The Dedicated and the Dabblers: How Social Interaction Propagates Active Investing

Rawley Z. Heimer* and David Simon†

Brandeis University, International Business School

August 2, 2011

Abstract

This research presents empirical evidence of the propagation of active strategies within a network of retail traders. We test a model based on Hirshleifer (2010) which demonstrates that social interactions contribute to the growth of active strategies. Using new proprietary data compiled through a social network for foreign exchange traders, we verify the model by showing that the willingness of traders to contact other traders is increasing in their short-term returns while trading intensity is increasing in the performance of those from whom they receive communications.

*Corresponding Author: Rawley Z. Heimer, Brandeis International Business School Mailstop 032, P.O. Box 549100, Waltham, MA 02454, USA. e-mail: rheimer@brandeis.edu.

†The authors are grateful for support from Brandeis University and to faculty for advice. We thank the operators of the social network for providing us with their data and Alex Dusenbery for assisting us with the database. We also thank Nicole S. Ballenger, Daniel B. Bergstresser, and Carol L. Osler for comments. This version is preliminary and incomplete. All errors are our own.

1 Introduction

This paper provides evidence that social interactions contribute to the growing popularity of active investment strategies. Our results support a model based on Hirshleifer (2010) in which strategies are transmitted by communications among investors in a social network. Traders with good short-term performance are more likely to initiate communications with others and share their investment activity. The better the initiators's recent performance, the more likely the recipients are to increase their own trading activity. Recipients, who only observe the sender's short-term returns, misguidedly adopt an approach to trading that is more intensive but not more profitable. All else being equal, these two results imply that over time the population of traders shifts towards active investing.

It is vital to develop a deeper understanding into the factors that encourage active investors, a type of market participant illuminated in Barber and Odean (2000), since their level of participation in financial markets can have profound effects on both their own welfare and market characteristics. Barber, Lee, Liu and Odean (2009) find that Taiwan's retail investors under-perform the market by 3.8 percent and accumulate losses that amount to 2.2 percent of Taiwan's GDP annually. With regards to financial market outcomes, if active investing is analogous to noise trading, models such as DeLong et al. (1990) point to its ability to impact liquidity and volatility. Among the empirical studies, Foucault, Sraer and Thesmar (2011) find that increasing the cost associated with active retail trading on Euronext Paris reduces the volatility of daily returns by about a quarter of its standard deviation, while Barber, Odean, and Zhu (2009), Kumar and Lee (2006), and Hvidkjaer (2008) document that trades of individual investors tend to be correlated and may affect asset prices.

To test the validity of our hypothesis that social interactions contribute to the growth of active strategies, we introduce new data from a sample of retail foreign exchange traders who are members of a social network we call myForexBook. Prior to joining the social network, users must have an open account on one of roughly 45 online brokerages from which myForexBook collects trading activity in real-time. The database contains the detailed trading history and social interactions of more than 5,500 traders. It includes over two-million time-stamped trades and over 140,000 time-stamped messages and friendships, the majority of which occurred between early-2009 and December, 2010, allowing us to identify clear links between trading and social activity.

The retail foreign exchange market, which did not exist even a decade ago, has tremendously grown since the advent of online trading.[1] According to King and Rime (2010), worldwide retail foreign exchange trading volume grew over seventy percent during 2007 to 2010 and now exceeds \$125 to \$150 billion per day, roughly the same as daily turnover on the entire NYSE family of stock exchanges (NYSE, Arca and Amex). The venue compares favorably to other asset classes, even the most liquid NYSE stocks, since it offers practically unlimited liquidity, tight spreads, and more than 100x leverage.

We first verify that the individuals in our dataset are suitable for testing the hypothesis that

[1] Since real-time margining is crucial for the leverage offered by retail foreign exchange trading, this type of trading was virtually impossible before electronic platforms and the internet. With regards to online trading and stock markets, Choi, Laibson and Metrick (2002) document that the internet has increased trading frequency among its participants. Similarly, Barber and Odean (2002) find that investors who switched to online trading trade more actively. Barber and Odean (2001) suggest that the internet lowers both fixed and variable costs of trading which would imply increased opportunity for entry.

social interactions stimulate active investing. We document heterogeneity in individual trading intensity, the frequency with which one trades, and classify traders into two groups: the "Dedicated" and the "Dabblers." Dedicated traders invest substantial time and resources in foreign exchange trading. They trade several times daily, sometimes even several hundred times per day. The Dabblers, by contrast, trade less than once a day on average. The greater commitment of Dedicated traders is also manifest in their larger initial capital base and their persistence in trading despite short-term losses. Notably, the Dedicated and the Dabblers do not differ in their trading experience which means that we observe individuals who have traded with low frequency for several years. Considering two types of market participants is further justified anecdotally. Whether or not one can earn a living by trading foreign exchange is often a topic of conversation among users of myForexBook and many claim to treat trading like a full-time job.

Despite clear differences in their commitment to trading, both groups are unprofitable losing on average \$6.20 per trade. While losses per trade are slightly smaller for the Dedicated than for the Dabblers, the Dedicated lose more overall because they make more trades. Nevertheless, both groups make more winning than losing trades. For Dedicated traders, 64.5 percent of the trades are profitable versus 57.4 percent for the Dabblers. The volatility of weekly returns is more than 50 percent greater for Dedicated traders.

We present novel empirical findings on the relationship between social interactions and investing. We find that traders are more likely to initiate communication with others after having experienced strong short-term gains. We suggest two candidate explanations for this relationship. As in Hirshleifer (2010), individuals may exhibit "self-enhancing transmission bias" or the tendency to broadcast one's successes while downplaying their failures. Secondly, traders may be rationally aware of the potential benefits of connections made through the social network and have incentive to signal only their best performance to others.

Traders who receive communications increase their trading intensity in response to hearing from those whom have had recent success. A one percent increase in returns by the initiator of communications is associated with about a two percent increase in the number of trades issued by the receiver. This result is in accordance with a body of literature suggesting that individuals tend to choose investments that have performed well in the recent past (Bernatzi, 2001, Choi, Laibson, and Madrian, 2010, and Barber and Odean, 2002). We also find that receiving communications from others reduces one's likelihood of quitting trading.

To analyze the long-run implications of these results we present a population evolution model which shows that this impies the average trader will adopt increasingly active strategies with a rising variance of returns. As predicted, we find that average trading intensity and the variance of returns have both increased over time among participants in the social network. We then address potential concerns over the channels of communication within the network and attempt to rule out other explanations for these trends. Taken as a whole, our analysis supports Hirshleifer's (2010) conclusion that social interactions propagate active investing.

There is substantial evidence that participation and investor behavior in financial markets are influenced by social interaction (Shiller, 1984, 1989, and, Shiller and Pound, 1989). Hong, Kubik, and Stein (2004) and Kaustia and Knüpfer (2011) show that social interactions promote stock market participation with the latter showing good returns stimulate entry. Among mutual fund managers, Hong, Kubik, and Stein (2005) demonstrate that portfolios exhibit higher correlation if they are from the same town while Cohen, Frazzini, and Malloy (2008) show that they place greater bets on firms whose board members are from their education network. Correlation across investments in retirement accounts are also observed by Madrian and Shea (2000) and Duflo and Saez (2002, 2003). Researchers document that investors are influenced by the investment decisions of others including famous investors like Warren Buffett (Sandler and Raghavan, 1996), insiders (Givoly and Palmaon, 1985), and readers of the Wall Street Journal's Dartboard column (Barber and Loeffler, 1993). Most similar to our research, Shive (2010) uses an epidemic model and data on Finnish stockholdings to study how social contact can predict investor trading. A common thread among the literature on the relationship between social interaction and investing is that it relies on proxies such as geographical proximity to infer variation in the level of communication about investments. This paper enhances the body of evidence by examining incidences of observed communications between investors.

Our study extends the analysis of social forces in two directions. First, it offers a new explanation for the over-trading puzzle documented in Barber and Odean (2000) and Barber et al. (2009) whereby individual investors trade actively and lose on averagerelative to passive benchmarks. The most commonly cited explanation of this phenomena is that they are overconfident.[2] (DeBondt and Thaler, 1995, and Bénabou and Tirole, 2002, among others). Second, our study explores more deeply the little-known world of day-traders as roughly 90 percent of positions in our sample are closed within a day. The literature on day-trading is limited due to a paucity of detailed datasets. Most recent papers confirm our finding that day-traders earn negative excess returns (Odean, 2009, Jordan and Diltz, 2003, and Linnainmaa, 2005, 2010). Among the few exceptions, Mizrach and Weerts (2009), relies on trades that were claimed by chatroom participants which likely adds significant upward bias. Harris and Schultz (1998) find that investors at two day-trading firms are profitable on aggregate yet their small sample may suffer from survivorship bias.

The paper is organized as follows. Section 2 presents a model that shows how social interactions can exacerbate trends towards active investing. Section 3 describes the social network and our proprietary data. Section 4 details our methodology for testing the model empirically and contains our results. Section 5 concludes.

2 The Model

This section presents a model of how interactions within a social network can influence the trading behavior of others and in turn shape the population of traders in a market. We highlight the key features of a model proposed by Hirshleifer (2010) in which communications between traders can exacerbate a trend towards trading or being dedicated to beating the market despite having no such success. Furthermore, we show that these trends hold on average so long as (1) the propensity to initiate communications is increasing in returns, (2) receivers of communications increase their trading intensity in response to hearing of higher returns, and (3) the volatility of returns for "Dedicated" traders are greater than those for "Dabblers".

