ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

Losing from Naive Reinforcement Learning: A Survival Analysis of Individual Repurchase Decisions

Peiran Jiao

Number 765
November 2015

Manor Road Building, Oxford OX1 3UQ

Losing from Naïve Reinforcement Learning: A Survival Analysis of Individual Repurchase Decisions

Peiran Jiao*

Department of Economics, University of Oxford

November 18, 2015

Abstract

This paper applies survival analysis to individual trading data from a discount brokerage firm, and documents significant individual-level repurchase bias, investors' tendency to disproportionately repurchase more previously sold winners than losers. Investor sophistication and experience mitigated the bias, but generated asymmetric effects: the most sophisticated/experienced investors' tendency to avoid prior losers were almost completely eliminated, but they were still over twice more likely to repurchase prior winners. Limited attention, chasing past performance and risk-adjusted returns could not justify the asymmetry. This suggests one reason for loss from frequent trading was persistent naïve reinforcement learning in repurchasing prior winners.

Keywords: Repurchase Bias, Reinforcement Learning, Sophistication, Experience.

JEL Classification Numbers: D10, D14, G10.

*Nuffield College and Department of Economics, University of Oxford. Address: New Road, Oxford, OX1 1NF, UK; Telephone: +44 01865 278993; E-mail: *peiran.jiao@economics.ox.ac.uk*. The author thanks Sir David Cox, Peyton Young, Paul Zak, Joshua Tasoff, Sean Flynn, and Heinrich Nax for helpful discussions, and participants at seminars in University of Oxford, Tsinghua University, Waseda University and Nankai University for comments and suggestions.

1 Introduction

Conventional portfolio choice theories hinge upon the assumption that individual investors form subjective beliefs about all future states, and maximize expected utility. Given the constraints they face and the complexity of the financial market, this task becomes so demanding that some form of learning is essential, regarding parameter values, private signal quality, personal skills, etc.[1] However, learning from experience can be suboptimal, because individual investors may overweight own experience relative to other information in a changing environment. This paper focuses on investors heterogeneous responses to own experience in common stock trading, and finds that their sophistication and experience help them but cannot stop them from losing by using naïve reinforcement strategies.

This work is partially motivated by an interest in understanding individual investor learning. The effect of experiential learning can manifest in the improvement of skills as experience accumulates. The finance literature documented that *more experience* can induce better trading performance (Nicolosi, Peng and Zhu, 2009) and reduce behavioral biases (Dhar and Zhu, 2006). To quantify experience, these studies use proxies such as the length of trading experience and the number of trades executed. What was missing in these measures is that good and bad experience may have differential impacts on behavior. For example, reinforcement learning predicts good experience with an action increases the probability of the action being retaken while bad experience decreases it, a robust finding in psychology and economics (see e.g. Thorndike, 1898; Erev and Roth, 1998; Charness and Levin, 2005). If skill, and/or private signal quality, is relatively constant, then investors with bad performance should exit (Seru, Shumway and Stoffman, 2010). Meanwhile, good experience can lead to overconfidence (Gervais and Odean, 2001), emergence and disappearance of investment styles (Barberis and Shleifer, 2003) and preferences towards certain assets, albeit unjustifiable given the nonstationarity of financial markets. Empirical evidence in finance suggests that *more*

[1] Learning can happen through either observation or personal experience, whereas this paper focuses on the latter. For a review of the literature on learning in financial markets, see e.g. Pastor and Veronesi (2009).

good experience with IPO auctions (Kaustia and Knüpfer, 2008), 401(k) portfolios (Choi et al., 2009) and common stocks (Strahilevitz, Odean and Barber, 2011, henceforth S2011) increases investors' subsequent demand for them.

Particularly in stock investments, S2011 document that investors used a reinforcement stragegy, repurchasing more stocks previously sold for gains than stocks sold for losses, which cannot be readily rationalized by investors' actual or perceived skill. This repurchase bias, on the one hand, is hard to justify on average given the the inherent randomness of the market. If skills improved with experience, investors should ultimately abandon the reinforcement strategy that is loss-inducing. In the meantime, it is also plausible that some investors carefully selected stocks and could sustain positive returns from repurchasing. S2011's analysis could not differentiate these cases. This paper addresses the issue by further analyzing the repurchase effect at the individual level, relating to investors' sophistication and experience, and comparing their repurchasing performances.

In the discount brokerage firm data in S2011, although significant repurchase bias was identified on average, only a small proportion of investors did repurchase: among the 19,059 households for whom the measure of repurchase bias can be calculated, 2,466 significantly repurchased more prior winners than losers and only 433 did the opposite at the 95% level. I use trading frequency, portfolio size, number of stocks in portfolio to proxy investor sophistication, and length of account tenure and number of trades executed to proxy experience. In the survival analysis with Weibull error distribution function, investors' sophistication and experience, especially trading frequency, significantly reduced the repurchase bias, but in an asymmetric way. Interestingly, for the most frequent investors, the tendency to avoid prior losers was almost completely eliminated, but they were still about twice more likely to repurchase prior winners. I further explore their reasons for doing so and find limited attention and risk adjusted returns could not explain the asymmetry. Thus, this suggests that naïve reinforcement learning could be a reason why investors lose from frequent trading.[2]

[2]There is a literature that documents individual investors' loss from trading too frequently. See e.g. Odean (1999).

Another motivation comes from the lack of empirical studies on individual buying decisions. Recently using individual-level trading data, some studies examine investors performance, trading intensity, portfolio composition, etc. (see e.g. Barber and Odean, 2000, 2001; Seru, Shumway and Stoffman, 2010), while others mostly investigate selling decisions. The focus on selling decisions is possibly because (1) the choice set of a selling decision can be somewhat unambiguously determined,[3] and (2) there is the interest in studying the disposition effect, which is fundamentally about selling (see e.g. Odean, 1998; Feng and Seasholes, 2005; Dhar and Zhu, 2006). Only a few individual-level empirical studies investigate buying decisions. Some compare the characteristics of stocks bought versus sold, evading the issue of an unclear choice set at the time of a purchase (see e.g. Grinblatt and Keloharju, 2001; Barber and Odean, 2008). S2011 and the present paper, study a special type of buying decision, i.e. repurchasing, which has a clear choice set that encompasses all stocks previously sold.

2 Data and Methodology

2.1 Data

The data used in this study contain trading records, position statements and demographic information of 78,000 households that had active accounts in a large US discount brokerage firm between January 1991 and December 1996.[4] The records include transactions of many different types of financial assets, whereas the focus here is only on common stocks. The data provide information of the stocks each account transacted each day: the commissions, transaction price and shares bought or sold. The position statements are the end-of-month portfolio compositions and their equity values. The price, return, volume, etc. of the corresponding stocks and of the market were retrieved from the Center for Research on Securities

[3]Not if short selling is permitted. However, this practice is considerably costly for individuals.
[4]A detailed description of this dataset can be found in Barber and Odean (2000).

Prices (CRSP) database.

This paper investigates the repurchase effect at the individual level, whereas the concept of an individual corresponds to a household in the dataset. Among the 78,000 households, 66,465 had common stock positions for at least one month, while on average 60 percent of their market value were invested in common stocks. For tax and other reasons, some households had more than one account, and each household had an average of roughly two accounts. The data also contain considerable amount of the demographic information. Various studies have demonstrated the effects of gender, income, occupation, etc. on trading behavior (see e.g. Dhar and Zhu, 2006).

