



LABORATORY FOR
FINANCIAL ENGINEERING

When Do Investors Freak Out? Machine Learning Predictions of Panic Selling

by Daniel Elkind, Kathryn Kaminski, Andrew W. Lo, Kien Wei Siah, and Chi Heem Wong

Authors' final manuscript as accepted for publication

Citation	Elkind, Daniel, Kathryn Kaminski, Andrew W. Lo, Kien Wei Siah, and Chi Heem Wong (2022), "When Do Investors Freak Out? Machine Learning Predictions of Panic Selling," <i>Journal of Financial Data Science</i> 4(1), 11–39.
As Published	https://doi.org/10.3905/jfds.2021.1.085
Publisher	Pageant Media Ltd.

When Do Investors Freak Out? Machine Learning Predictions of Panic Selling^{*}

Daniel Elkind^a, Kathryn Kaminski^a, Andrew W. Lo^{a,b,c,d,1}, Kien Wei Siah^{a,c,d},
Chi Heem Wong^{a,c,d}

^a*MIT Laboratory for Financial Engineering, 100 Main Street, E62-611, Cambridge, MA
02142*

^b*MIT Sloan School of Management, 100 Main Street, E62-618, Cambridge, MA 02142-1347*

^c*MIT Computer Science & Artificial Intelligence Laboratory, 32 Vassar St, Cambridge, MA
02139*

^d*MIT Department of Electrical Engineering and Computer Science, 50 Vassar St,
Cambridge, MA 02142*

*Research support from the MIT Laboratory for Financial Engineering is gratefully acknowledged. The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of any other organizations, any of their affiliates or employees, or any of the individuals acknowledged.

Email address: alo-admin@mit.edu (Andrew W. Lo)

¹Corresponding author

When Do Investors Freak Out? Machine Learning Predictions of Panic Selling

Abstract

Using a novel dataset of 653,455 individual brokerage accounts belonging to 298,556 households, we document the frequency, timing, and duration of panic sales, which we define as a decline of 90% of a household account's equity assets over the course of one month, of which 50% or more is due to trades. We find that a disproportionate number of households make panic sales when there are sharp market downturns, a phenomenon we call 'freaking out.' We show that panic selling and freak-outs are predictable and fundamentally different from other well-known behavioral patterns such as overtrading or the disposition effect.

Highlights

- We define a heuristic to identify investors who sell off a substantial proportion of their risky assets abruptly, or panic sell. We find that investors panic sell more often when there are large market movements.
- Panic selling is suboptimal if executed in an improving market, but it is beneficial as a stop-loss mechanism in rapidly deteriorating markets.
- Artificial intelligence techniques can identify individuals at risk of panic selling in the near future when given the demographic characteristics of the investor, their portfolio histories, and current and past market conditions.

Keywords: Deep learning, Freaking out, Panic selling, Stop-loss, Tactical asset allocation
JEL Classification: G40

1. Introduction

Financial advisors have long advised their clients to stay calm and weather any passing financial storm in their portfolios. Despite this, a percentage of investors tend to ‘freak out’ and sell off a large portion of their risky assets in certain adverse market environments. This situation is often discussed in the financial press and media², but it is rarely defined or quantified. In this paper, we develop a method to identify panic selling and apply it to a novel large dataset of brokerage account information from 2003 to 2015 to examine panic selling and freak-out behavior.

We begin by characterizing the aggregate behavior of investors who make panic sales. First, we document the frequency and timing of panic selling. We see that, although panic sales are infrequent, with only 0.1% of the investors panic selling at any point in time, they occur at up to three times the baseline frequency when there are large market movements.

Second, we find that 30.9% of the investors who panic sell never return to reinvest in risky assets; however, of those who do, more than 58.5% reenter the market within half a year.

Third, we analyze the demographics of the investors who tend to ‘freak out’ under our definition (i.e., they make panic sales during periods of sharp market downturns), and find that a higher proportion of investors who are male, above the age of 45, married, have a greater number of dependents, or have declared themselves to have excellent investment experience or knowledge tend to freak out.

Fourth, we find that the median investor earns a zero to negative return after panic selling. Calculating the opportunity cost of panic selling over time reveals that panic selling is suboptimal if executed in an improving market, but it is beneficial as a stop-loss mechanism in rapidly deteriorating markets.

Finally, we develop machine learning models to predict when investors might panic sell in the near future. Our set of predictive features includes the demographic characteristics of the investor, the investor’s portfolio histories, and current and past market conditions. This task is made difficult by the extreme rarity of panic sales. Nonetheless, our best-performing deep neural network achieves a 69.5% true positive accuracy rate and an 81.2% true negative accuracy rate, demonstrating that artificial intelligence techniques can assist in identifying individuals at risk of panic selling in the near future.

²Consider the typical CNBC headline, “The market may be swinging, but the last thing you should do is freak out: Wall Street trading coach”. Source: <https://www.cnbc.com/2018/02/09/the-market-may-be-swinging-but-dont-freak-out-says-trading-coach.html>

2. Background

Behavioral finance has documented a wide range of stylized actions of investors, including loss aversion, regret aversion, the snake-bite effect, overtrading, and the disposition effect (see, e.g., Odean (1998); Weber and Camerer (1998)). There has been renewed interest in these behavioral patterns since the financial crisis of 2007–2008, both within the academic community and among the general public. We summarize several documented investor behaviors, some that are related to panic selling and others that are inconsistent with the phenomenon.