[2]Although it is possible that social interactions contribute to overconfidence or visa versa.

2.1 Investor Communications

Our model consists of a population of traders that enter as one of two types, Dedicated or Dabblers, whom we denote DE and DA respectively. Dedicated traders are the more active, hands-on traders and their higher trading intensity is also associated with strategies that are of higher variance. The positive correlation between trading intensity and variance of returns is not a certainty; however, there are many examples for which this would be the case empirically, such as in comparing one investor who holds a fund which tracks the S&P 500 and another who day-trades equities.[3]

When two traders interact one may reveal their investments and if they do, they place emphasis on their greatest successes. This relationship can be generalized by the following linear sending function:

$$s(R_i) = aR_i + b \tag{1}$$

in which s is the probability that an individual discusses their strategies and returns, R is the return of the strategy they transmit, and $i \in \{DE, DA\}$ is the trader's type. We assume $a > 0$ and since b is the baseline probability of transmission it therefore must be that $b \in [0, 1]$. The positive relationship between short-term returns and revealing strategies can be justified in a few ways. For one, Hirshleifer (2010) and Bénabou and Tirole (2002) draw from psychological research in which individuals have a tendency to attribute their successes to their own skill while blaming their failures on poor luck. This motivates them to broadcast their successes while remaining mute about their failures. Another justification is that individuals recognize the advantages to maintaining strong placement in the topology of investor networks and have incentive to signal only their best performance to others. For example, a bad tip on a stock may lead to being blackballed from the inner circle of an investment club. The relationship need not always be positive. For one, new traders are likely to engage socially with others (such as joining online social networks like myForexBook) with the intention of learning from the more experienced. This would suggest that the worse they perform the more likely they are to seek advice. Another example is that investors may be more tight-lipped about their strategies if doing so has an impact on prices and thus their profitability.

When a trader learns of the returns of the person with whom they are in contact, they exhibit some probability, $r(R_i)$ of adopting the sender's strategy and being converted to the sender's type:[4]

$$r(R_i) = cR_i + d \tag{2}$$

Here c is positive if individuals are more likely to be swayed by higher returns and d is the baseline probability.[5] There is strong empirical evidence that investors, faced with the extremely difficult task of having to forecast security returns, choose to extrapolate past returns into the future and invest in securities that have recently performed well. Benartzi (2001) finds that the willingness of employees to invest in their own firm in their retirement accounts is increasing in the performance of

[3] This relationship is confirmed in our data. Our sample of Dedicated traders have a significantly higher standard deviation in weekly log returns.

[4] Hirshleifer (2010) uses a quadratic form for the receiver function in order to reflect a greater emphasis on hearing about extreme returns. For simplicity we use a linear form. It does not change the predictions of the model so long as $r'(R_i) > 0$.

[5] It is important to note that the parameters of the model a, b, c and d do not vary by trader type. This is the case so long as (1) traders do not care about or (2) are unaware of the sender's type.

its stock but does not predict future returns. In an experimental study, Choi, Laibson, and Madrian (2010) find that investors choose high-fee over low-fee index funds based on annualized returns. In an analysis of online trading, Barber and Odean (2002) find that early adopters switched to online trading after initial good performance, even if they later traded more actively but with weaker performance.

The probability of the strategy transmitting from a Dedicated sender to a Dabbler receiver is the joint probability of the sender and receiver functions assuming independence:

$$T_{DE,DA}(R_{DE}) = r(R_{DE})s(R_{DE}) \tag{3}$$

By symmetry,

$$T_{DA,DE}(R_{DA}) = r(R_{DA})s(R_{DA}) \tag{4}$$

If the assumptions behind the sender and receiver functions hold and s', $r' > 0$, then it is straightforward to show that $T'_{DE,DA}$, $T'_{DA,DE} > 0$ as well.

2.2 Population Dynamics

In each period, two randomly drawn traders of type i meet and have the opportunity to share their strategies. The population of traders is finite and equal to n, and the fraction of traders f who are of type DE is:

$$f = \frac{n_{DE}}{n} \tag{5}$$

For simplicity, we do not allow traders to exit the market so the fraction of DE and the fraction of DA traders sums to one in every period.

Since homogeneous pairings do not impact the strategies of the traders, we seek to define the probability of drawing one DE and one DA at random. If the probability of first choosing a DE is $\frac{n_{DE}}{n}$ then the probability of drawing a DA is $\frac{n-n_{DE}}{n-1}$. Likewise, the probability of first choosing a DA is $\frac{n-n_{DE}}{n}$ and the probability of following that with a DE is $\frac{n_{DE}}{n-1}$. Together, they yield the total probability χ that a DE/DA pairing is drawn:

$$\chi = \left(\frac{n_{DE}}{n}\right)\left(\frac{n-n_{DE}}{n-1}\right) + \left(\frac{n-n_{DE}}{n}\right)\left(\frac{n_{DE}}{n-1}\right) \tag{6}$$

or,

$$= \frac{2nf(1-f)}{(n-1)} \tag{7}$$

The probability that the number of traders of either type increases by one in any given period is a function of both the probability of drawing a cross-pairing and the probability that the strategy transmits from one trader to the other. We denote the following period with a $*$ and therefore:

$$Pr(n^*_{DE} = n_{DE} + 1, n^*_{DA} = n_{DA} - 1) = \left(\frac{\chi}{2}\right) T_{DE,DA}(R_{DE})$$

$$Pr(n^*_{DA} = n_{DA} + 1, n^*_{DE} = n_{DE} - 1) = \left(\frac{\chi}{2}\right) T_{DA,DE}(R_{DA}) \tag{8}$$

The change in the fraction of DE traders can be defined as: $\triangle f = f^* - f$, which is equal to the set $\frac{1}{n}$ with probability $\frac{\chi T_{DE,DA}(R_{DE})}{2}$, $-\frac{1}{n}$ with probability $\frac{\chi T_{DA,DE}(R_{DA})}{2}$, and 0 with probability $1 - \frac{\chi T_{DE,DA}(R_{DE})}{2} - \frac{\chi T_{DA,DE}(R_{DA})}{2}$. The expected change in the fraction of DE traders in a given period is thus:

$$\left(\frac{2n}{\chi}\right) E\left[\triangle f\right] = E\left[T_{DE,DA}(R_{DE})\right] - E\left[T_{DA,DE}(R_{DA})\right] \tag{9}$$

The intuition behind Equation 9 is that so long as the transition rate from DA to DE is greater than the transition rate from DE to DA, then on average the fraction of Dedicated traders in the market will be increasing.

2.3 Expected Population Trends, Communication, and Idiosyncratic Volatility

In this section, we diverge from Hirshleifer (2010) and present a condition necessary for the population to trend towards Dedicated traders. We show that the average change in the fraction of active investing is positive so long as: (1) on average, the propensity to initiate communications is increasing in returns, (2) receivers of communications increase their trading intensity in response to hearing of higher returns, and (3) the volatility of DE returns is greater than that of DA traders. This setup incorporates the realistic assumption in item (3) above. It suggests that recipients of communications are responding to the right tail of the sender's distribution of returns. Accordingly, DE's are more persuasive since they have more opportunities to broadcast extreme returns.

Suppose that DE and DA traders share some common component to their returns , $\bar{R}$, with $E\left[\bar{R}\right] = 0$ and variance, $\sigma^2_{\bar{R}}$ (as mentioned by Hirshleifer (2010), this could be the market portfolio). Strategies may differ in their sensitivity to the common factor, β_i. There is also an idiosyncratic component to their strategies, ε_i, which is mean zero as well, $E\left[\varepsilon_i\right] = 0$.[6] The variance of the idiosyncratic portion of their trading activities is assumed to be greater for the DE's, $\sigma^2_{DE} > \sigma^2_{DA}$. Therefore, if we assume that these components are uncorrelated and there is no penalty to being a DE trader (see Section 2 of the Appendix for further discussion of the consequences of including a penalty to trading actively), realized returns are as follows:

$$R_{DE} = \beta_{DE} * \bar{R} + \varepsilon_{DE}$$

$$R_{DA} = \beta_{DA} * \bar{R} + \varepsilon_{DA} \tag{10}$$

Substituting the returns structuredepicted in Equation 10, into Equation 9, gives an expression for the expected change in the fraction of DE traders (see Section 1 of the Appendix for the derivation):

$$\left(\frac{2n}{\chi}\right) E\left[\triangle f\right] = ac\left((\sigma^2_{DE} - \sigma^2_{DA}) + (\beta^2_{DE} - \beta^2_{DA})\sigma^2_{\bar{R}}\right) \tag{11}$$

This expression is positive so long as $\sigma^2_{DE} > \sigma^2_{DA}$, $|\beta_{DE}| \geq |\beta_{DA}|$ (or the linear combination of the differences in the expression are greater than zero) , and a and c, the coefficients in the sender and receiver functions, are positive.

The model implies that the fraction of Dedicated traders in a market will increase on average

[6] Note that both $\bar{R}$ and ε_i do not have to be drawn from a normal distribution. As in Hirshleifer (2010), their distributions may be skewed.

provided that their returns have a wider variance than the average market participant and high realized returns bring about conversation and conversion. This result is intuitively appealing: those with more extreme positive outcomes to discuss will be more influential. It further implies that individuals respond to the positive tail of a distribution. They may falsely attribute a few observations as representing the mean of the entire sample or simply have preferences towards these sorts of gambles (Kumar, 2009). Hirshleifer (2010) points out a variety of phenomena which can be explained by the relationship between social interactions and volatility. In the following sections of this paper we examine whether this theory applies in communications between investors.