2.2 Methodology

As the first step to individual analysis of the repurchase bias, I measure the tendency of each investor repurchasing stocks previously sold for gain (loss). To quantify this tendency, I follow S2011 in calculating the proportion of stocks repurchased out of the total opportunities to repurchase stocks previously sold for gain (loss), with the resulting measures called PWR (PLR). S2011 did this across all investors, while I generate these measures for each individual whenever the data permit. Two issues are noteworthy here. For one, when counting the number of opportunities to repurchase, the analysis restricts attention to sales on which return can be calculated, discarding instances where purchases were made before the start of the records, and thus the resulting number of opportunities to repurchase is potentially smaller than the actual number. But this should not systematically bias the measures of interest here. The other issue is that I focus on repurchases made within a year so that all stocks sold would have the same one-year window of repurchasing opportunities.

Specifically, to calculate PWR (PLR), on each day when an investor buys a stock, I count the number of sales at gain (loss) within the previous year whose return can be calculated, then aggregate across all trading days to get her total number of opportunities to repurchase stocks sold for gain (loss); PWR (PLR) is the total number of repurchases of stocks previously sold for gain (loss) divided by the total number of opportunities to do so. Equations (1) and (2) show how these measures are calculated for household i.

$$PWR_i = \frac{\text{Number of Prior Winners Repurchased}_i}{\text{Number of Opportunities to Repurchase Prior Winners}_i} \tag{1}$$

$$PLR_i = \frac{\text{Number of Prior Losers Repurchased}_i}{\text{Number of Opportunities to Repurchase Prior Losers}_i} \tag{2}$$

Secondly, I use survival analysis with a binary dependent variable accounting for whether a stock previously sold was repurchased (1=Repurchased, 0=Not Repurchased), and independent variables (or covariates) including gain/loss from previous sale, investors sophistication and experience, and other control variables that may also influence the propensity to buy. To model binary dependent variables, logit regressions are widely adopted: for example, Grinblatt and Keloharju (2001) study sell versus hold, and buy versus sell decisions using Finnish data, uncovering evidence of disposition effect and tax-loss selling; Ranguelova (2001) and O'Connell and Teo (2009) study the effect of firm size and institutional investors on disposition effect. Survival analysis recently became popular in finance, such as in studying the disposition effect (see e.g. Ivkovic, Poterba and Weisbenner, 2005; Seru, Shumway and Stoffman, 2010); Feng and Seasholes (2005) use this method and find that investors' sophistication and trading experience together eliminate their reluctance to realize losers, but not their propensity to sell winners. One distinct feature of survival analysis, compared with logit regression, is that it explicitly models the time to an event, such as death, failure, collapse of financial institutions, etc. To study the disposition effect, the event of interest is the stock's sale. S2011 and the present paper both study repurchasing decisions, with each stock entering the data when it was sold, the start date being the sales date, and the event of interest being the repurchase.

The survival analysis in S2011 is the Cox proportional hazard model (Cox, 1972), a non-parametric survival model, assuming when all the covariates are zero, the baseline hazard rate, i.e. the baseline repurchase rate, is the same on each day subsequent to the sale. This assumption seems too restrictive, because holding all covariates constant, the probability of repurchase potentially decreases over time, for reasons such as limited memory/attention. To account for this, I adopt the parametric survival analysis: the proportional hazard model with a Weibull error distribution function, for monotone hazard rates that increases or decreases exponentially with time.[5] In this model, the binary repurchase decision is regressed on the baseline hazard rate and other covariates, with the Weibull hazard and survival functions as the following:

$$h(t, p, X) = pt^{p-1}e^{X\beta}, \tag{3}$$

$$S(t, p, X) = e^{-e^{X\beta}t^p}, \tag{4}$$

where t is time, X are the covariates, and p the Weibull parameter for the shape of the baseline hazard function. The baseline hazard rate is $h_0(t) = pt^{p-1}$. An intercept β_0 is estimated, which serves as a scaler for the baseline hazard function $h_0(t)\exp(\beta_0)$. If the estimated p is 1, it implies constant baseline hazard rate over time, so the model reduces to the proportional hazard model with exponential error distribution. The model will be estimated using the maximum likelihood method, with standard errors clustered by households to account for correlation of observations within households. The marginal effect of each covariate will be a hazard ratio, i.e. the effect on repurchase rates per unit change of the covariate relative to the baseline *ceteris paribus*.

3 Results

3.1 Investor Sophistication and Experience

A closer examination of the data reveals that among the 66,465 households that held common stocks, 60,241 had at least one trading record; 5,355 had no sales, i.e. no repurchase opportunity, for whom we cannot calculate PWR or PLR. The rest of the households had at

[5] Feng and Seasholes (2005) use this method and find the baseline probability of selling a stock in ones portfolio decreases monotonically with length of holding period.

least one record of selling, but no sales return could be calculated for 29,760 households. This yields a restricted sample of 31,350 households. In the restricted sample, 17,084 households never repurchased even given the opportunity.

In order to measure investors sophistication and experience, I generate two sets of explanatory variables: the fixed and time-dependent covariates, according to whether the measure changes over time. The fixed covariates are the investors average monthly trading frequency ($AvgFreq$), average number of stocks in end-of-month portfolios ($AvgNum$), and average size of end-of-month portfolios in common stocks measured in thousands of dollars ($AvgSize$). Instead of using the initial portfolio values and number of stocks as in Feng and Seasholes (2005), I use monthly averages because (1) all accounts were opened on or after the start of their sample, which is not the case here, and (2) their had a two-year sample, which is relatively short compared to six years here, and within a longer time horizon, investors' skills may improve. Trading frequency was used as a measure of sophistication in Dhar and Zhu (2006). Another fixed covariate is the length of an investors trading experience, measured in years of account tenure ($Tenure$). Table 1 reports the summary statistics of these measures for the whole and restricted samples. Compared to the whole sample, the restricted sample households had higher trading frequency and were more diversified, but managing smaller portfolios.

(Insert Table 1 approximately here.)

I use two time-dependent covariates: on each day an investor trades, I measure the length of account tenure from account opening ($tdTenure$), and the number of trades executed since the beginning of the sample ($tdTrades$). Account tenure, number of positions and trades are popular measures of trading experience (see e.g. Feng and Seasholes, 2005; Dhar and Zhu, 2006).[6]

[6]Due to the small proportion of investors who repurchased given the opportunity, I used a logit regression to first test what kind of investors are more likely to repurchase. The results indicate that frequent traders repurchase more; tenure of account has no significant effect.

3.2 Repurchase of Prior Winners and Losers

Within the restricted sample, PWR can be calculated for 28,116 households, and PLR for 22,289 households. For 19,059 households, both measures and their difference can be calculated. Table 1 reports the summary statistics.[7] Figure 1 shows the histogram of $PWR-PLR$ across individuals, excluding those who never repurchased. The majority of these households had $PWR-PLR$ clustered around zero, but clearly the mass of the distribution lies above zero, suggesting a positive repurchase bias.[8]

(Insert Figure 1 approximately here.)

I calculate standard error of the measures as in equation (5), following Odean (1998). S2011 report the statistics of PWR and PLR for the aggregate data, and find their difference significantly positive using a standard error of $PWR - PLR$ calculated by a conservative measure, replacing the denominators in equation (5) with the number of repurchases of prior winners and losers. Although the conservative measure intends to alleviate the problem caused by violation of the independence assumption, using it requires nonzero repurchases, which is quite restrictive at the individual level. The test statistics indicate that for 2,466 households, the null hypothesis that $PWR-PLR <= 0$ can be rejected at the 5% significance level; and the opposite is true for only 433 households.

$$\text{S.E.} = \sqrt{\frac{PWR(1-PWR)}{\text{\# of opportunities to repurchase prior winners}} + \frac{PLR(1-PLR)}{\text{\# of opportunities to repurchase prior losers}}} \tag{5}$$

[7] Here I treat taxable and tax-exempt accounts together because S2011 showed that they suffered from similar repurchase bias.

[8] The spike immediately to the right of zero is not caused by households who had $PWR - PLR = 0$ sorted into that bin. There were only 72 such households and the histogram does not change much after removing them.