2.1. Panic Selling

Although widely discussed in the financial industry (see, e.g., Rotblot (2004)), little of the available literature discusses the concept of panic selling during a period of lowered market performance. This is most likely due to the limited availability of datasets that cover a wide range of selling events and market environments. Using price and volume information as well as data from Chinese stock markets, Shi et al. (2011) provided a theoretical model based on conditioning to explain investor behavior. Their model shows that investors can be either overconfident or panicked based on price momentum. The strongest positive correlation in behavior occurs during price reversals, when many investors are more likely to sell their risky assets in a panic.

In contrast to panic selling, however, Statman et al. (2006) found that share turnover is positively correlated with lagged returns, which suggests overconfidence is a dominant factor. Barber et al. (2009) demonstrated that investors tend to buy stocks with strong recent performance and higher trading volume.

Our paper presents evidence that investors occasionally panic and sell off a large portion of their portfolio. It attempts to address the aforementioned issues by using a larger and more fine-grained dataset over a longer time horizon than earlier studies. In this way, we hope to capture a broader range of circumstances under which investors may make panic sales.

2.2. Overtrading and the Disposition Effect

Overtrading is in some ways the opposite of panic selling, which causes the investor to leave the market, either temporarily or permanently. Several authors have documented that some investors tend to overtrade. For example, Benos (1998) and Odean (1998) suggested that overconfidence causes investors to trade too frequently. Using trading account data, Barber and Odean (2000) documented overtrading and demonstrated that it is detrimental to the wealth of those investors who trade too frequently.

In addition, much of the behavioral finance literature has focused on the disposition effect (see Shefrin and Statman (1985)), which is the tendency for investors to buy stocks

with strong recent performance and hold onto their losing investments. This can also be considered as another near-opposite to panic selling.

2.3. Stop-loss

A parallel to panic selling can be found in the use of stop-loss strategies by investors. Stop-loss strategies are rules used by investors to reduce their holdings in risky assets should the value of their holdings reach a certain predetermined threshold. Kaminski and Lo (2014) and Lo and Remorov (2017) examined the value of these rules under different market conditions. In some situations (for example, if market prices exhibit momentum), stop-loss strategies may outperform buy-and-hold strategies over certain time horizons. This, however, depends on the condition that investors return to the market at some point (see Klement (2013)). Because the performance of a stop-loss rule is dependent on investor reentry, in our empirical analysis, we also examine investor reentry after a freak-out.

2.4. Stock Market Crashes and Investor Overreaction

Many authors have studied stock market crashes. From an asset pricing perspective, several authors have demonstrated that rare disaster risk can explain the equity risk premium and other puzzles in macro-finance (e.g., Barro (2006, 2009); Bollerslev and Todorov (2011); Gabaix (2012); Pindyck and Wang (2013)). Many others examined the impact of tail risk on total market returns (see, e.g., Bali et al. (2009); Kelly and Jiang (2014)).

The more relevant question for panic selling, however, is how investors behave during a stock market crash. De Bondt and Thaler (1985, 1987) argued that investors tend to overreact to large economic events. Chopra et al. (1992), Rozeff and Zaman (1998), Bauman et al. (1999), Wang et al. (2013) and others provided empirical evidence for investor overreaction. There are many different explanations for overreaction during a market crash (for a summary, see Amini et al. (2013)). These include changes in investor sentiment (see Baker and Wurgler (2006)), herding behavior (see Puckett and Yan (2008)), market constraints (see, e.g., Park (1995); Kaul and Nimalendran (1990)), and changing risk preferences (see Brown et al. (1993)).

Similarly, many studies of investor overreaction focus on price changes, but only in some cases do they use survey information to document these possible behavioral factors, and few studies have access to the entire portfolio of investors to consider their actual portfolio decisions.

3. Data Summary

We analyze the financial activity of 653,455 anonymous accounts corresponding to 298,556 households from one of the largest brokerage firms in the United States. These accounts are

drawn at random from the population of U.S. brokerage accounts active as of December 31, 2015, and have had their account identification numbers fully anonymized.

Our dataset consists of (i) monthly snapshots of positions and balances held in every sampled portfolio, (ii) every trade made through these accounts, and (iii) the demographic information of the account holder as reported on the initial application form, including age, income, and self-declared levels of experience and knowledge. The kinds of assets contained in the accounts include stocks, mutual funds, options, fixed-income securities, and cash securities. Details of the composition of our dataset are included in the Supplementary Materials for the sake of brevity in the main text of this article. We have also been given a map from accounts to households that allows us to aggregate activities of related accounts. In this study, we analyze panic selling at the household level.

A household will consist of one or more individual accounts opening and closing at different points in time. We consider the time when the first individual account is opened as the account opening date of the household. Because the data given to us only record activities starting from January 2003, all households that were active prior to this date will be reported as though they were started on January 2003. Although all the household accounts in our sample have at least one account that is open at the time of the study, there are some households who sold their assets and decided not to return to the market. We call these ‘inactive’ households. The number of active household accounts at time t , N_t , can be computed recursively with the following formula:

$$N_t = N_{t-1} + n_t^o - n_t^p \tag{1}$$

where n_t^p and n_t^o denote the number of households that panic sold and opened at time t , respectively. Exhibit 1 shows the cumulative number of household accounts that opened, exited the market, and were active over time.

4. Methodology

4.1. Identifying panic sells

It is typically understood that an investor panic sells when he or she *intentionally* sells off a substantial proportion of his risky assets abruptly. We develop rules in order to systematically capture such behavior.