3 The Data

3.1 A Social Network for Retail Traders

The data were compiled by a social networking website that, for privacy purposes, we call myForexBook. Registering with myForexBook – which is free – requires a trader to have an open account with one of roughly 45 retail specific foreign exchange brokers. Once registered, myForexBook can access a trader's complete trading record at those brokers, even the trades they made before joining the network. New trades are entered via the retail brokerages but they are simultaneously recorded in the myForexBook database and are time-stamped to the second. myForexBook, which began registering users in January 2009, had 5,693 individuals who made at least one trade during our sample period, which extends to December 2010.[7] The database includes daily account balances per user and, after cleaning, 2,149,083 opened positions of which 2,144,357 had been closed.[8] For roughly half (1,041,658) of these trades – those submitted to specific brokers – the data includes order types and unfilled limit orders.

In addition to providing a forum for communication between investors, several of myForexBook's features have the potential to aid trader performance. A trader registered with myForexBook has access to a "Dashboard" web-browser window which shows the news plus information specific to the social network, specifically a "Sentiment Index" which compiles the aggregate positions of the entire network in a given currency pair. Furthermore, once establishing a bidirectional friendship with another member, both users are able to view each others' trading activity in real-time. Both features are portrayed in Figure 1 and the latter ensures that the vast majority of communication between two users in the network allows for the sharing of returns and strategies.[9]

Our data also includes a complete record of activities within the social network, including the times of logins, friendships established, and messages sent. The median user has made 11.0 logins while the mean has 30.8. Similarly, the median user has 8.0 friends while the mean has 20.9 (Table 1). Care should be taken when referring to these numbers owing to the fact that users enter the

[7]In addition to the 5,693 users whose trades we have records for, there are a few thousand additional users of myForexBook who have not made any trades. These users have either found loopholes through which to register with the network such as using a brokerage practice account or they have not issued any trades on their account. These users will sometimes be involved in the social aspects of the network such as sending messages to other users and posting on forums. They are excluded in all analysis involving trading.

[8]Our initial dataset began with 2,177,747 positions opened. We dropped all duplicate observations and what we believed to be mis-entered data. Observations that we considered to be mis-entered were ones in which the size of the position was negative, the position was closed before it opened, or prices that were not consistent with the historical range of the currency pair.

[9]It is important to note that traders are unable to place orders with their broker from myForexBook's website; rather, it may be useful to view simultaneously while trading.

database (and potentially quit trading) at uneven times.

The database also contains information on the characteristics of its members. This information is offered voluntarily, but the non-response rate is only around ten percent on any given question. The median trader in our database is 36.2 years old, has one to three years of trading experience, calls herself a technical trader and lives in either the USA or Western Europe.

With respect to their trading activities, myForexBook users have short holding periods in comparison to equity traders. Roughly half of all positions are closed within an hour and only around ten percent last longer than a day. They tend to concentrate on the most liquid pairs with the most frequently traded pair, the EUR/USD, constituting 34.3 percent of all trades. The mean trade size is US$34,580 and they use 34x leverage on average.

3.2 The Dedicated and the Dabblers

In this section we confirm that the population of traders in our data is suitable for testing Hirshleifer's hypothesis, namely, that they differ in their level of trading intensity and that there are no gains to active trading. We however expand on the theory of Hirshleifer (2010) in that we document other key differences between those classified as the more time-intensive or "active" traders and "passive" traders who are less so . Their differences, namely the size of their initial capital and reluctance to quit after short-term losses, speak to their commitment to trading which motivates us to classify traders as being either "Dedicated" or "Dabblers".

Trading volume ranges widely among myForexBook users (Table 2). Some registered users made only a few trades in total while others traded almost non-stop. A few users placed several hundred trades a day – even occasionally a few thousand trades (presumably using algorithms). Anecdotal evidence confirms that there is substantial heterogeneity in the level of commitment to trading among myForexBook participants. A frequent topic of conversation on the myForexBook discussion forum is whether it is possible to earn a living by trading. The responses vary from those who claim they do so, others who claim they would be able to if they possessed sufficient capital, and others who say it is unrealistic. For the purposes of illustrating and examining their differences we partition the sample into two groups, the Dedicated and the Dabblers, who differ in their level of trading intensity.

Distinguishing these two groups involves a careful balance. Relying solely on the number of trades per individual biases the sample towards those who entered the dataset at an earlier date. Relying instead on the frequency with which individuals trade over-samples individuals who made several trades quickly and then quit. In order to address these concerns, we restrict the Dedicated group according to two criteria: (1) total trades by an individual must exceed the median (128); (2) and the frequency with which they trade during a given week must also exceed the median (32.1).[10] The resulting partition of the sample involves 2,012 Dedicated individuals who made 1,642,262 trades and 3,681 Dabblers who made 506,821 trades.

[10]This is calculated by taking the total number of trades per individual divided by the number of weeks that pass between their first and last trade. This measure incorporates any lengthy absences from trading making those who take them more likely to be Dabblers.

3.2.1 Distribution of Profits

Those who trade the most are not more successful, consistent with the existing literature on active investing. Trades made by myForexBook users are unprofitable on average, losing $6.20 each roundtrip trade. The Dedicated lose slightly less per trade, but more than make up for it in trading activity so they end up losing more overall. However, the median trade books a $0.22 profit since 63.4 percent are profitable after execution costs. The Dedicated do however have a much higher hit ratio per trade than the Dabblers with positive gains on 65.1 percent of their trades versus only 57.8 percent.

In examining profitability per trader we find that 21.0 percent of the total sample and only 17.8 percent of the Dedicated are profitable as of December 2010. The average trader has accumulated $2,335 in losses while the average Dedicated trader has lost $4,776. The 95 percent confidence interval for cumulative profits of individual traders is [$-11,751; $1,382]. Consistent with our vision of the traders as aggressive risk-takers, the standard deviation of weekly returns to Dedicated traders, 61.2 percent, is statistically higher than the corresponding variance for the Dabblers, 47.4 percent. 75.7 percent of Dedicated traders have negative skewness of weekly returns versus 64.0 percent of the Dabblers.

3.2.2 Starting Capital

We find that at least some differences between the groups can be accounted for by different levels of initial investment. As shown in Table 4, the median starting balance among myForexBook traders is US$983. This is substantially lower than Finnish day traders in Linnainmaa (2003) where the median is €17,525, or approximately US$25,000. Dedicated traders have a median starting account balance of $1,938, compared to $612 for Dabblers. The mean for both groups is substantially higher, $8,512 for the Dedicated and $1,101 for the Dabblers. A student's t-test indicates that the difference is significant at the one percent level.

3.2.3 Trader Lifespan

Another substantial difference between the two groups is their reaction to large losses. Table 5 displays results from estimating Cox-proportional hazard models in which the regressors are zero-one indicators for the decile of weekly returns. Consistent with Linnainmaa (2005), we proxy for having left the market if a trader has been inactive for the last month of the dataset. If a user is found to have quit trading then we say they quit at the time of their last observable activity in the dataset. According to this definition, roughly 75 percent of all participants in our sample quit trading. This fact is not surprising considering that the mean trade is unprofitable regardless of user type. Overall, the Dedicated are slightly more likely to continue trading than the Dabblers, but all of this difference is eliminated if the trader makes it past two weeks.[11]

The results from our tests suggest that for both Dedicated and Dabbling traders a week of good performance reduces the probability of quitting; a week in the highest decile of returns reduces the probability by roughly 40 percent. The Dedicated and the Dabblers however react differently to poor performance. While a performance in the lower deciles for the Dabblers increases the likelihood

[11] When plotting hazard rates we find very little difference between Dedicated and Dabblers in their underlying probability of quitting over time. All of the difference is eliminated when excluding traders who failed to last past two weeks.

of quitting by anywhere from 20 to 60 percent, it has little to no effect on the likelihood of quitting for the Dedicated. Attempts to account for this difference by including proxies for sunk costs such as their initial balances failed to change these results.

4 Empirical Analysis

In section 2, we present a model that shows how communication among investors can shift the average market participant towards strategies that are characterized as being active and of higher variance. We show that this holds, all else equal, so long as (1) the propensity to initiate communications is increasing in returns, (2) receivers of communications increase their trading intensity in response to hearing of higher returns, and (3) the volatility of returns for those who are characterized as being active or dedicated traders are greater than those for whom are not. In the previous section, we introduce data suitable for testing our hypothesis that social interactions propagate active investing. Since we have already confirmed (3) when introducing the data, we are tasked with verifying (1) and (2) in the section that follows.

4.1 The Sending Function

In order to confirm that traders are more likely to initiate communications the greater their returns, we use our data to generate weekly[12] returns per individual and indicator variables for whether or not individual i initiated communication with another member of the social network via a user-message. Weekly returns R in time t are defined as:

$$R_{i,t} = \log\left(\frac{V^e}{V^b}\right) \tag{12}$$

where V^b is the balance at the beginning of the week and V^e is the end of week balance (excluding net deposits) sampled between consecutive Saturdays at midnight, GMT.

Table 6, column I, displays results from estimating a logit model of the form in equation 1 in which the dependent variable is an indicator variable for having sent a message during the week and the dependent variable is weekly returns. The model predicts that a 5.9 percent increase in weekly returns is associated with about a ten percent increase in the probability of sending a message to another user. This increases the baseline probability by about 20 percent. Since the standard deviation in weekly returns is 54.4 percent, it implies a significant variation in messages sent by performance.