3.3 Survival analysis

3.3.1 The Baseline

This section uses survival analysis to first demonstrate the existence of a disproportionately larger propensity to repurchase stocks that brought previous gains than those that incurred losses. Then the measures of individual investor sophistication and experience will be added to the baseline model. The sample of this analysis encompasses all sales from January 1991 to December 1995 for which return can be calculated, whereas the event of interest occurs whenever a stock is repurchased within a year following its sale. Sales made after 1995 are dropped because they do not have at least one-year window for repurchases. There were a total of 62,821 repurchases (49,576 for prior winners and 13,245 for prior losers), out of 2,127,326 opportunities to repurchase (1,304,157 for prior winners and 823,169 for prior losers). As in S2011, the regressions below control for some usual factors that influence buying decisions, including the inverse holding period ($InvHoldP$) and log of relative trading volume ($LogRelVol$). The holding period is the number of days a stock was held before its sale; and relative trading volume is the volume of all trades of a stock measured in dollars in the previous year divided by the total volume of all trades of all stocks measured in dollars in the same year. These may trigger repurchases: a stock held longer before the previous sale could be a stock that the investor knew more about or felt more emotionally attached; a stock with higher relative trading volume could catch more attention.

The survival models also control for returns subsequent to sale: Grinblatt and Keloharju (2001) find that individual investors are net sellers of stocks with high past returns at least up to a week; Kaniel, Saar and Titman (2008) find individuals buy (sell) after price decreases (increases) in the previous month; S2011 demonstrate a disproportionate tendency to repurchase more stocks that declined in value after previous sale than those that increased in value, which they proposed to explain using counterfactual learning.[9] Though this is not

[9]They found the proportion of repurchases of stocks previously sold and declined in value roughly doubled the proportion of stocks previously sold and increased in value.

the focus of the present paper, I control for the performance of a stock subsequent to its sale using a set of performance measures. From the CRSP daily stock file, I calculate the 5, 10, 15, 21, 42-day and 63 trading day returns of each stock in excess of market return after the sales date, i.e. the stock return minus the CRSP value-weighted market portfolio return within each time range.

To justify the use of the Weibull error distribution instead of other parametric or non-parametric specifications in the survival model, Figure 2 plots the number of repurchases against the number of days between the sale and the repurchase. It clearly shows that as time elapsed, fewer stocks were repurchased, suggesting that the probability of repurchase declined monotonically over time. The majority of the repurchases were made within a short period after the original sale: 11.92% repurchases were made within five days after the sale, and only 15.73% were made more than a year later.

(Insert Figure 2 approximately here.)

Table 2a and 2b report results separately for the survival models that test effects of positive and negative previous sales returns. The first regressions in these tables report the effect of the main explanatory variable, the dummy variable *Pos* (*Neg*) which is equal to 1 when the previous sales return was positive (negative). This is similar to Feng and Seasholes (2005) in studying the disposition effect, but different from S2011, who generated a set of dummy variables for different brackets of sales returns with the purpose of capturing the complete locus of return effects. I use the positive and negative sales return dummies for the convenience of studying their interaction with investor characteristics. The results are consistent with S2011, showing that a positive sales return increased the likelihood of the stock's subsequent repurchase, and a negative return decreased it. The hazard ratio of approximately 1.16 on *Pos* indicates that, *ceteris paribus*, stocks previously sold at gain were 16% more likely to be repurchased later compared to stocks sold at negative or zero return. Similarly, stocks previously sold at loss were almost 40% less likely to be repurchased later. The set of control variables are included in all regressions, but their hazard ratios not reported. They consistently exert significant effect on repurchases. The return measures subsequent to sales reveal their inverse relationship with repurchases. The value of parameter p is significantly smaller than one, indicating that holding all covariates at zero, the baseline hazard rate, i.e. the chance of repurchase, decreased monotonically over time.

(Insert Table 2a and 2b approximately here.)

3.3.2 The Effects of Fixed Covariates

Now I add the fixed covariates ($AvgFreq$, $AvgSize$, $AvgNum$ and $Tenure$) to the survival analysis. But note that they are not added as continuous variables, which may complicate interpretation of the results. For each fixed covariate, I generate 3 dummy variables ($F2$, $F3$, $F4$) respectively for values falling into the 2nd, 3rd, and 4th quartile across all observations of the variable, and 0 otherwise. The 1st quartile is omitted to avoid multicollinearity.[10] The baseline hazard rate of each regression is obtained when all covariates are set to zero, and that specifically corresponds to each fixed covariate included in that regression falling in its first quartile. The product of the sales return dummy and the quartile dummies can reveal if investor sophistication/experience influenced the propensity to repurchase previous winners and losers.

The hazard ratio estimates for the quartile dummies alone show the tendency to repurchase stocks sold, regardless of it being a previous winner or loser. For example, the hazard ratios for the quartile dummies of $AvgFreq$ are all significant at 1% level and and they reveal that households in higher quartiles of trading frequency were at least three times more likely to repurchase stocks previously sold than households in the first quartile. The main interest of this analysis lies in the hazard ratios associated with the interaction variables. The results show that frequent traders repurchased more, and they were less likely to repurchase prior winners but more likely to repurchase prior losers, compared with less frequent traders. For

[10]Note that the quartile dummies for $AvgFreq$ are not exactly defined according to the quartiles, because as shown in Table 1, the 1st quartile for average trading frequency does not differ a lot from the 2nd quartile. Therefore, the dummies for $AvgFreq$ are for average trading frequencies between 1 and 2, between 2 and 3, and above 3 respectively, with the upper bound inclusive in each range.

example, in Table 2a regression (2), the hazard ratio for Pos is 6.6946, and that for $Pos \times F2$ is 0.5688, which means investors in the first quartile of trading frequency were 6.6946 times more likely to repurchase prior winners than other stocks sold; investors in the second quartile were only $6.6946 \times 0.5688 = 3.8079$ times more likely; investors in the fourth quartile were only $6.6946 \times 0.3801 = 2.5446$ times more likely. On the other hand, investors in the first quartile of trading frequency were about 53% less likely to repurchase a prior loser than if the stock yielded positive or zero return; the number is 65.89%, 77.82% and 89.63% for the second, third and fourth quartiles, respectively.

Investors managing larger portfolios repurchased more, and had smaller tendency to avoid prior losers; investors who had more stocks in portfolios had smaller tendency to repurchase prior winners; the length of trading experience has weakly signifiant effect on reducing the tendency to avoid prior losers. Together, the results point to the same direction for the influence of investor sophistication and experience on repurchase decisions, in that investors who were more sophisticated or had longer experience tended to repurchase more prior losers and fewer prior winners, leading to a smaller repurchase bias. Trading frequency exerted the largest effect among the variables tested. In a cross-sectional comparison, simply longer trading experience had little effect, which is consistent with Seru, Shumway and Stoffman (2010), who found that extra trades, but not extra length of experience, alleviated behavioral biases such as the disposition effect.

It is understandable that more frequent investors repurchased more, if the stocks they originally purchased were selected carefully, and their repurchases brought superior performances. But this does not justify any bias towards prior winners. The bias towards buying more prior winners is reduced at much slower rate than the bias towards avoiding prior losers. This asymmetry can possibly be rationalized if winners repurchased on average outperform losers repurchased, even for sophisticated traders. In what follows, I will investigate the performance of winners and losers repurchased for different groups of investors, and whether the bias can be eliminated by investor sophistication and experience.

3.3.3 The Effects of Time-Dependent Covariates

Now let's consider the time-dependent covariates ($tdTenure$ and $tdTrades$) to the survival analysis. Similar to the foregoing analysis, I create quartile dummy variables for each of them. The length of account tenure is represented by the number of years of account existence. The number of trades is a good indicator of the experience investors accumulate through trading. An investor who had an account for a long time was not necessarily an active trader. I first tried to test the effects of time-dependent covariates alone, but did not report the results. Basically, investors who executed larger number of trades exhibited similar behavioral patterns as those more sophisticated investors in the foregoing analysis, i.e. they had larger propensity to repurchase previously sold stocks in general, and smaller repurchase bias than their less experienced counterparts.