Consider a situation in which we are given monthly snapshots of portfolios and a view of every trade. Let V_t be the value of an investor portfolio at month t and $x_t = \frac{V_t - V_{t-1}}{V_{t-1}}$ be the percentage change in value of the portfolio between month t and $t + 1$. Let T_t be the sum of value of all the trades in month t . Then $t_t = \frac{T_t}{V_{t-1}}$ is the proportion of the portfolio traded in this month. A positive t_t denotes a net buy and a negative t_t denotes a net sell. An investor is said to have made a panic sell in month t when

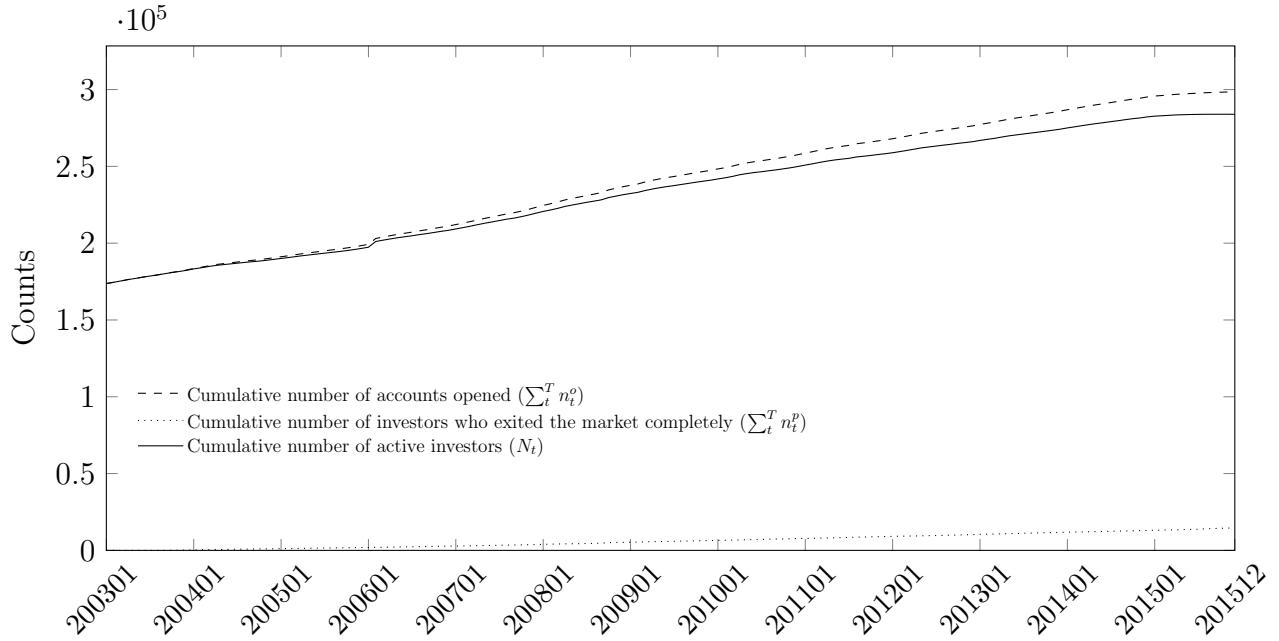


Exhibit 1: Number of household accounts versus time (YYYYMM format).

- Condition (1): The value of the portfolio declines by at least p_1 over one month (i.e. $x_t \leq -p_1$ for some $p_1 > 0$); and
- Condition (2): The investor makes a net sell of p_2 of the portfolio within the same period (i.e. $t_t \leq -p_2$ for some $p_2 > 0$).

Condition (1) states that the value of the portfolio falls substantially between two monthly snapshots of the portfolio. This is a necessary condition for any liquidation. However, it is not a sufficient condition, because large changes in a portfolio may be induced by natural market movements without any action by the investor. To determine that the investor *intentionally* reduced his or her holding of risky assets, we impose Condition (2). At first glance, it may seem that Condition (2) alone would be effective in detecting panic selling. This is untrue, however, because a portfolio may have depreciated substantially before an investor sells. For example, suppose that we let p_2 take a value of 0.8 in order to capture a large change in the portfolio. However, if the value of the investor’s portfolio falls 25% from \$100,000 to \$75,000 due to market conditions before the investor liquidates the rest of it, it is impossible to fulfill the condition of p_2 of 0.8. Hence, just using Condition (2) by itself will cause us to miss this liquidation event; however, a lower value of p_2 can be used if we impose Condition (1).

In addition to identifying panic selling, we identify cases in which investors who exited from their risky positions decide *intentionally* to take on risk again. We call such an event a ‘return to market.’ The following rules identify such events:

- Condition (3): The value of the portfolio must reach at least p_3 of the pre-liquidation

value; and

- Condition (4): The investor must have a cumulative net buy of p_4 of the amount that they sold during liquidation.

We set p_1 (the monthly portfolio decline), p_2 (the monthly portfolio net sell), p_3 (the portfolio rebound), and p_4 (the cumulative net buy) as 0.9, 0.5, 0.5, and 0.5, respectively, in this study. Although setting both p_1 and p_2 to lower values will increase the number of panic sales identified, this does not change our analysis or the general pattern exhibited by household investors (see Section Appendix A.6 of the Supplementary Materials).

4.2. Identifying risk factors for liquidations

To understand which groups of investors (G) are more susceptible to panic selling or freak-out events (E), we compute the relative prevalence of the group given an event: $P(G|E)/P(G)$. A number greater than 1 indicates that the group is more likely to have the event compared to other groups, whereas a value less than 1 signals the opposite.