We also confirm the presence of a positive relationship between sender returns and the *number* of messages sent. In the second column of Table 6, we use OLS to regress the number of messages sent on log returns of the sender conditional on having sent at least one message. The relationship is positive, but insignificant. The lack of statistical significance may be caused by using OLS to predict count data. To account for this bias, in the third column of Table 6, we present results from estimating the same relationship using a zero-truncated Poisson regression. The coefficient in this specification is again positive, but is now strongly statistically significant implying that the better an individual's returns the more communications they issue.

[12]Considering that much of the activity in this market centers around the release of economic news and that weekends are comparably silent, we believe that week-to-week returns best capture the mindset of these traders.

Upon examining how the relationship between returns and communications varies by trader experience, we find evidence that the structure of the network plays an important role in the propensity to communicate. The most experienced traders – those who claim to have been trading for at least four years – display the strongest tendency to initiate social contact following weeks of good performance (Table 5). Traders at the center of the distribution (one to three years) also have a positive coefficient, but it is smaller and only significant at the ten percent level. Those with the least amount of self-declared experience (zero to one years) when joining the network have a negative and insignificant coefficient. Since the least experienced traders perform significantly worse than other groups, this might be a sign that beginner traders send messages seeking advice. On the other hand, the more experienced traders may be attempting to gain a following within the network and thus strategically communicate only after good returns.

4.2 The Receiving Function

In this section we verify that traders increase their trading intensity in response to hearing from individuals who have had good returns. In the model presented in Section 2, we include the simplification that there are only two types of traders, Dedicated and Dabblers, who differ in their level of trading intensity. Conversion between the two types occurs through communications that are instigated by good short term performance. Identifying incidences of conversion from Dedicated to Dabblers in our dataset is cumbersome owing to the fact that trading intensity of individuals is not a binary variable and highly dependent on our ad hoc criteria for distinguishing between trader types. We therefore proxy for conversion to active investing by calculating the number of trades issued in a given week by the recipient of communication.

We are confronted with two challenges when attempting to identify the empirical relationship between sender returns and recipient activity: (1) how does an individual respond to receiving messages from more than one individual in a given period and (2) what unit of measurement for presenting one's returns does an individual respond to most? In order to address the first issue we calculate the max, mean, and sum of sender returns and estimate the effect of each separately. To combat the second complication, we calculate dollar returns rather than the specification for log returns presented in equation 12. Conversations about dollar returns are presumably more salient to individuals. Furthermore, responding to log returns requires a recipient of communications to have prior knowledge of the initiator's opening balances, a proposition we assume unlikely.

Table 8 displays results from estimating the receiver function, equation 2. In this model, the dependent variable is the log number of trades issued in the week (or the week after) the trader received and read at least one user message. The independent variable is log dollar returns of the sender. We consider the possibility that there may be considerable lag between receiving the message and reading it; therefore, the number of trades issued by the receiver refers to the week in which it was read.

We find that a one percent increase in sender returns is associated with about a two percent increase in the number of trades issued in that week. The effect may appear small, but since the mean number of trades per week is 25 (excluding weeks of inactivity) then a two percent increase extrapolated out to an entire year of trading is an additional 26 trades per individual. This result holds in both the week the individual receives the message and the week after with the strongest effect being on the latter. It is interesting to note that the max and sum of the sender's returns are associated with increased trading by the receiver while the mean of the sender's returns is less strongly correlated. The coefficient in this specification (Table 8, column IV) is smaller than all others and only significant at the ten percent level. Since low returns bring down the weekly averages we calculate, this result may indicate that receiving communication from individuals with poor performance can offset some of the increase in trading intensity brought about by hearing of good returns.

We find that not only are traders more likely to increase their trading intensity upon receiving contact from traders with recent strong performance, but they are also less likely to quit trading when contacted even after controlling for realized returns. In Table 9, we present hazard rates from including indicator variables in the analysis described in Section 3.2.3. The event in the survival analysis is an indicator variable for whether or not an individual quit trading. The independent variables are indicators for whether or not the trader initiated or received and read communications from another individual during the week in question.

We find that those who receive communications are less likely to quit trading while those who initiate them are more likely. The latter result appears to be evidence against the relationship between social interactions and investing; however, this empirical finding does not imply causality and may reflect other factors. We suspect that the decision to quit trading is associated with a greater propensity than the average week to contact other traders since individuals may be motivated to maintain ties when leaving the network. For example, a trader may be planning a move to a different asset class and wishes to remain in contact should the receiver change as well. She may also wish to maintain contact should she decide to return to trading at some point beyond our sample.

The model further suggests that individuals are about 30 to 40 percent less likely to quit trading after receiving communications. This result reinforces the findings of Hong, Kubik, and Stein (2004) and Kaustia and Knüpfer (2011) that social interactions promote market participation. While their results specifically refer to market entry, our results bolster their argument by examining the rate of attrition. In considering its implications on the average market participant, if we were to consider a dynamic setting of the model proposed in Section 2, in which the population of traders includes entry and exit, incorporation of this finding would point towards further exacerbation of trends towards active investing.

4.3 Social Networking and Active Investing Over Time

The model presented in Section 2 suggests that communications between investors can lead to the growth of active investing. We find that features consistent with the predictions of the model are present in our data.

First, we verify that the average participant in the social network has increased their trading intensity. This requires us to determine which individuals in our sample are participants in the market at any given time t. Accordingly, we define a user's time of entry as their first observable action in the dataset and quitting is defined as in Section 3.2.3. This means that the total number of surviving users in our dataset at any given time t is derived as follows:

$$survivors_t = \sum_{t=1}^{T} (entrants_t - quitters_{t-1}) \tag{13}$$

We then calculate the average number of trades issued per surviving user each month as: $\frac{\#\, trades_t}{survivors_t}$. This measure, rather than the number of trades over the the number of users who issued them in a given month, incorporates individuals who take breaks from trading. The corresponding time series is plotted in Figure 2. We find that the average trading intensity per myForexBook user has increased over the course of the sample from roughly 40 trades per month for most of 2009 to roughly double that by late 2010.

We also find that average volatility of returns has increased among participants. We regress the standard deviation of log weekly returns against time (Figure 6) finding that it increases by about 0.2 percent (statistically significant at the one percent level) per week over the life of the social network. This implies an increase in the standard deviation of around 20 percentage points in less than two years.

4.4 Caveats and Discussion

In this section we address three potential concerns that would either offer an alternative explanation of our findings that the average market participant possesses more active strategies or weaken our assertion that social interaction is contributing to the trend: (1) does communication in the network travel along a channel that would promote active strategies, (2) can uneven entry and exit explain the empirical finding that the average trader has increased their activity over time, and (3) is the level of social networking activity sufficient to sustain these trends?

A key consideration necessary to confirm that social interactions are contributing to the growth of active strategies is to establish that the channels of communication travel in directions that would promote this trading behavior. In the model, communications between individuals of different type, Dedicated and Dabbler, leads to the transmission of active strategies. The probability of the two types communicating with one another is a function of the percentage of each type in the population, but is otherwise random. In reality, individuals make choices about whom to communicate with and if there is a high degree of homophily – the tendency of individuals to bond with those who possess like characteristics – among myForexBook participants then strategies are unlikely to spread. In Figure 4, we plot against time the number of new user friendships established among participants in the social network. While the number of friendships made by the users of the social network is roughly constant over time, the prevalence of Dedicated/Dabbler pairings, roughly half of all friendships, is striking considering that Dedicated traders constitute only one-third of the population. This finding implies a network structure in which the Dedicated establish a central location within the social network and encourage the Dabblers to adopt more active strategies.[13]

Another concern is that uneven entry and exit may explain the time series, Figure 2, showing that the average trader has increased their activity over time. In particular, an influx of Dedicated, high activity traders at the end of the sample period could explain this empirical finding. Contrary to this argument, bias is more likely to run in the opposite direction because Dedicated traders are surely more attuned to media intended to improve trading performance (a notion reinforced by the fact that the average Dedicated is more involved in the social networking aspects of myForexBook) and therefore more likely than the Dabblers to be among the early participants in the network. Secondly, a sharp decrease in Dabblers at the end of the sample could explain the time series, but

[13] Another unexplored possibility is that Dedicated/Dedicated pairings simply reinforce and aggravate bad trading behavior.

this is likely to be offset by new participants. Our belief that entry and exit of individuals is not at the heart of our findings is reinforced when, in Figure 5, we plot entrants and exits of each type and the number of surviving users in the dataset over time. The ratio of Dedicated to Dabblers remains roughly constant over time and while there is a spike in exits among Dabblers at the end of the sample period it is unlikely to discount much of our findings.

One last consideration is that unless the impact on one's trading activity caused by receiving communications about high returns is extremely persistent, then the trend towards active investing will stagnate. Therefore, social networking usage must also have increased over the time frame in question. In Figure 6, we plot the number of logins per user to the social network on a monthly basis, a key proxy for social networking usage, and find that it has nearly doubled over the course of the sample from around five to close to ten.

5 Concluding remarks

Our analysis of a new dataset on the activities of retail foreign exchange traders who are participants in a social network supports our hypothesis that social interactions promote the growth of active investment strategies. We apply a population evolution model in which strategies are transmitted through communications between investors and their adoption is motivated by the promise of high returns. The model predicts that the average individual employs increasingly active strategies so long as (1) the propensity to reveal one's strategies is increasing in realized returns, (2) receivers of communication increase their trading intensity in response to hearing of higher returns, and (3) the volatility of returns for those who are characterized as being active traders are greater than those for whom are not. We confirm the predictions of the model by documenting two novel empirical findings: on average, individual investors are more likely to initiate communications with other investors the greater their returns and they increase their trading intensity upon hearing of good returns.

