the estimation. ‘Old’ denotes a sample that contains all traders in the oldest two quintiles of the full sample, while ‘Young’ denotes a sample that contains the traders in the youngest two quintiles in the sample. Standard errors are given in parentheses below coefficient values. All standard errors are clustered by trader and are robust to heteroscedasticity. All regressions also include month-year dummies, lagged monthly trading volume and lagged monthly number of trades together with separate inverse Mills ratios estimates for each month in the sample.

Figure 1: The Effect of Success on Propensity to Quit



Notes: This figure plots the sensitivity of the quit decision to a successful trade estimated from a cross-sectional model. The sensitivities are estimated for all traders using their i -th trading day performance. The horizontal axis gives the experience i of the traders, and the vertical axis gives the point estimate of the sensitivity. We plot the evolution of the sensitivity for all traders and for younger and older traders.

Figure 2: The Effect of Profit and Loss on the Probability of Increased Volume



Notes: This figure plots the mean probability that the trading volume on trading day $t + 1$ increased relative to day t (vertical axis) against the profit and loss in dollars recorded on day t (horizontal axis).