We can test our null hypothesis of $P(G|E) = P(G)$ against the alternative hypothesis of $P(G|E) \neq P(G)$ using the two-proportion Z-test. Let $p_1 = P(G|E)$, $P(G) = p_2$, and the number of investors in each group be n_1 and n_2 , respectively. The test statistic is given by $z = \frac{p_1 - p_2}{SE}$, where $SE^2 = p(1 - p)(\frac{1}{n_1} + \frac{1}{n_2})$ and $p = \frac{n_1 p_1 + n_2 p_2}{n_1 + n_2}$.

5. Results

We counted 36,374 panic sells by 26,852 household investors (9.0% of all households) across a period of 13 years between January 2003 and December 2015, endpoints inclusive. A heat map for panic sales and returns to the market is given in Exhibit 2, and Exhibit 3 shows the distribution of panic sales per household. Of households with at least one panic-selling event, 21,706 of them (80.8%) did so once within our sample period, and 3,081 (11.4%) did so twice. The mean and standard deviation of the number of panic sells per investor are 1.35 and 0.98, respectively. These numbers suggest that we are seeing a behavioral pattern that is different from overtrading.

5.1. When do the investors panic sell?

As can be seen from Exhibit 4, panic sales occur regularly, with a base level around 0.10%. By overlaying the change in S&P 500 value against the plot, however, we noticed that the spikes in the proportion of households panic selling coincide with sharp falls in the stock market. Looking at the top 10 months with the highest proportion of active investors panic selling, we see they include significant stock market events (see Exhibit 5), confirming the common idea that investors freak out in times of market uncertainty. In the rest of the




Exhibit 2: Heat map of panic-selling events and returns to market. Row entries are unique to households. The horizontal axis denotes time in the YYYYMM format. Red denotes a panic sale, and green denotes a return to market for the household.

paper, we collectively refer to these months as ‘crisis’ periods, and panic selling specifically in these months as ‘freak-outs.’

5.2. Returning to the market

We ask the question, “What happens to an investor after he or she panic sells?” As shown in Exhibit 6, as of December 31, 2015, 30.9% of these investors have not taken on risky assets since they freaked out. Of the freak-outs that concluded with the investor reentering the market, 58.5% and 13.1% lasted 1 to 5 months and 6 to 10 months, respectively.

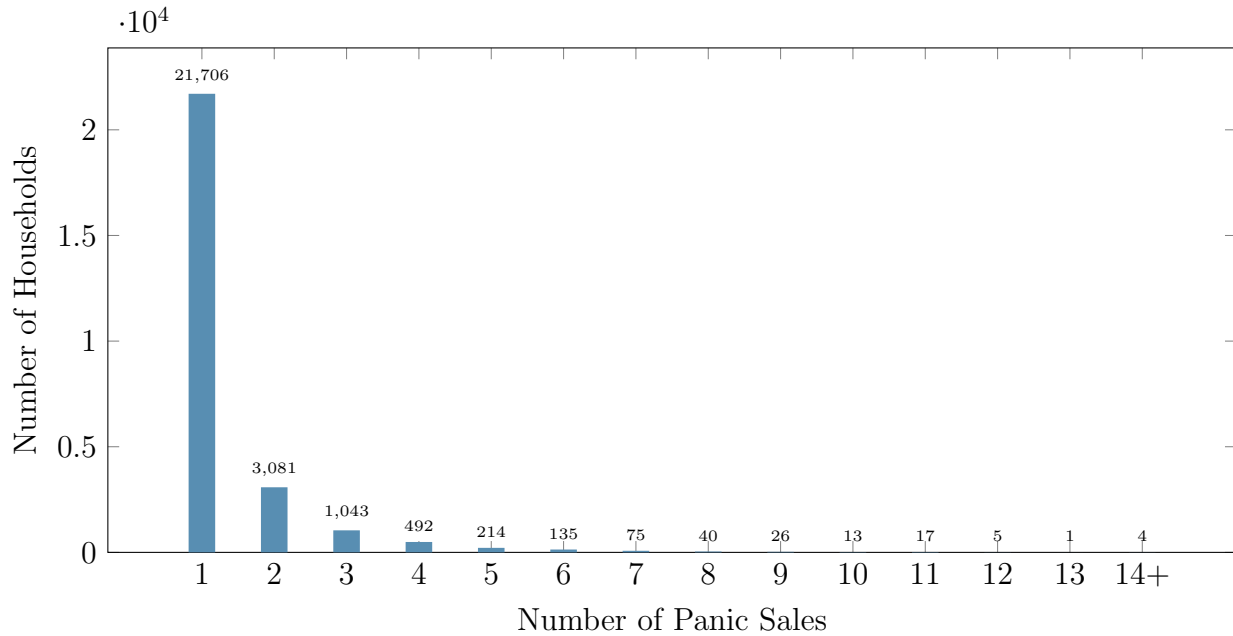


Exhibit 3: Frequency of panic sales. A total of 80.8% and 11.4% of all investors made panic sales once and twice, respectively.