Up to this point, my analysis investigated the effect of each sophistication/experience measure separately, so it is still unclear whether the most sophisticated/experienced investors could completely get rid of the repurchase bias. I try to address this question by examining the combined effects of sophistication and experience on responses to previously sold winners and losers. The results for the best specification of such a model are reported in Table 3. This model includes the fixed measure of individual trading frequency, and the time-dependent measure of the length of trading experience. Number of trades executed is not included because it is correlated with trading frequency; portfolio size and number of stocks do not have additional explanatory power when trading frequency is controlled for.[11] First of all, the coefficients on the quartile dummies alone reaffirms that more frequent investors repurchased more. As to the interaction effects, apparently the most robust effect still comes from average trading frequency alleviating the bias in both domains of gains and losses, whereas the time-varying experience measure had some effect in reducing the tendency of avoiding prior losers. Controlling for trading frequency, the length of experience had no effect on repurchasing propensity on prior winners, which suggests that experience quantity is insufficient to help

[11]Other specifications are tried but fail to generate an improvement based on those shown in the table.

them completely avoid the naïve reinforcement strategy.

(Insert Table **??** approximately here.)

Individual investors' sophistication and experience have been shown to diminish behavioral biases (see e.g. List, 2003; Feng and Seasholes, 2005). Chiang et al. (2011) show that more sophisticated investors were less prone to the reinforcement learning bias in a cross-sectional comparison in the Taiwanese IPO market, but investors' auction selection ability and performance deteriorated with experience due to naïve reinforcement learning. To demonstrate the effects here visually, Figure 3 depicts the relationship between the average trading frequency quartiles and the hazard ratios associated with each quartile, showing different sensitivities to prior winners (Panel A) and to prior losers (Panel B). Both curves are converging to 1, which represents no bias in repurchasing prior winner or loser. Generally, there is a tendency for higher trading frequency alone to eliminate a large proportion of the repurchase bias: Investors in the fourth quartile of average trading frequency were 61.92% $(1-0.3808)$ less likely to repurchase previously sold winners, and 89.93% $(1.8993-1)$ more likely to repurchase previously sold losers. However, there was an asymmetry: the bias of avoiding prior losers was almost entirely eliminated, but the bias of repurchasing prior winners still existed. And this is the case even for the investors in the most sophisticated/experienced quartile were still more than twice ($6.8408 \times 0.3808 = 2.6050$, significantly different from 1 with $p < 0.01$) more likely to repurchase a prior winner, but had a sensitivity to prior losers of $0.4190 \times 1.8993 \times 1.1826 = 0.9411$ (not significantly different from 1).

(Insert Figure 3 approximately here.)

The drop in naïve reinforcement learning can be caused by the larger sample of personal direct experience which helped them debias. In addition to direct experience, even inactive investors could learn from the indirect experience of others. The additional effect from the length of experience in reducing the tendency to avoid prior losers could be due to this observational learning. It is worth noting that debiasing might be more difficult in common stock trading, because in the IPO market, each new IPO auction belongs to a different company, which makes it easier to realize that past and future performances are not necessarily correlated, but this is not the case here, in that stocks previously sold are always in the choice set when investors make subsequent buying decisions. Thus, this might be a reason why the sophistication and experience variables combined fail to completely eliminate the bias given the abundant opportunities to learn in the stock market.

Although sophistication and experience together failed to completely eliminate the propensity to repurchase previously sold winners, they almost completely eliminate the propensity to avoid previously sold losers. That is, one component of naïve reinforcement learning, which is to repeat actions that brought pleasant experience, was more stubborn than the component of avoiding those that brought unpleasant experience. [12] This asymmetry in the gain versus loss domain is comparable to Feng and Seasholes (2005) who find an asymmetric effect where sophistication and experience together eliminated the tendency to hold losers but only dampened the tendency to sell winners, reducing the disposition effect. They try to use mental accounting to explain this but do not find supporting evidence.[13] Mental accounting suggests that losses should be combined and gains should be treated separately, but this is with regard to how investors treat stocks in their portfolio when they try to close a mental account, whilst repurchasing opens new mental accounts. So the same logic does not apply well to repurchasing decisions. This asymmetry can be rationalized if winners repurchased on average performed better than losers repurchased, especially for high frequency traders. To test this hypothesis, the performances of stocks repurchased by traders of different trading frequency are presented in the following section.

3.4 Why and What did Investors Repurchase?

The foregoing analyses have demonstrated that more sophisticated investors, especially those with higher trading frequency, compared with their less sophisticated counterparts,

[12] Other sophistication and experience measures were tried, but the sensitivity to previous gains cannot be reduced to below 2.

[13] I propose that simple reinforcement learning might explain their results but not mine, although careful evaluations of this hypothesis is needed.

repurchased fewer prior winners but more prior losers, although this does not completely eliminate their tendency of naïve reinforcement learning. Up to this point, it is still unclear why they repurchased, and whether there is any performance-related motivation that triggered their repurchase pattern. This section will address these questions.

The first possibility is that individual investors, including sophisticated ones, were constrained by their limited attentional budget, so that they monitored stocks previously sold because it is harder to search for new stocks than to focus attention on a particular subset of stocks, especially those prior winners.[14] The trading accounts in my sample belonged to individual investors, who were supposed to have limited attention issue and more prone to behavioral biases than institutional investors (see e.g. Barber and Odean, 2008) . Suppose repurchases were caused by attention constraints, then in order to purchase a stock never owned before investors should require the stock to be more attention-catching than a repurchased stock. The above-documented asymmetry could possibly emerge if attention is an issue also for frequent traders.

Table 4 presents a comparison of the attention-catching characteristics of previously unowned stocks purchased, and stocks repurchased that brought positive, negative or zero return on previous sale, summarized separately for investors in the top and bottom half of trading frequency distribution. Three categories of stock characteristics are compared. Firstly, market variables that may attract attention are compared in the upper panel of Table 4: $AbVol$, abnormal volume, is the number of standard deviations from the current trading volume to the average trading volume of the stock within the previous year; $HistH$ and $HistL$ are indicators of the prior day price being a new historical high or new historical low in the past month, so averages of these variables represent the proportion of cases where the stock bought had a new historical high or low price; $Volatility$, measured by average squared returns during the previous week, is known to drive disagreements among investors and to be positively correlated with trading volume (Karpoff, 1987).

[14] Feng and Seasholes (2005) also find some evidence that investors continued monitoring the stocks they sold.

If investors faced attentional constraints, they should be more willing to purchase a stock previously sold, even with zero return, than a stock never owned before; and due to naïve reinforcement learning, they should have a bias towards previously sold winners. Therefore, within each trading frequency category, the attention-grabbing factors should be higher for stocks never owned than previously sold stocks, and the stocks previously sold with positive returns should be the least attention-catching; on the other hand, the differences should be less salient for more frequent traders. The predicted patterns manifested clearly in abnormal trading volumes. Moreover, investors chose to purchase never-owned stocks that were more likely to have historical high prices but less likely to have historical low prices than repurchased stocks.[15] The volatility measures did not generate much difference. Comparing between the two categories of investors, many attention-grabbing factors were lower for infrequent traders than for frequent traders when they repurchased.