Our research is the first to use detailed data on communications between investors rather than proxies to document its impact on financial behavior, thereby strengthening the empirical literature on the role of social interactions in financial markets. It provides greater insight into the process of diffusion of strategies and news about returns. Our findings also contribute to the disagreement over how increased flow of information contributes to efficient outcomes. While in most standard theory the flow of information within networks leads to better performance among market participants, we find that communications between investors may reinforce and even promote reckless trading behavior. This is largely driven by bias found among traders in which they develop forecasts of future returns that are merely extrapolations of the recent performance of assets. This leads them to follow strategies with occasional outstanding results, but that are less profitable on average.

A final thought to discuss is that while our analysis considers the influence of peer-to-peer communications the traders we studied are participants in an entire network, one that contains over one-hundred thousand direct linkages between traders. There may be substantial network effects that we fail to account for in this research. For example, our findings may stem from "group-think" among clusters of traders whose activities have become correlated. Traveling down this road may answer questions about the contribution of social interactions within networks to many puzzles of asset pricing including the formation of bubbles, propagation of herding, and attention-grabbing.

Section 1

In this section we derive the result in equation 11.

Substituting equations 3 and 4 into equation 9 yields:

$$\left(\frac{2n}{\chi}\right) E\left[\triangle f\right] = E\left[r(R_{DE})s(R_{DE})\right] - E\left[r(R_{DA})s(R_{DA})\right] \tag{14}$$

$$= E\left[(aR_{DE}+b)(cR_{DE}+d)\right] - E\left[(aR_{DA}+b)(cR_{DA}+d)\right] \tag{15}$$

and since, a, b, c, and d are constants,

$$= acE\left[R_{DE}^2\right] + (ad+bc)E\left[R_{DE}\right] - acE\left[R_{DA}^2\right] - (ad+bc)E\left[R_{DA}\right] \tag{16}$$

Further substituting the returns structure from the equations in 10 into the equation above yields:

$$= acE\left[(\beta_{DE} * \bar{R} + \varepsilon_{DE})^2\right] + (ad+bc)E\left[\beta_{DE} * \bar{R} + \varepsilon_{DE}\right] - acE\left[(\beta_{DA} * \bar{R} + \varepsilon_{DA})^2\right] - (ad+bc)E\left[\beta_{DA} * \bar{R} + \varepsilon_{DA}\right] \tag{17}$$

and since $E\left[\bar{R}\right] = E\left[\varepsilon_i\right] = 0$,

$$= ac\left(E\left[(\beta_{DE} * \bar{R} + \varepsilon_{DE})^2\right] - E\left[(\beta_{DA} * \bar{R} + \varepsilon_{DA})^2\right]\right) \tag{18}$$

After expanding out the expressions in parentheses and zeroing out any term with $E\left[\bar{R}\right]$or $E\left[\varepsilon_i\right]$:

$$= ac\left(E\left[\varepsilon_{DE}^2\right] - E\left[\varepsilon_{DA}^2\right] + (\beta_{DE}^2 - \beta_{DA}^2)E\left[\bar{R}^2\right]\right) \tag{19}$$

Since $E\left[\varepsilon_i^2\right] = \sigma_i^2$ and $E\left[\bar{R}_i^2\right] = \sigma_{\bar{R}}^2$,

$$\left(\frac{2n}{\chi}\right) E\left[\triangle f\right] = ac\left((\sigma_{DE}^2 - \sigma_{DA}^2) + (\beta_{DE}^2 - \beta_{DA}^2)\sigma_{\bar{R}}^2\right) \tag{20}$$

which is what we wanted to show.

Section 2

In this section, we modify the returns structure to include a penalty (or premium) to being a Dedicated trader exactly as suggested in Hirshleifer (2010).

$$\begin{gathered} R_{DE} = \beta_{DE} * \bar{R} + \varepsilon_{DE} - D \\ R_{DA} = \beta_{DA} * \bar{R} + \varepsilon_{DA} \end{gathered} \tag{21}$$

Following the same set of steps as in Section 1 of the Appendix brings us to the result:

$$\left(\frac{2n}{\chi}\right) E\left[\triangle f\right] = ac\left((\sigma_{DE}^2 - \sigma_{DA}^2) + (\beta_{DE}^2 - \beta_{DA}^2)\sigma_{\bar{R}}^2\right) + \left(acD^2 - (ad+bc)D\right) \tag{22}$$

Having already discussed the first term on the right hand side of equation 22, we turn our attention to the second set of outermost parentheses. This term governs how the change in the fraction of Dedicated traders responds to the return penalty (or premium) to being a Dedicated trader and it has the potential to offset any movement in the population towards Dedicated trading. Holding all else equal, since it is quadratic in D, the average change in the fraction is as follows:

$$\begin{gathered} E\left[\triangle f\right] \geq 0 \; if \; D \leq 0 \\ E\left[\triangle f\right] < 0 \; if \; 0 < D < \frac{(ad+bc)}{ac} \\ E\left[\triangle f\right] \geq 0 \; if \; D \geq \frac{(ad+bc)}{ac} \end{gathered} \tag{23}$$

The first line of above is straightforward to explain: if there is a return premium to being a Dedicated trader then the fraction of that type grows. This region of the function has $\lim_{D\to-\infty} E\left[\triangle f\right] = \infty$. The second line defines a positive range for D in which the average fraction of DE traders is trending downwards. It makes sense that if there is a penalty to trading there will be fewer DE's, but when traveling along the function there is a point, $D = \frac{(ad+bc)}{2ac}$, at which the penalty works increasingly less against the trend towards DE trading. Since this is the positive sloped region of the function, we consider an explanation that also includes the last line of 23. This range, $D > \frac{(ad+bc)}{2ac}$, suggests that when D grows larger, $E\left[\triangle f\right]$ does as well. The only appealing explanation is that as D grows larger it becomes prohibitively costly to enter the market in the first place. This is because an increase in D results in a downward shift in the specification for returns, $R_{DE} = \beta_{DE} * \bar{R} + \varepsilon_{DE} - D$. Incorporating market entry and exit could be accomplished by defining some minimum threshold for t period returns above which DE traders participate. It also requires a non-constant population, n, which is beyond the scope of the modeling efforts of this paper.

Regardless of the potential modeling issues surrounding D, we believe that the market in question empirically, retail foreign exchange, is one in which there is a relatively low penalty to being a Dedicated trader and thus unlikely to confound our results. Unstated in Hirshleifer (2010) is that, since there are costs associated with being a trader of any type, D is a relative term which defines the penalty (or premium) associated with being a DE rather than a DA. The term could account for a difference in risk-bearing, total transaction costs (for instance, the spread paid per trade times the number of trades or the account start-up fee), or even opportunity cost. With regards to risk-bearing, since the traders in our dataset chose to enter the market for foreign exchange they are all likely to have preferences towards risk. Transaction costs are also extremely low since retail brokerages usually charge the half-spread which is rarely more than one or two pips per trade on the most frequently traded pairs.

Section 3

If we assume that realized returns are achieved as indicated in equation 10, then it is sufficient to show empirically that $Var\left[R_{DE}\right] > Var\left[R_{DA}\right]$, to demonstrate that $\left((\sigma^2_{DE} - \sigma^2_{DA}) + (\beta^2_{DE} - \beta^2_{DA})\sigma^2_{\bar{R}}\right)$ in equation 11 is positive.

$$Var\left[R_{DE}\right] > Var\left[R_{DA}\right] \tag{24}$$

Substituting in the returns structure from equation 10 into the equation above yields:

$$Var\left[\beta_{DE} * \bar{R} + \varepsilon_{DE}\right] > Var\left[\beta_{DA} * \bar{R} + \varepsilon_{DA}\right] \tag{25}$$

$$\beta_{DE}^2 \sigma_{\bar{R}}^2 + \sigma_{DE}^2 > \beta_{DA}^2 \sigma_{\bar{R}}^2 + \sigma_{DA}^2 \tag{26}$$

$$(\sigma_{DE}^2 - \sigma_{DA}^2) + (\beta_{DE}^2 - \beta_{DA}^2)\sigma_{\bar{R}}^2 > 0 \tag{27}$$

which is what we wanted to show.

References

[1] Werner Antweiler and Murray Z. Frank, *Is all that talk just noise? the information content of internet stock message boards*, Journal of Finance **59** (2004), no. 3, 1259–1294.

[2] Brad M. Barber, *The internet and the investor*, Journal of Economic Perspectives **15** (2001), no. 1, 41–54.

[3] Brad M. Barber, Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, *Just how much do individual investors lose by trading?*, Review of Financial Studies **22** (2009), no. 2, 609–632.

[4] Brad M. Barber and Douglas Loeffler, *The "dartboard" column: Second-hand information and price pressure*, Journal of Financial and Quantitative Analysis **28** (1993), no. 02, 273–284.

[5] Brad M. Barber and Terrance Odean, *Trading is hazardous to your wealth: The common stock investment performance of individual investors*, Journal of Finance **55** (2000), no. 2, 773–806.

[6] Brad M. Barber, Terrance Odean, and Ning Zhu, *Do retail trades move markets?*, Review of Financial Studies **22** (2009), no. 1, 151–186.