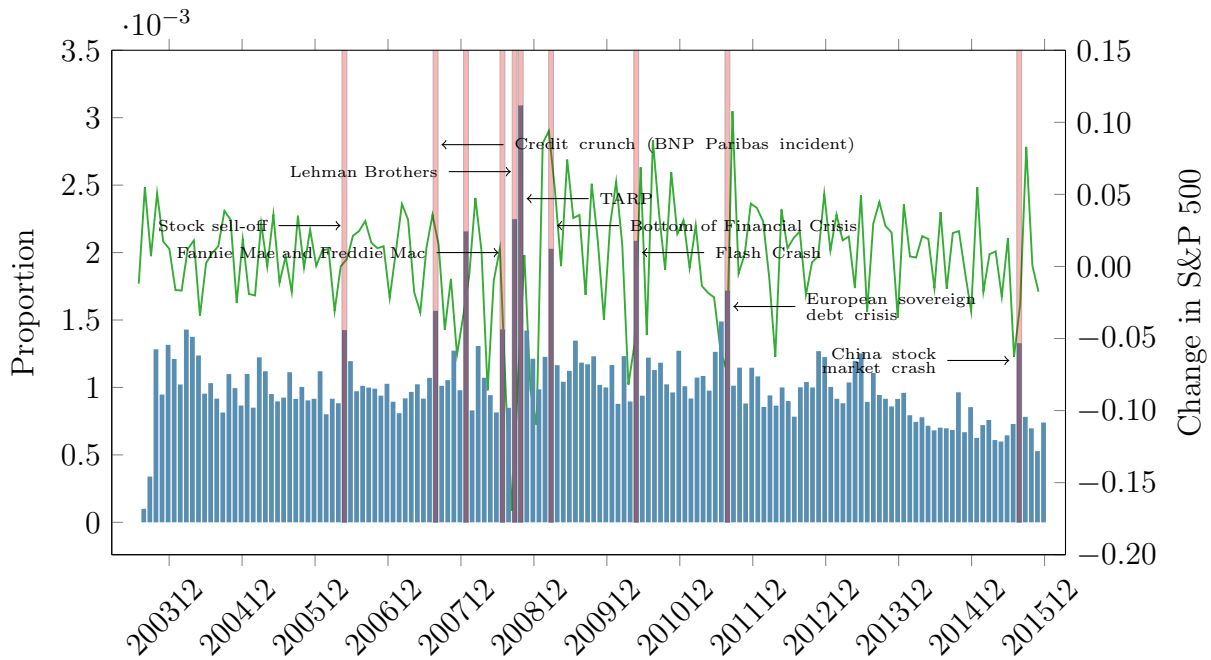


Exhibit 4: The proportion of active households who panic sold in each month (YYYYMM). The green line depicts the month-to-month change in S&P 500 over time.

5.3. Portfolio characteristics of investors who panic sold

Exhibit 7 tabulates the distribution of the value of portfolios just prior to panic sales, and 43.2% of the portfolios are less than \$20,000 in value. The 25th, 50th, 75th and 90th per-

YYYYMM	Counts	Active Accts	Pct Of Active Accts (%)	Known Event
200605	290	203,534	0.142	Stock and commodity sell-off
200708	338	215,628	0.157	Credit crunch (BNP Paribas incident)
200801	476	220,646	0.216	
200807	324	226,614	0.143	Fannie Mae and Freddie Mac
200809	513	228,220	0.225	Lehman Brothers
200810	710	229,703	0.309	TARP
200903	475	234,268	0.203	Bottom of Financial Crisis
201005	512	245,319	0.209	Flash Crash
201108	440	256,098	0.172	European sovereign debt crisis
201508	377	283,900	0.133	China stock market crash

Exhibit 5: Months with the highest relative percentage of liquidations and the corresponding events.