The second comparison, reported in the middle panel of Table 4, is about relative performance before the (re)purchase, measured by the return to the stock in excess of the CRSP value-weighted market portfolio return within the 5, 15, 21, 63 or 126 trading days. Recent good performance should also be attention-catching and should attract momentum traders to buy the stock. In general the frequent traders consistently chose better past performers than the infrequent traders in all horizons tested in both repurchases and new purchases. Both categories of investors apparently required significantly better prior performance before they bought new stocks or repurchased stocks previously sold with zero return than when they made repurchases of prior winners or losers, which was especially prominent within short horizons up to 21 trading days before the purchase. Additionally, the less frequent traders were contrarian buyers of previously sold stocks. The more frequent traders seemed to be short run (less than a month) contrarian traders but momentum in the medium run when repurchasing prior losers, which should be a profitable strategy suggested by the literature on momentum and contrarian strategies (Jegadeesh, 1990; Jegadeesh and Titman,

[15]This is also consistent with counterfactual learning in Strahilevitz, Odean and Barber (2011).

1993). However, they were doing the opposite when repurchasing prior winners. Together with the results from the upper panel, these suggest that limited attention seemed to be an issue for these individual investors, and it was more severe for infrequent traders; frequent traders were more careful in selecting previous losers to repurchase, but not in repurchasing prior winners.

And lastly, the lower panel of Table 4 compares relative performances in the 5, 15, 21, 63 or 126 trading days after the (re)purchase. If investors (re)purchased due to reasons such as private information or superior stock-selection skills, then the stocks they (re)purchased should have good performance; and because the frequent traders still had a large propensity to repurchase prior winners, we would expect they made positive return from doing so. However, the results suggest that on average both categories of investors were losing from their (re)purchases, indicating poor stock-picking abilities. Overall, these investors could have been better off purchasing stocks never owned before than repurchasing. The previously sold winners underperformed previously sold losers in almost (but not) all horizons. The previously losers repurchased by frequent traders outperformed those repurchased by infrequent traders. These can potentially justify why frequent traders repurchased more losers, but not why they could not resist repurchasing prior winners, leaving naïve reinforcement learning a plausible explanation.

(Insert Table 4 approximately here.)

To further explore the performance of stocks repurchased and whether frequent traders were better at selecting stocks to repurchase, next I test the risk-adjusted returns of the repurchased stocks. I construct portfolios that use 1 dollar to purchase the previous winners/losers investors repurchased and hold each position for one calendar year, and calculate the portfolio risk-adjusted returns using the Capital Asset Pricing Model (CAPM). The four-factor model (Fama and French, 1993; Carhart, 1997) is also estimated to take into account the effects of investment styles on returns, such as small or large market capitalization, value or growth stock and price momentum. These portfolios are constructed for all investors, those in the top quartile and those in the bottom quartile of trading frequencies respectively. Equation (6) is the four-factor model to be estimated,

$$(R_t^p - R_t^f) = \alpha + \beta(R_t^m - R_t^f) + sSMB_t + hHML_t + uUMD_t + \epsilon_t, \tag{6}$$

where R_t^p is the return to the constructed portfolio; R_t^f is the one-month treasury bill rate; R_t^m is the return to the value-weighted market portfolio; SMB_t is the average return on the three small-cap portfolios minus the average return on the three big-cap portfolios; HML_t is the average return on the two value stock portfolios minus the average return on the two growth stock portfolios; and UMD_t is the average return on the recent winner portfolio minus the return on the recent loser portfolio. The estimation results are reported in Table 5.

(Insert Table 5 approximately here.)

Across all investors in general, the winners repurchased underperformed the market by 0.58% per day while the losers repurchased outperformed by 0.27% and these cannot be explained by investment styles, so the investors would be better off repurchasing more previously sold losers. For the infrequent traders, neither the CAPM nor the four-factor model could explain the performance of the stocks they repurchased: the losers they repurchased did not generate significant abnormal return, but the winners they repurchased were much worse, underperforming by 0.74% per day. For the frequent traders, it did seem that they were more careful in picking which stocks to repurchase, generating higher return in both portfolios of repurchased winners and losers than the whole sample, but their repurchased winners still considerably underperformed, which cannot rationalize their propensity to repurchase prior losers. In choosing the losers to repurchase, they were picking some recently underperforming low-beta small value stocks. But the four factors cannot explain the abnormal return to their prior losers portfolio: the intercept terms were the same as in the CAPM, significantly positive, suggesting that they probably exhibited some skills when choosing which losers to repurchase.

4 Conclusion

This paper contributes to our understanding of individual investors' repurchasing decisions, a special form of buying decisions for which the choice set is determined at the time the decisions are made. Consistent with previous findings at the aggregate level, at the individual level, there was a significant bias towards repurchasing more stocks previously sold at gain than those sold at loss, as if they were following the reinforcement learning strategy, repeating the action that brought pleasant experience while avoiding those that brought unpleasant experience.

Additionally, this paper addresses the issue of investor learning, by investigating the relationship between investors sophistication/experience and the repurchase bias. Trading frequency or the number of trades executed, but not the length of trading experience, helped investors reduce the repurchase bias. However, sophistication and experience had an asymmetric effect: even for the most frequent and experienced traders, although the tendency to avoid prior losers was almost entirely eliminated, they were still about twice more likely to repurchase prior winners.

By further investigating the characteristics of stocks (re)purchased by investors with different trading frequencies, I find that limited attention was an issue for them, especially the infrequent traders, that they were not simply chasing past returns but using counterfactuals when repurchasing, and that although the most frequent traders had some skills in choose prior losers to repurchase, the risk-adjusted returns could not justify their repurchase of prior winners. None of these could explain the asymmetry mentioned above, leaving reinforcement learning motives still a plausible explanation. So by trading frequently, individual investors lose from naïve reinforcement learning by repurchasing previously sold winners, which can be a potential reason for the bad performance documented for frequent traders.

The asymmetry in the effect of sophistication/experience calls for further research efforts. In the meantime, a test of reinforcement learning is a test of the joint hypothesis regarding the choice of the source of reinforcement and the learning pattern. Here we assume investors received reinforcements from absolute gains and losses, while other potential sources of reinforcements can be validated.

References

Barber, Brad M, and Terrance Odean. 2000. "Trading is hazardous to your wealth: The common stock investment performance of individual investors." *The Journal of Finance*, 55(2): 773–806.

Barber, Brad M, and Terrance Odean. 2001. "The internet and the investor." *Journal of Economic Perspectives*, 41–54.

Barber, Brad M, and Terrance Odean. 2008. "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors." *Review of Financial Studies*, 21(2): 785–818.

Barberis, Nicholas, and Andrei Shleifer. 2003. "Style investing." *Journal of Financial Economics*, 68(2): 161–199.

Carhart, Mark M. 1997. "On persistence in mutual fund performance." *The Journal of Finance*, 52(1): 57–82.

Charness, Gary, and Dan Levin. 2005. "When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect." *American Economic Review*, 1300–1309.

Chiang, Yao-Min, David Hirshleifer, Yiming Qian, and Ann E Sherman. 2011. "Do investors learn from experience? evidence from frequent ipo investors." *Review of Financial Studies*, hhq151.

Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick. 2009. "Reinforcement learning and savings behavior." *The Journal of Finance*, 64(6): 2515–2534.

Cox, David R. 1972. "Regression models and life-tables." *Journal of the Royal Statistical Society*, 34: 187–220.

Dhar, Ravi, and Ning Zhu. 2006. "Up close and personal: Investor sophistication and the disposition effect." *Management Science*, 52(5): 726–740.

Erev, Ido, and Alvin E Roth. 1998. "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria." *American Economic Review*, 848–881.

Fama, Eugene F, and Kenneth R French. 1993. "Common risk factors in the returns on stocks and bonds." *Journal of Financial Economics*, 33(1): 3–56.

Feng, Lei, and Mark S Seasholes. 2005. "Do investor sophistication and trading experience eliminate behavioral biases in financial markets?" *Review of Finance*, 9(3): 305–351.