[7] Roland Bénabou and Jean Tirole, *Self-confidence and personal motivation*, The Quarterly Journal of Economics **117** (2002), no. 3, 871–915.

[8] Shlomo Benartzi, *Excessive extrapolation and the allocation of 401(k) accounts to company stock*, Journal of Finance **56** (2001), no. 5, 1747–1764.

[9] Werner F.M. De Bondt and Richard H. Thaler, *Chapter 13 financial decision-making in markets and firms: A behavioral perspective*, Finance (V. Maksimovic R.A. Jarrow and W.T. Ziemba, eds.), Handbooks in Operations Research and Management Science, vol. 9, Elsevier, 1995, pp. 385 – 410.

[10] James J. Choi, David Laibson, and Brigitte C. Madrian, *Why does the law of one price fail? an experiment on index mutual funds*, Review of Financial Studies **23** (2010), no. 4, 1405–1432.

[11] James J. Choi, David Laibson, and Andrew Metrick, *How does the internet affect trading? evidence from investor behavior in 401(k) plans*, Journal of Financial Economics **64** (2002), no. 3, 397–421.

[12] Lauren Cohen, Andrea Frazzini, and Christopher Malloy, *The small world of investing: Board connections and mutual fund returns*, Journal of Political Economy **116** (2008), no. 5, 951–979.

[13] Bradford J. Delong, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, *Noise trader risk in financial markets*, Journal of Political Economy **98** (1990), no. 4, 703–38.

[14] Esther Duflo and Emmanuel Saez, *Participation and investment decisions in a retirement plan: the influence of colleagues' choices*, Journal of Public Economics **85** (2002), no. 1, 121–148.

[15] ———, *The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment*, The Quarterly Journal of Economics **118** (2003), no. 3, 815–842.

[16] Thierry Foucault, David Sraer, and David J. Thesmar, *Individual investors and volatility*, The Journal of Finance **66** (2011), no. 4, 1369–1406.

[17] Kenneth A. Froot, *Consistent covariance matrix estimation with cross-sectional dependence and heteroskedasticity in financial data*, Journal of Financial and Quantitative Analysis **24** (1989), no. 03, 333–355.

[18] Dan Givoly and Dan Palmon, *Insider trading and the exploitation of inside information: Some empirical evidence*, Journal of Business **58** (1985), no. 1, 69–87.

[19] Jeffrey H. Harris and Paul H. Schultz, *The trading profits of soes bandits*, Journal of Financial Economics **50** (1998), no. 1, 39 – 62.

[20] Harrison Hong, Jeffrey D. Kubik, and Jeremy C. Stein, *Social interaction and stock-market participation*, Journal of Finance **59** (2004), no. 1, 137–163.

[21] ———, *Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers*, Journal of Finance **60** (2005), no. 6, 2801–2824.

[22] Soeren Hvidkjaer, *Small trades and the cross-section of stock returns*, Review of Financial Studies **21** (2008), no. 3, 1123–1151.

[23] Douglas J. Jordan and J. David Diltz, *The profitability of day traders*, Financial Analysts Journal **59** (2003), no. 6, 85–94.

[24] Markku Kaustia and Samuli Knüpfer, *Peer performance and stock market entry*, Journal of Financial Economics **In Press, Accepted Manuscript** (2011).

[25] Michael R King and Dagfinn Rime, *The $4 trillion question: what explains fx growth since the 2007 survey?*, BIS Quarterly Review (2010).

[26] Alok Kumar, *Who gambles in the stock market?*, Journal of Finance **64** (2009), no. 4, 1889–1933.

[27] Alok Kumar and Charles M.C. Lee, *Retail investor sentiment and return comovements*, Journal of Finance **61** (2006), no. 5, 2451–2486.

[28] Juhani T. Linnainmaa, *The individual day trader*, mimeo (2005).

[29] ———, *Do limit orders alter inferences about investor performance and behavior?*, Journal of Finance **65** (2010), no. 4, 1473–1506.

[30] Brigitte C. Madrian and Dennis F. Shea, *The power of suggestion: Inertia in 401(k) participation and savings behavior*, The Quarterly Journal of Economics **116** (2001), no. 4, 1149–1187.

[31] Bruce Mizrach and Susan Weerts, *Experts online: An analysis of trading activity in a public internet chat room*, Journal of Economic Behavior & Organization **70** (2009), no. 1-2, 266–281.

[32] Terrance Odean, *Do investors trade too much?*, American Economic Review **89** (1999), no. 5, 1279–1298.

[33] Carol L. Osler, *Foreign exchange microstructure: A survey of the empirical literature*, Encyclopedia of Complexity and Systems Science, 2009.

[34] Linda Sandler, *Salomon holders watch for a possible buffet!-ing*, Wall Street Journal (1996), C1.

[35] Robert J. Shiller, *Stock prices and social dynamics*, Brookings Papers on Economic Activity **15** (1984), no. 2, 457–510.

[36] ———, *Market volatility*, vol. 1, MIT Press Books, no. 0262691515, The MIT Press, 1989.

[37] Robert J. Shiller and John Pound, *Survey evidence on diffusion of interest and information among investors*, Journal of Economic Behavior & Organization **12** (1989), no. 1, 47–66.

[38] Sophie Shive, *An epidemic model of investor behavior*, Journal of Financial and Quantitative Analysis **45** (2010), no. 01, 169–198.

Table 1: **Social Networking Summary Statistics**

	Logins per user			Friends per user		
	Total	Dedicated	Dabbler	Total	Dedicated	Dabbler
Mean	30.8	36.7	27.6	20.9	31.9	18.5
Median	11.0	14.0	10.0	8.0	13.0	8.0
STdev	75.7	70.5	78.4	63.3	93.6	46.3
Max	2,723	913	2,723	1,801	1,801	1,004
Min	1	1	1	1	1	1
N users	5,597	1,981	3,616	3,871	1,456	2,415

Note: These statistics are conditional on having made at least one login or, in the second panel, at least one friendship.

Table 2: **Trading Volume per User**

	Number of Positions Opened per User		
	Total	pre-myForexBook	post-myForexBook
Mean	377.5	197.3	276.2
Median	128	65	70
STdev	1,541.7	478.2	1,526.1
Max	97,448	9,202	93,732
Min	1	1	1
N users	5,693	3,913	4,985

Note: This document presents summary statistics on the number of trades issued per user in our dataset. In columns two and three, we partition the data into trades made before and after the user joined myForexBook. All statistics are conditional on having made at least one trade.

Table 3: **Profitability per Trade (US$)**

	Total	Dedicated	Dabbler
Mean	-6.20	-5.49	-8.50
Median	0.22	0.24	0.13
STdev	1,109.7	899.7	1,612.8
Max	32,825	32,825	26,190
Min	-59,300	-59,300	-37,510
N trades	2,149,083	1,642,262	506,821

Note: This table presents summary statistics on the profitability of individual trades in our dataset. In columns two and three, we partition the data into trades made by those classified as Dedicated and as Dabblers.

Table 4: **Initial Account Balances (US$)**

	Total	Dedicated	Dabbler
Mean	2,773	8,512	1,101
Median	983	1,938	612
STdev	7,975	10,536	3,273
Max	185,650	185,650	95,458
Min	16	100	16
N users	5,361	1,885	3,476

Note: The number of users in this sample, 5,361, is less than the total number of traders we studied, 5,693, because the data was unavailable when coming from certain brokerages. In these instances we were unable to use the existing data to construct realistic estimates for their initial balance.

Table 5: **The Decision to Quit Trading**

	Total Deciles (deflated)			Within Group Deciles		
Decile	Baseline	Dabbler	Dedicated	Baseline	Dabbler	Dedicated
(Lowest) 1st	1.263***	1.545***	1.086	1.399***	1.481***	1.263***
	(0.069)	(0.114)	(0.088)	(0.070)	(0.091)	(0.113)
2nd	1.135**	1.213***	1.136	1.086	1.158**	0.963
	(0.058)	(0.081)	(0.090)	(0.056)	(0.072)	(0.087)
3rd	1.118**	1.202***	1.039	1.162***	1.225***	1.044
	(0.055)	(0.072)	(0.089)	(0.057)	(0.073)	(0.090)
4th	1.228***	1.196***	1.274***	1.265***	1.293***	1.216**
	(0.057	(0.067)	(0.106)	(0.059)	(0.074)	(0.097)
5th	1.505***	1.355***	1.734***	1.391***	1.286***	1.607***
	(0.065)	(0.070)	(0.134)	(0.062)	(0.073)	(0.116)
6th	1.257***	1.157***	1.316***	1.197***	1.134**	1.331***
	(0.057)	(0.062)	(0.115)	(0.056)	(0.067)	(0.103)
7th	0.816***	0.719***	1.003	0.805***	0.779***	0.871
	(0.044)	(0.047)	(0.095)	(0.044)	(0.054)	(0.078)
8th	0.581***	0.587***	0.569***	0.617***	0.647***	0.566***
	(0.037)	(0.045)	(0.065)	(0.038)	(0.049)	(0.062)
9th	0.529***	0.541***	0.549***	0.547***	0.542***	0.547***
	(0.036)	(0.048)	(0.060)	(0.037)	(0.045)	(0.063)
(Greatest) 10th	0.598***	0.598***	0.625***	0.591***	0.556***	0.630***
	(0.045)	(0.067)	(0.063)	(0.043)	(0.050)	(0.076)
Subjects	5,358	3,430	1,928			
Observations	77,307	39,354	37,953			
Quitters	4,135	2,693	1,442			

Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Description: This table displays hazard rates from estimating a Cox-proportional hazard model. The event in question is whether or not a trader quit in a given week. We generate independent variables by sorting the entire sample space of weekly returns and giving the observation a "1" if it is part of a given decile, "0" otherwise. In some specifications we deflate returns by the individual's median trade size in an attempt to capture individual wealth. We also try sorting the entire sample of weekly returns into deciles and in other specifications just the subset belonging to a trader's type. In all estimations we include controls for trader experience and age as well as the number of trades issued by the trader in each week. All three controls are associated with a decreased probability of quitting trading. Furthermore, we examine but do not report the prior week's and monthly performance and found similar results. Both cases yeild similar coeficients, but the results are of lower significance. We also computed, but do not report standard errors when clustering by trader and by week using the method outline in Froot (1989). This did not change the statistical significance of our results.