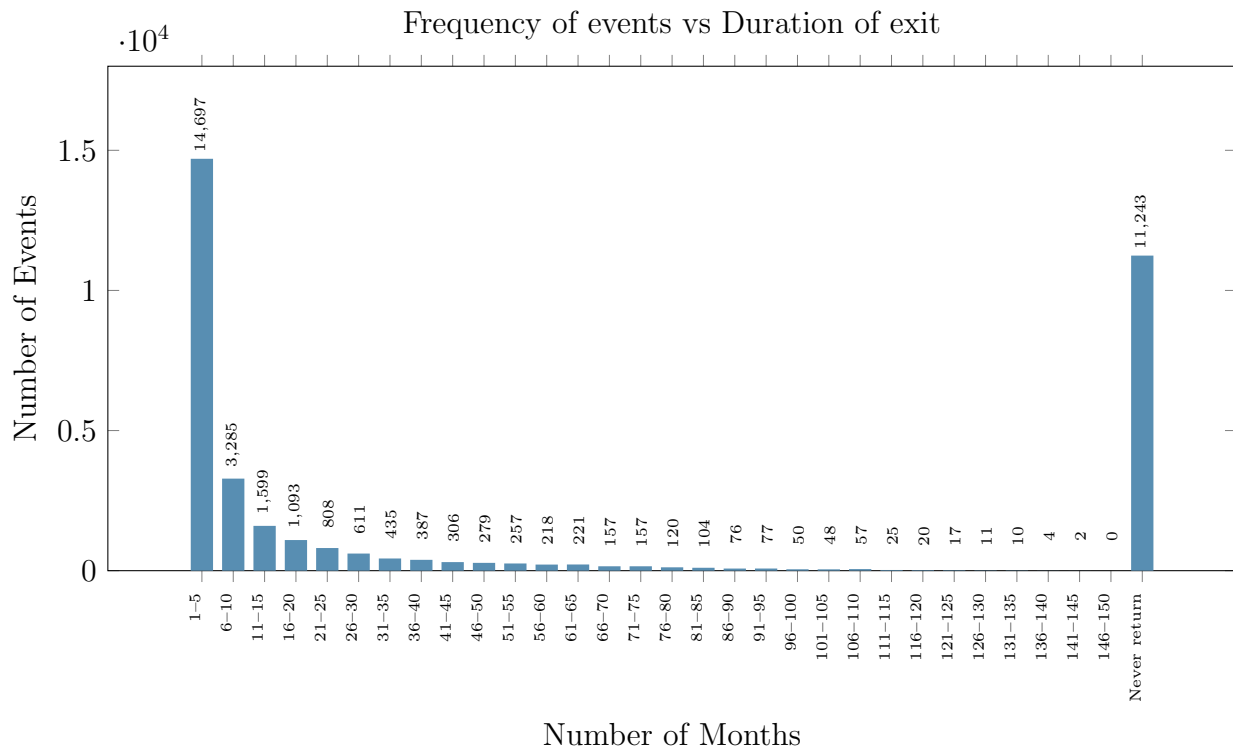


Exhibit 6: Frequency of duration of exit between panic selling and returning to the market.

centile portfolio values are \$7,688.78, \$27,605.35, \$96,387.94, and \$277,986.65, respectively.

5.4. Is panic selling optimal?

Are investors wise to liquidate most of their risky assets over a short period of time? On the one hand, one may subscribe to the view that investors are rational actors who are optimally changing the composition of their portfolio. This behavior can be observed in stop-loss or trailing-stop trades, in which trades are executed to limit further losses when

Portfolio Value	Count	Percentage
0–20,000	15,714	43.20
20,000–40,000	5,284	14.53
40,000–60,000	3,049	8.38
60,000–80,000	1,945	5.35
80,000–100,000	1,549	4.26
100,000–200,000	3,796	10.44
200,000–400,000	2,625	7.22
400,000–600,000	976	2.68
600,000–800,000	479	1.32
800,000–1,000,000	261	0.72
1,000,000– ∞	696	1.91
Total	36,374	100.01%

Exhibit 7: Distribution of portfolio value immediately prior to panic sales. Percentages do not total exactly 100.00 due to rounding.

the market is plunging, or to lock in profits when the market is on the rise. On the other hand, it is possible that investors are panicked by changing market conditions and therefore sell, despite knowing that it is not in their best interest to do so.

5.4.1. *Opportunity cost of panic selling*

We examine whether panic selling is an optimizing behavior by first calculating hypothetical returns over various time horizons in which the investor did not panic sell. That is, we assume that the panicked investor did not sell off risky assets, and track the hypothetical returns of this portfolio until the investor actually returned to the market. We then average the returns to get an aggregate estimate of the potential returns that were forgone. If these hypothetical returns are negative, we conclude that panic selling is an optimizing behavior because it prevented further losses. On the other hand, if the hypothetical returns are positive, this implies that investors could have profited if they simply left the money in their accounts.

Exhibit 8 shows the average hypothetical returns over 20-, 100-, 200-, 600- and 1,000-day periods of investors who liquidated in the tabulated month. Negative values indicate that the average investor would have lost money, whereas positive values indicate that average investors would have been better off had they not liquidated.

We found that the average hypothetical returns are highly correlated with the returns of the S&P 500 over the same time horizons. Our results also suggest that the experience of the individual investor depends on market conditions when they exited, and the duration of the exit.

This point is more obvious in Exhibit 9, where we compute the median of the hypothetical