Gervais, Simon, and Terrance Odean. 2001. "Learning to be overconfident." *Review of Financial Studies*, 14(1): 1–27.

Grinblatt, Mark, and Matti Keloharju. 2001. "What makes investors trade?" *The Journal of Finance*, 56(2): 589–616.

Ivkovic, Zoran, James Poterba, and Scott Weisbenner. 2005. "Tax-Motivated Trading by Individual Investors." *American Economic Review*.

Jegadeesh, Narasimhan. 1990. "Evidence of predictable behavior of security returns." *The Journal of Finance*, 45(3): 881–898.

Jegadeesh, Narasimhan, and Sheridan Titman. 1993. "Returns to buying winners and selling losers: Implications for stock market efficiency." *The Journal of Finance*, 48(1): 65–91.

Kaniel, Ron, Gideon Saar, and Sheridan Titman. 2008. "Individual investor trading and stock returns." *The Journal of Finance*, 63(1): 273–310.

Karpoff, Jonathan M. 1987. "The relation between price changes and trading volume: A survey." *Journal of Financial and Quantitative Analysis*, 22(01): 109–126.

Kaustia, Markku, and Samuli Knüpfer. 2008. "Do investors overweight personal experience? Evidence from IPO subscriptions." *The Journal of Finance*, 63(6): 2679–2702.

List, John A. 2003. "Does market experience eliminate market anomalies?" *Quarterly Journal of Economics*, 118(1): 41–72.

Nicolosi, Gina, Liang Peng, and Ning Zhu. 2009. "Do individual investors learn from their trading experience?" *Journal of Financial Markets*, 12(2): 317–336.

O'Connell, Paul GJ, and Melvyn Teo. 2009. "Institutional investors, past performance, and dynamic loss aversion." *Journal of Financial and Quantitative Analysis*, 44(01): 155–188.

Odean, Terrance. 1998. "Are investors reluctant to realize their losses?" *The Journal of Finance*, 53(5): 1775–1798.

Odean, Terrance. 1999. "Do Investors Trade Too Much?" *American Economic Review*, 89(5): 1279–1298.

Pastor, Lubos, and Pietro Veronesi. 2009. "Learning in Financial Markets." *Annual Review of Financial Economics*, 1(1): 361–381.

Ranguelova, Elena. 2001. "Disposition effect and firm size: New evidence on individual investor trading activity." *Available at SSRN 293618.*

Seru, Amit, Tyler Shumway, and Noah Stoffman. 2010. "Learning by trading." *Review of Financial Studies*, 23(2): 705–739.

Strahilevitz, Michal Ann, Terrance Odean, and Brad M Barber. 2011. "Once burned, twice shy: How naïve learning, counterfactuals, and regret affect the repurchase of stocks previously sold." *Journal of Marketing Research*, 48(SPL): S102–S120.

Thorndike, Edward L. 1898. "Animal intelligence: An experimental study of the associative processes in animals." *Psychological Monographs: General and Applied*, 2(4): i–109.

Table 1: Summary Statistics

This table reports the summary statistics of investors sophistication and experience measures for the whole sample of 60,241 households and the restricted sample of 31,350 households, respectively. $AvgFreq$ is the average trading frequency, $AvgSize$ is the average size of end-of-month portfolios in thousands of dollars; $AvgNum$ is the average number of stocks in end-of-month portfolios; $Tenure$ is the length of household trading experience in years from first account opening to December 1996. The table also reports summary statistics of PWR, PLR and $PWR - PLR$.

		N	Mean	Std Dev	1st Q	2nd Q	3rd Q
AvgFreq	Whole	60,241	1.74	1.42	1.00	1.40	1.91
	Restricted	31,350	2.18	1.80	1.40	1.75	2.37
AvgSize	Whole	60,241	57.51	187.50	9.72	20.57	47.06
	Restricted	31,350	43.49	189.43	7.10	15.30	34.87
AvgNum	Whole	60,241	4.48	5.69	1.55	2.83	5.27
	Restricted	31,350	5.66	6.88	2.10	3.72	6.77
Tenure	Whole	60,241	10.09	3.28	7.10	9.90	12.61
	Restricted	31,350	10.00	3.37	6.90	9.78	12.54
PWR		28,116	0.0678	0.1699	0	0	0.0512
PLR		22,289	0.0389	0.1392	0	0	0.0094
PWR-PLR		19,059	0.0161	0.1452	0	0	0.0263

Table 2a: The Effects of the Fixed Covariates on Repurchasing Prior Winners

This table reports the survival analyses results regarding the effects of fixed covariates on repurchasing prior winners. Regressions (1) to (4) respectively use $AvgFreq$, $AvgSize$, $AvgNum$ and $Tenure$ as the fixed covariate of interest, introducing the second, third and fourth quartile dummies ($F2$, $F3$ and $F4$), as well as their interaction with dummy variable Pos. The control variables are included but not reported. Hazard ratios of the covariates are reported. Clustered robust standard errors are in parentheses. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$)

	(1)	(2) *F=AvgFreq*	(3) *F=AvgSize*	(4) *F=AvgNum*	(5) *F=Tenure*
Pos	1.1632***	6.6946***	2.7987***	3.0260***	2.7340***
	(0.0155)	(1.1564)	(0.1696)	(0.1713)	(0.1108)
F2		3.3033***	0.9932	1.1706*	1.1390*
		(0.5301)	(0.0940)	(0.1018)	(0.0793)
F3		5.6030***	1.2827***	1.0840	1.0504
		(0.9004)	(0.1183)	(0.0883)	(0.0680)
F4		10.7093***	1.9220***	1.0727	1.1110
		(1.7314)	(0.1765)	(0.0936)	(0.0961)
Pos×F2		0.5688***	1.0764	0.8720**	0.9543
		(0.0993)	(0.0766)	(0.0585)	(0.0508)
Pos×F3		0.4355***	1.0401	0.8658**	0.9895
		(0.0761)	(0.0696)	(0.0550)	(0.0497)
Pos×F4		0.3801***	0.9347	0.8757**	0.9548
		(0.0663)	(0.0611)	(0.0563)	(0.0572)
Control Variables	Yes	Yes	Yes	Yes	Yes
p	0.4332	0.4376	0.4357	0.4345	0.4345
	(0.0042)	(0.0047)	(0.0047)	(0.0047)	(0.0047)
N	599,735	596,854	596,854	596,854	596,854

Table 2b: The Effects of the Fixed Covariates on Repurchasing Prior Losers

This table reports the survival analyses results regarding the effects of fixed covariates on repurchasing prior losers. Regressions (1) to (4) respectively use $AvgFreq$, $AvgSize$, $AvgNum$ and $Tenure$ as the fixed covariate of interest, introducing the second, third and fourth quartile dummies ($F2$, $F3$ and $F4$), as well as their interactions with dummy variable Neg. The control variables are included but not reported. Hazard ratios of the covariates are reported. Clustered robust standard errors are in parentheses. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$)

	(1)	(2) *F=AvgFreq*	(3) *F=AvgSize*	(4) *F=AvgNum*	(5) *F=Tenure*
Neg	0.5972***	0.4678***	0.7038***	0.8560**	0.8272***
	(0.0106)	(0.0819)	(0.0429)	(0.0577)	(0.0324)
F2		2.0406***	1.0332	1.0760	1.0619
		(0.1507)	(0.0564)	(0.0484)	(0.0450)
F3		2.7066***	1.2650***	1.0035	1.0140
		(0.2018)	(0.0668)	(0.0432)	(0.0399)
F4		4.3661***	1.6090***	0.9817	1.0694
		(0.3307)	(0.0854)	(0.0458)	(0.0504)
Neg×F2		1.4086*	1.0040	1.0276	1.0994*
		(0.2492)	(0.0719)	(0.0782)	(0.0554)
Neg×F3		1.6635***	1.0767	1.0000	1.0848*
		(0.2943)	(0.0722)	(0.0734)	(0.0535)
Neg×F4		1.9159***	1.3389***	1.0350	1.0617
		(0.3382)	(0.0870)	(0.0757)	(0.0613)
Control Variables	Yes	Yes	Yes	Yes	Yes
p	0.4346	0.4319	0.4303	0.4294	0.4294
	(0.0042)	(0.0046)	(0.0046)	(0.0047)	(0.0047)
N	599,735	596,854	596,854	596,854	596,854