Table 6: **The Sending Function**

	I	II	III
	logit	OLS	Poisson
Returns	0.101***	0.064	0.100***
	(0.038)	(0.048)	(0.007)
Observations	48,102	3,350	3,350
(Pseudo) R-squared	0.003	0.010	0.059
Prob > F		0.000	0.000
LR Chi2(6)	59.74		
Prob > chi2	0.000		

Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Description: This table describes results from estimating the relationship between the returns of the sender and the number of trades issued by the receiver. The dependent variables in each specification and the estimation method are listed below. The independent variable in all regressions is log weekly returns. Regressions include controls for receiver age and experience. In other regressions we included time fixed effects as well as standard errors clustered by trader and by time, all of which had no effect on our results.

I: Logit, the dependent variable is an indicator for having sent a message.
II: OLS, the dependent variable is the number of trades issued, conditional on having sent at least one message.
III: Zero-Truncated Poisson Regression, the dependent variable is the number of trades issued, conditional on having sent at least one message.

Table 7: **The Sending Function by Trader Experience**

	Trading Experience (years)				
	none specified	0 - 1	1 - 3	4 - 5	5 - up
Returns	-1.062**	-0.035	0.113*	0.280***	0.187**
	(0.026)	(0.073)	(0.618)	(0.109)	(0.074)
Observations	481	14,642	22,719	4,578	5,682

Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Description: This table describes results from estimating a logit model, as in Table 1, column I, in which the dependent variable is an indicator variable for having sent a user message in a given week and the dependent variable is log weekly returns. We estimate each model separately for users binned into different experience levels. New myForexBook users are asked to specify their level of trading experience when registering and setting up their profile. They are allowed to choose one of the four options listed above, 0 - 1, 1 - 3, 4 - 5, or 5 - above years, or can bypass the question (none specified).

Table 8: **The Receiving Function**

log number of trades	I t+1	II t	III t+1	IV t+1
Returns (sum)	0.030***	0.022***		
	(0.008)	(0.007)		
Returns (max)			0.022***	
			(0.007)	
Returns (mean)				0.014*
				(0.008)
Observations	4,632	5,027	5,879	4,632
R-squared	0.014	0.011	0.012	0.012
Prob > F	0.000	0.000	0.000	0.000

Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Description: This table describes results from using OLS to estimate the relationship between the weekly returns of the sender, conditional on them being positive, and the number of trades issued by the receiver. Regressions include controls for receiver age and experience. In other regressions we included time fixed effects as well as standard errors clustered by trader and by time, all of which had no effect on our results.

I: The lagged one week forward number of receiver trades on the sum of sender returns.
II: The same week receiver trades on the sum of sender returns.
III: The lagged one week forward number of receiver trades on the max of sender returns.
IV: The lagged one week forward number of receiver trades on the mean of sender returns.

Table 9: **User Messages and Quitting Trading**

	Total	Dabbler	Dedicated
sent message	2.128***	1.765***	2.626***
	(0.103)	(0.109)	(0.207)
received message	0.675***	0.766***	0.569***
	(0.040)	(0.055)	(0.059)
Subjects	5,693	3,681	2,012
Observations	126,212	73,730	52,482
Quitters	4,587	3,051	1,536

Standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%.

Description: This table displays hazard rates from estimating a Cox-proportional hazard model. The event in question is whether or not a trader quit in a given week. The independent variables are indicators for whether or not a trader sent a message to another individual or received and read a message. We also computed, but do not report standard errors when clustering by trader using the method outline in Froot (1989). This did not change the statistical significance of our results.

Figure 1: **myForexBook "Dashboard"**

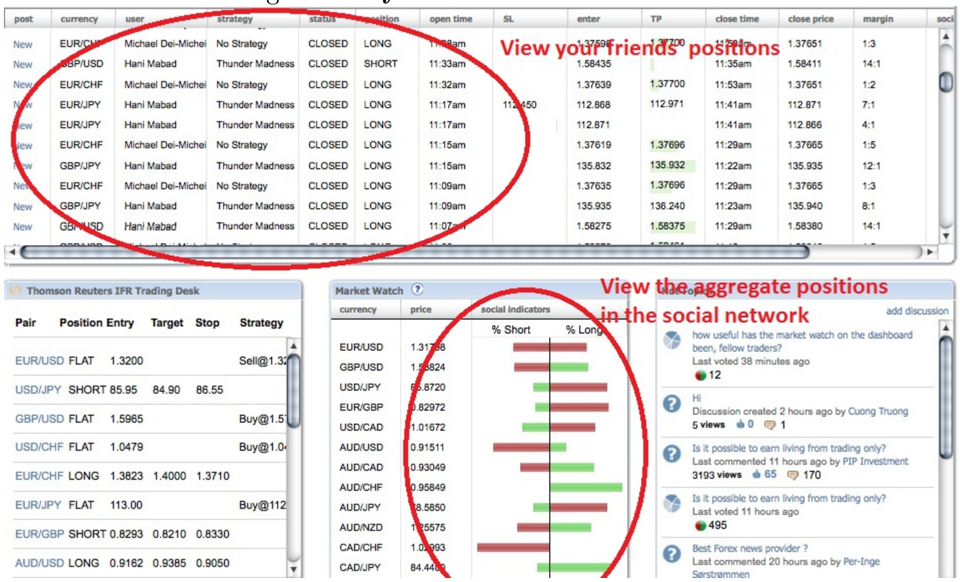

Note: This image displays the contents of a web browser that would be viewed by a myForexBook trader.