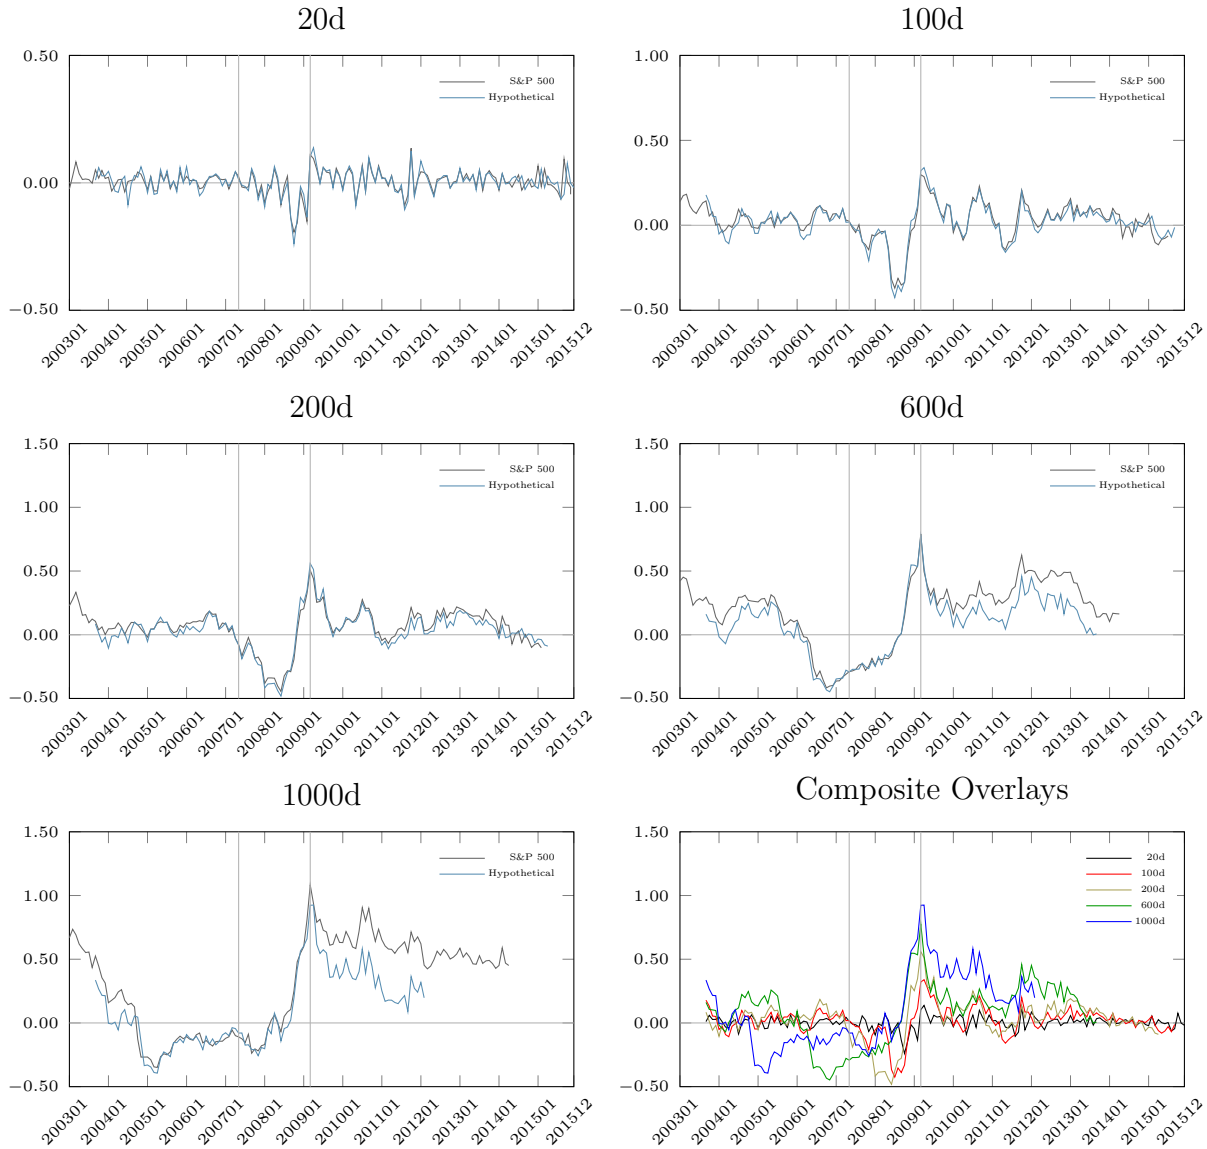


Exhibit 8: Median hypothetical returns of investors who liquidated in a particular month (YYYYMM) over d days. This is constructed by assuming that the investor did not panic sell and held the portfolio for d days.

returns conditioned on the time of their liquidation and the duration of their exit. We plot the kernel regressions to smooth out variations over the time horizons. As can be seen, during the financial crisis, it was typically wise to liquidate one's entire portfolio of risky assets over the short to medium term (less than 35 months). Individuals who liquidated at the start of the crisis and then left the market for 15 months at that time saved themselves from losing another 17%. Holding out for more than 34 months after liquidation, however, would have caused the investor to miss the post-2009 market rally and forgo potential profits.

The reverse is true after the financial crisis. Notably, the persistent rally of the financial markets after the financial crisis ensured that investors who liquidated then would pay a

high price in terms of opportunity cost.