Table 3: COMBINED EFFECTS OF THE FIXED AND TIME-DEPENDENT COVARIATES

This table reports the survival analyses results regarding the combined effects of the fixed and time-dependent covariates on repurchasing. Regressions (1) and (2) test their effects on repurchasing prior winners and losers respectively. D represents the dummy variable for the sign of previous sales return. $D = Pos$ in regression (1) and $D = Neg$ in regression (2). The explanatory variables included here are $AvgFreq$ and $tdTenure$. The second, third and fourth quartile dummies, as well as their interactions with Pos or Neg, are added to the survival analyses. The control variables are included but not reported. Hazard ratios of the covariates are reported. Clustered robust standard errors are in parentheses. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$)

	(1) *D=Pos*		(2) *D=Neg*	
	Haz. Ratio	Std. Err.	Haz. Ratio	Std. Err.
D	6.8408***	(1.1934)	0.4190***	(0.0740)
AvgFreq2	3.3150***	(0.5321)	2.0495***	(0.1514)
AvgFreq3	5.6276***	(0.9047)	2.7236***	(0.2033)
AvgFreq4	10.7886***	(1.7411)	4.4177***	(0.3343)
D×AvgFreq2	0.5684***	(0.0993)	1.4046*	(0.2486)
D×AvgFreq3	0.4357***	(0.0762)	1.6569***	(0.2931)
D×AvgFreq4	0.3808***	(0.0664)	1.8993***	(0.3352)
tdTenure2	0.9558	(0.0459)	0.8739***	(0.0226)
tdTenure3	0.9635	(0.0449)	0.9068***	(0.0277)
tdTenure4	0.8916	(0.0679)	0.8124***	(0.0347)
D×tdTenure2	0.9690	(0.0429)	1.1697***	(0.0473)
D×tdTenure3	0.9935	(0.0377)	1.1458***	(0.0437)
D×tdTenure4	0.9480	(0.0483)	1.1826***	(0.0580)
Control Variables	Yes	Yes	Yes	Yes
p	0.4376	(0.0047)	0.4319	(0.0047)
N	596,854		596,854	

Table 4: Characteristics of Stocks Purchased and Repurchased

This table reports the means and standard deviations (in parentheses) of the characteristics of stocks purchased with positive, negative or zero return from prior sales and those that were never owned before, for investors in the least ($Freq = 1, 2$) and most ($Freq = 3, 4$) average trading frequencies. The characteristics include the following: $Volatility$ is the squared return of the previous trading day; $HistH$ and $HistL$ are indicators of whether the price was a new historical high or historical low; $AbVol$ is the difference between the volume of a stock and the previous-year average measured in standard deviations. The $ExRet$ variables are the excess returns of the stock compared to the CRSP value-weighted market portfolio in the 5, 15, 21, 63 and 126 days before and after the purchase.

	Positive		Negative		Zero		Never Owned	
	Freq=1,2	*Freq*=3,4	*Freq*=1,2	*Freq*=3,4	*Freq*=1,2	*Freq*=3,4	*Freq*=1,2	*Freq*=3,4
N	11,209	56,207	2,780	24,284	4,493	18,185	251,189	467,612
AbVol	1.7471	1.9307	1.8987	2.0954	1.8446	2.7257	2.6644	2.7300
	(4.7936)	(3.3609)	(2.9076)	(3.6093)	(2.6989)	(6.4530)	(51.158)	(22.715)
HistH	0.0960	0.1414	0.0641	0.1067	0.0904	0.1481	0.1130	0.1549
	(0.2947)	(0.3484)	(0.2450)	(0.3088)	(0.2867)	(0.3553)	(0.3166)	(0.3618)
HistL	0.2050	0.1698	0.2844	0.2151	0.2221	0.1779	0.1723	0.1478
	(0.4038)	(0.3755)	(0.4512)	(0.4109)	(0.4157)	(0.3825)	(0.3776)	(0.3549)
Volatility	0.0021	0.0021	0.0016	0.0017	0.0015	0.0018	0.0019	0.0019
	(0.0041)	(0.0076)	(0.0037)	(0.0062)	(0.0073)	(0.0099)	(0.0111)	(0.0112)
$ExRet_{-5}$	-0.0089	0.0054	-0.0335	-0.0141	-0.0163	0.0171	0.0088	0.0204
	(0.1626)	(0.1492)	(0.1168)	(0.1976)	(0.2658)	(0.4248)	(0.4132)	(0.3795)
$ExRet_{-15}$	-0.0215	0.0007	-0.0451	-0.0060	-0.0358	0.0135	0.0074	0.0301
	(0.2350)	(0.2165)	(0.1825)	(0.2600)	(0.3595)	(0.4038)	(0.4082)	(0.3911)
$ExRet_{-21}$	-0.0261	-0.0015	-0.0449	0.0025	-0.0447	0.0126	0.0074	0.0352
	(0.2666)	(0.2439)	(0.2275)	(0.2938)	(0.3680)	(0.4153)	(0.4269)	(0.4136)
$ExRet_{-63}$	-0.0534	-0.0013	0.0021	0.0829	-0.0745	0.0231	0.0124	0.0740
	(0.3868)	(0.4427)	(0.4044)	(0.5137)	(0.3841)	(0.5128)	(0.5140)	(0.5406)
$ExRet_{-126}$	-0.0895	0.0283	0.0901	0.2135	-0.0835	0.0518	0.0447	0.1554
	(0.5250)	(0.7075)	(0.6386)	(0.8402)	(0.4834)	(0.8752)	(0.8149)	(0.9149)
$ExRet_{+5}$	-0.0025	-0.0029	-0.0028	-0.0008	-0.0031	-0.0041	-0.0027	-0.0018
	(0.1040)	(0.0921)	(0.0843)	(0.0830)	(0.0802)	(0.0853)	(0.0835)	(0.0897)
$ExRet_{+15}$	-0.0024	0.0000	-0.0011	-0.0006	0.0009	0.0002	-0.0007	0.0001
	(0.0895)	(0.0824)	(0.0768)	(0.0773)	(0.0758)	(0.0727)	(0.0741)	(0.0778)
$ExRet_{+21}$	-0.0050	-0.0034	-0.0034	-0.0021	-0.0024	-0.0011	-0.0013	-0.0011
	(0.0903)	(0.0891)	(0.0835)	(0.0837)	(0.0795)	(0.0820)	(0.0794)	(0.0826)
$ExRet_{+63}$	-0.0058	-0.0028	-0.0054	-0.0074	-0.0024	-0.0030	-0.0055	-0.0052
	(0.1582)	(0.1631)	(0.1460)	(0.1504)	(0.1325)	(0.1389)	(0.1371)	(0.1437)
$ExRet_{+126}$	-0.0072	-0.0135	-0.0067	-0.0051	-0.0087	-0.0030	-0.0129	-0.0113
	(0.2746)	(0.2876)	(0.2516)	(0.2733)	(0.2286)	(0.2529)	(0.2376)	(0.2546)

Table 5: CAPM and Four-Factor Model Estimations for Portfolios of Repurchased Stocks