Figure 2: **Average Trades per Surviving User**

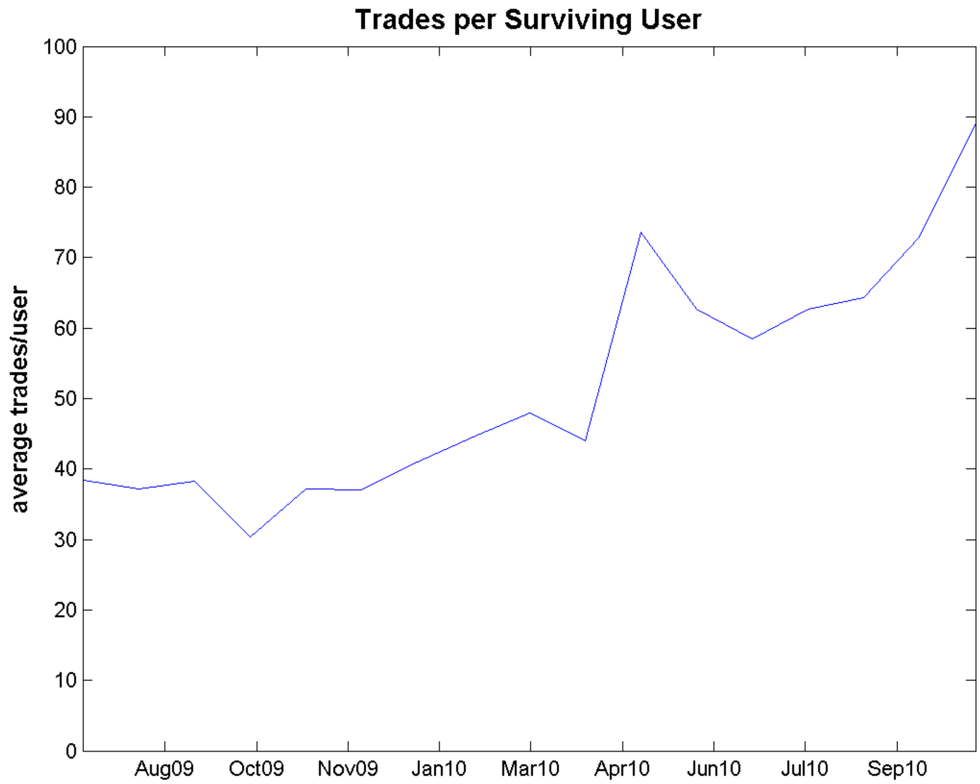

Figure 3: **Average Standard Deviation of Returns per User**

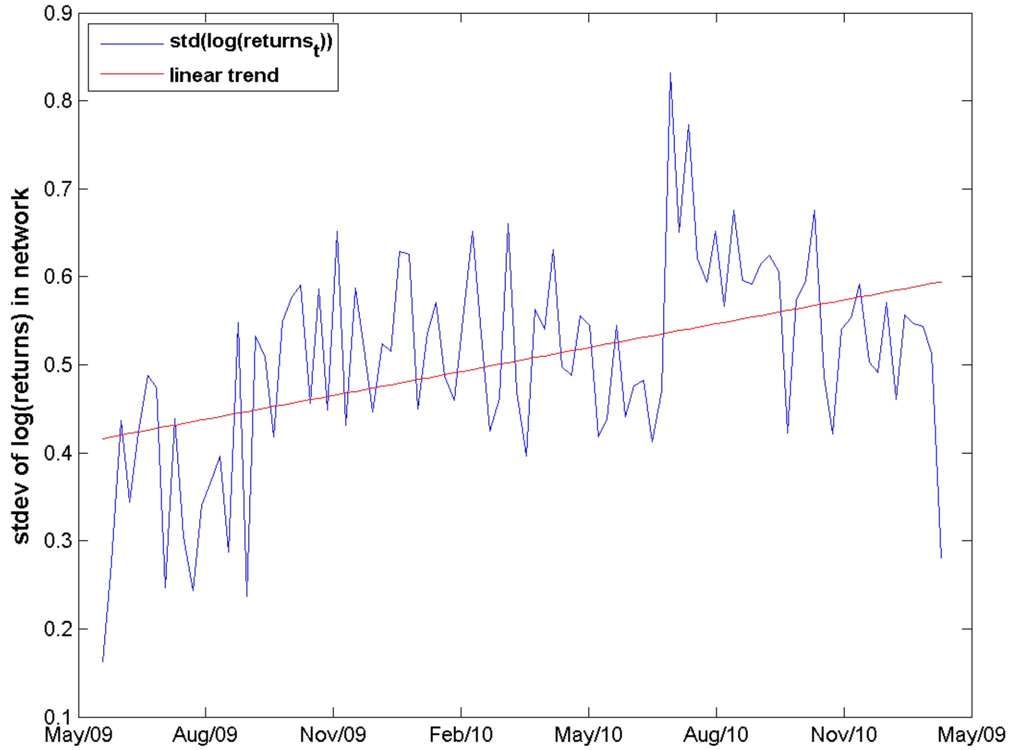

Figure 4: **Friendships Made in myForexBook**

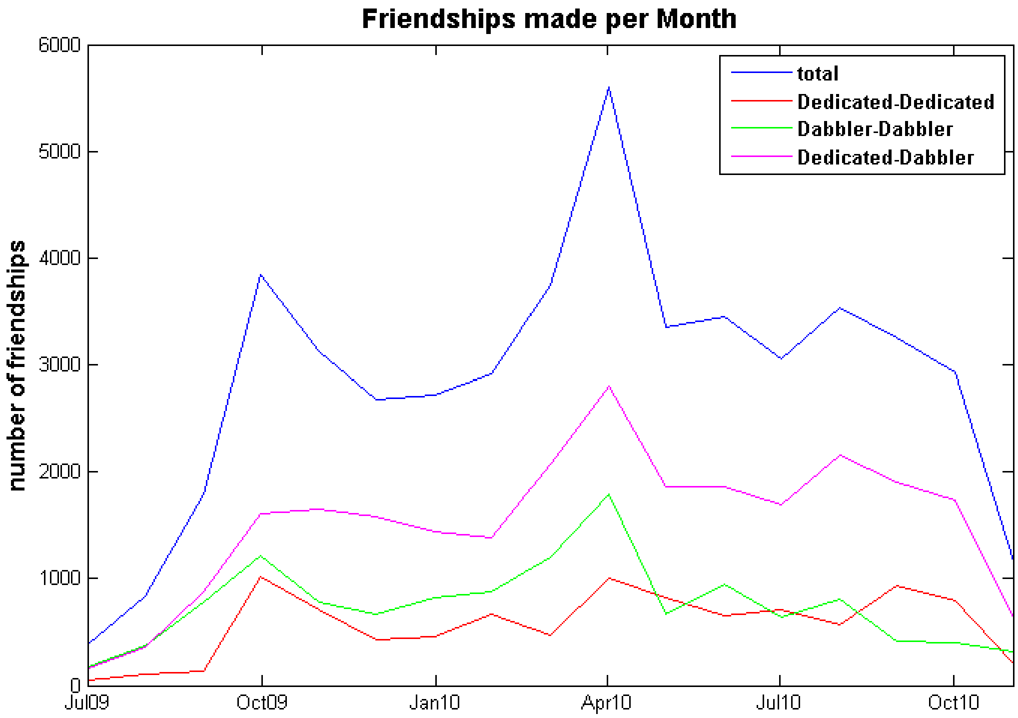

Figure 5: **Entries, Exits, and Survivors**

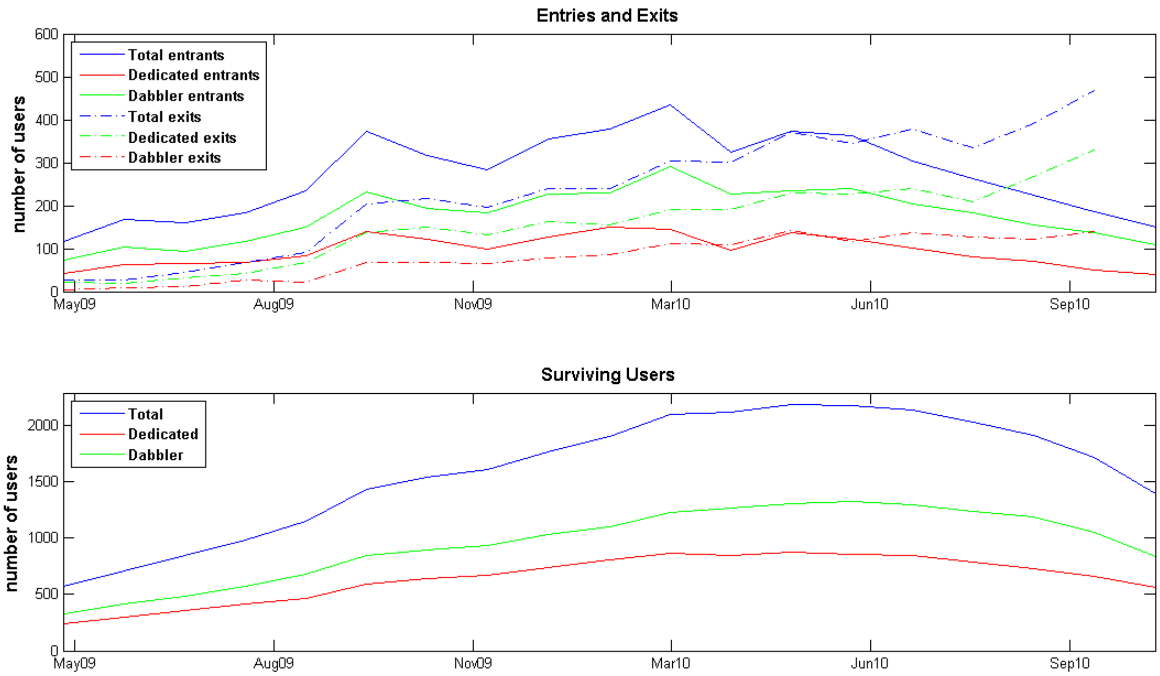

Figure 6: **Logins per Surviving User**

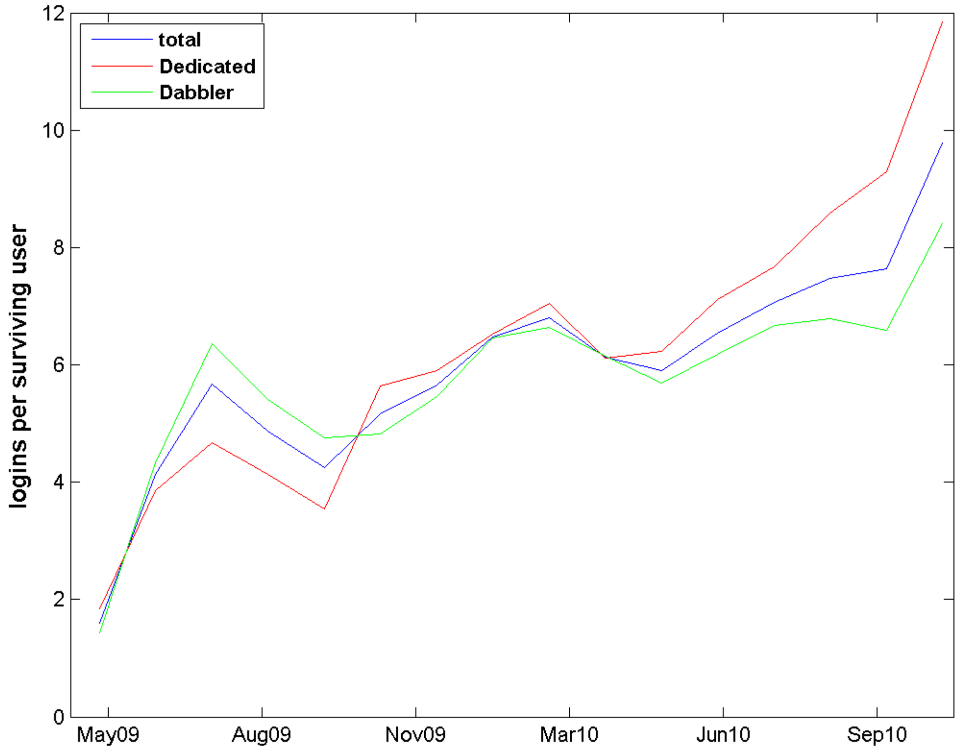

Note: This figure plots the number of logins made in a given month divided by the number of surviving users present in said month. A surviving user is defined as one who has had activity in the final month of the dataset. If the user did not survive, then they are said to have quit trading at the time of their last observable activity.