This table contains the estimation results of the CAPM and Four-Factor Model for portfolios constructed using 1 dollar to purchase each stock repurchased, prior winners and losers respectively, and hold the position for one year. Regressions (1) to (4) are for all investors, (5) to (8) for the investors in the first quartile of average trading frequency, (9) to (12) for the investors in the fourth quartile of average trading frequency. R_t^m is the CRSP value-weighted portfolio return; R_t^f is the one-month T-bill rate; SMB_t is the average return on the small cap portfolio minus the average return on the big cap portfolio; HML_t is the average return on the value stock portfolio minus the average return on the growth portfolio; UMD_t is the average return on the recent winner portfolio minus the average return on the recent loser portfolio. Standard errors are in parentheses. (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$)

		Coefficients					
		Constant	R_t^m-R_t^f	SMB_t	HML_t	UMD_t	F stat
All Investors	(1)Losers	0.0027***	0.1000***				8.49
	CAPM	(0.0002)	(0.0343)				
	(2)Losers	0.0026***	0.2603***	0.2736***	0.1130**	-0.2434***	7.99
	4-Factor	(0.0001)	(0.0539)	(0.0546)	(0.0457)	(0.0664)	
	(3)Winners	-0.0058***	0.0463**				4.22
	CAPM	(0.0001)	(0.0225)				
	(4)Winners	-0.0058***	0.0764**	0.0563*	-0.0151	-0.0900***	3.04
	4-Factor	(0.0001)	(0.0302)	(0.0336)	(0.0266)	(0.0327)	
AvgFreq 1st Quartile	(5)Losers	0.0012	0.1723				0.9
	CAPM	(0.0009)	(0.1817)				
	(6)Losers	0.0012	0.3396	0.3062	0.0436	-0.3101	0.79
	4-Factor	(0.0009)	(0.2418)	(0.2759)	(0.3021)	(0.2496)	
	(7)Winners	-0.0074***	0.0273				0.3
	CAPM	(0.0002)	(0.0496)				
	(8)Winner	-0.0074***	0.0208	0.0236	-0.0112	0.0362	0.24
	4-Factor	(0.0002)	(0.0609)	(0.0765)	(0.0755)	(0.0582)	
AvgFreq 4th Quartile	(9)Losers	0.0036***	0.1177***				7.14
	CAPM	(0.0002)	(0.0441)				
	(10)Losers	0.0036***	0.3302***	0.3589***	0.1392**	-0.3410***	7.14
	4-Factor	(0.0002)	(0.0715)	(0.0737)	(0.0581)	(0.0903)	
	(11)Winners	-0.0049***	0.0531**				4.99
	CAPM	(0.0001)	(0.0238)				
	(12)Winners	-0.0049***	0.0816**	0.0573	-0.0116	-0.0766**	2.58
	4-Factor	(0.0001)	(0.0323)	(0.0361)	(0.0280)	(0.0341)	

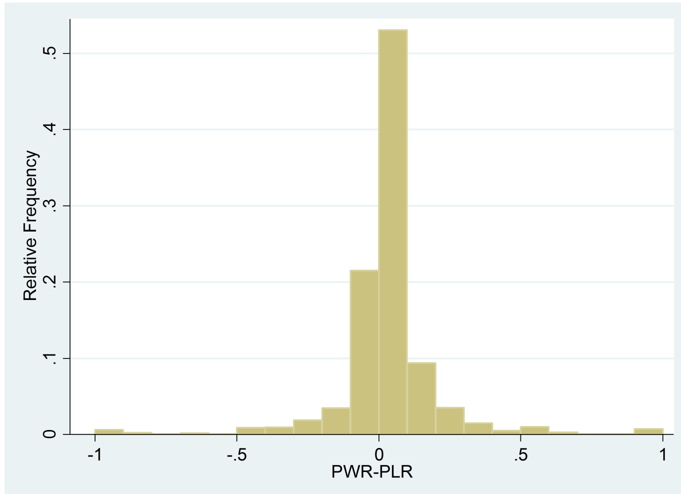

Figure 1: HISTOGRAM OF $PWR - PLR$

Note: This figure plots the histogram of $PWR - PLR$ of the 19,059 households within the sample for whom the difference between PWR and PLR can be calculated. This difference ranges from -1 to 1, with the width of each bin being 0.1 in the histogram.

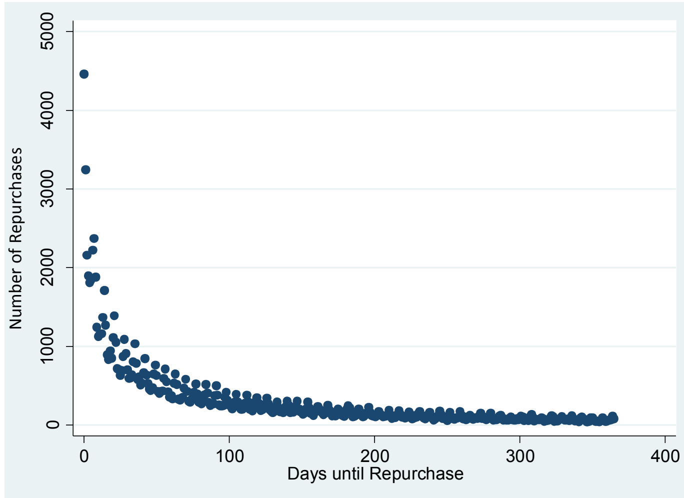

Figure 2: Days to Repurchases

Note: This figure is a scatter plot of the number of repurchases against the number of days till repurchase. The repurchases are grouped according to the number of days elapsed between the original sale and the repurchase. Then the number of repurchases is counted for each group. The number of days ranges from 0 to 365 with 1-day increments. Therefore, only repurchases made within a year subsequent to the sale are included.

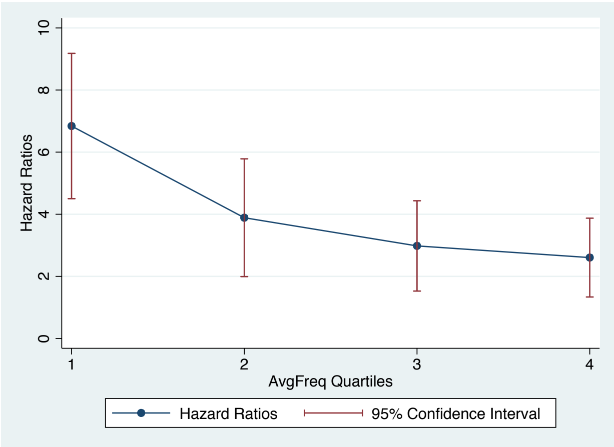

Panel A Sensitivity to Prior Winners

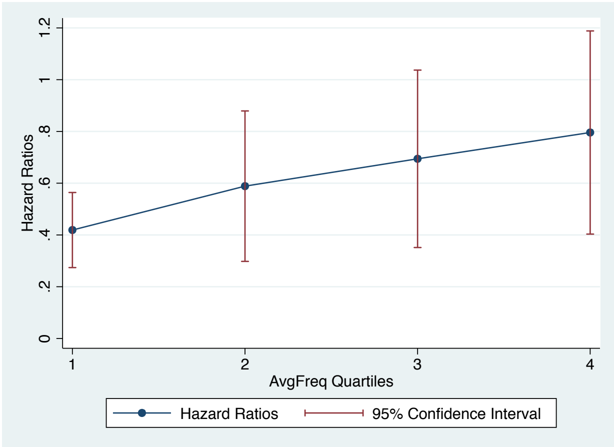

Panel B Sensitivity to Prior Losers

Figure 3: The Effects of Trading Frequency on Repurchasing

Note: This figure shows the relationship between trading frequency and the investors sensitivity to prior winners and losers in Panel A and Panel B respectively. The horizontal axis corresponds to the quartiles of average trading frequency, and their effects on repurchases are measured by the survival analysis hazard ratios on the vertical axis. The 95% confidence interval bars were added.