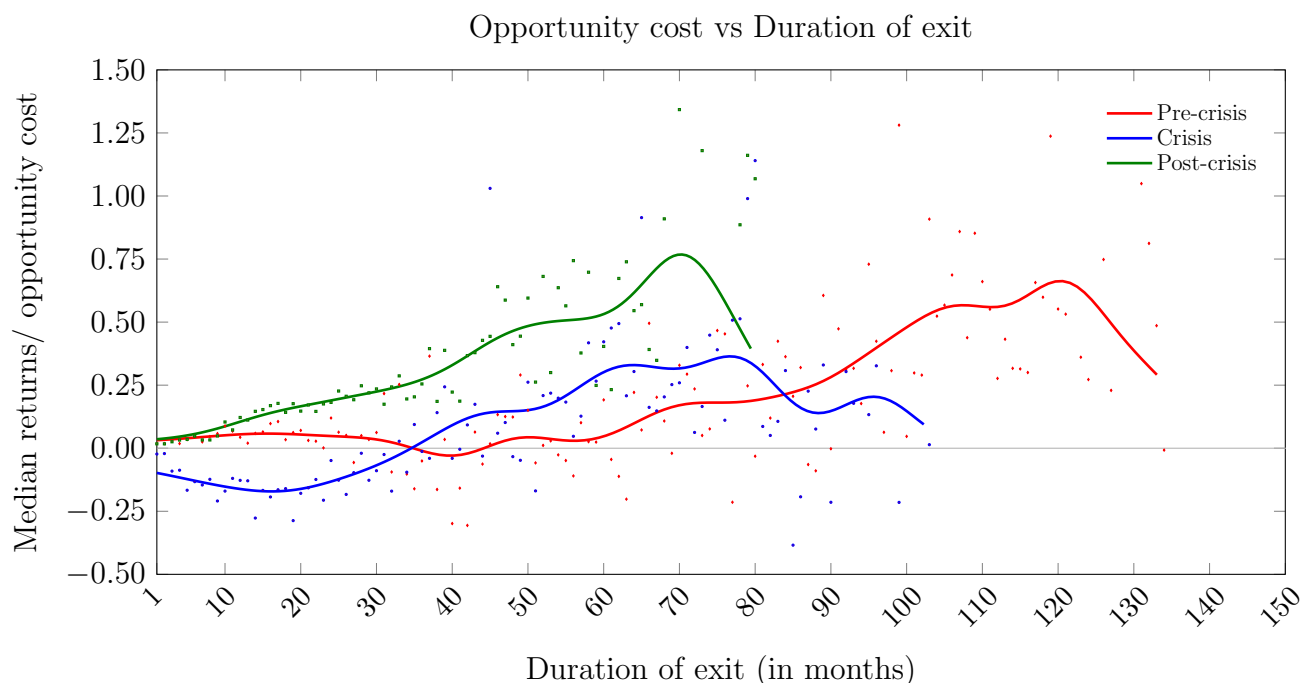


Exhibit 9: Median return of investors under the assumption that they held their portfolio over the duration of their exit. We define the pre-crisis, crisis, and post-crisis periods to be Jan 2003–Apr 2007, May 2007–Feb 2009, and Mar 2009–Dec 2015, respectively. The smoothed lines are kernel regressions of the individual series. The number of data points drops exponentially with the duration of staying out (see Exhibit 6). Thus, values for a duration > 60 months are based on only a few data points.

5.4.2. Performance during liquidation

It might be argued that investors who made panic sales did so strategically, which in turn gave them better returns than the market. For example, they could have kept their outperforming stocks while selling the bulk of the underperforming ones, or invested the proceeds of the sales in assets with higher returns. Exhibit 10 shows that this is typically not the case. The median investor trades infrequently and makes zero to negative returns when out of the market for periods between one month and five years.

5.5. Demographic profile of investors

In this section, we profile the demographic characteristics of investors who liquidated significant parts of their portfolio.

All demographic information in our dataset, with the exception of age, reflects the customer profile at the time the brokerage accounts were opened and is not updated over time. Although this may produce inaccuracies in our analysis, we believe that it can still generate

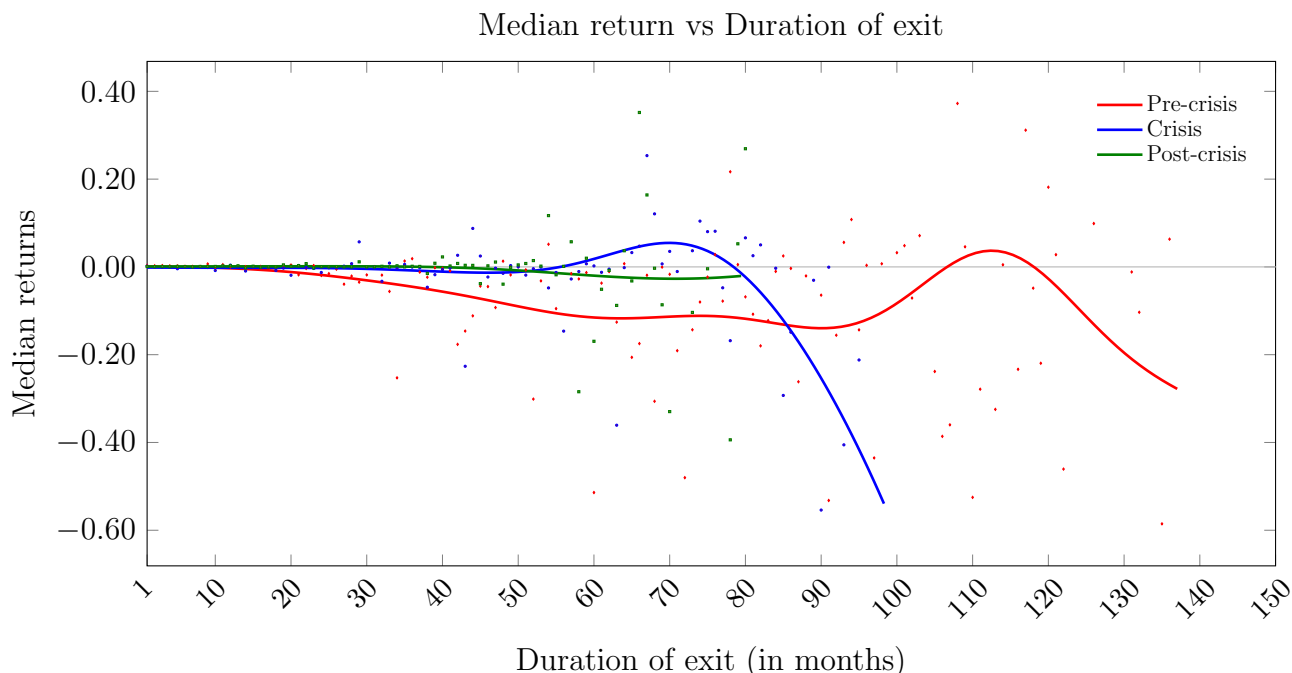


Exhibit 10: Median returns of investors based on their investment actions over the duration of their exit. We define the pre-crisis, crisis, and post-crisis periods to be Jan 2003–Apr 2007, May 2007–Feb 2009, and Mar 2009–Dec 2015, respectively. The smoothed lines are kernel regressions of the individual series. The number of data points drops exponentially with the duration of exit (see Exhibit 6). Values for a duration > 60 months are thus based on only a few data points.

insights as to which investors are more likely to panic sell. Certain fields are missing for some investors because demographic information is collected on a voluntary basis.

Owing to the structure of the data, in which a household can contain multiple investing accounts and multiple investing accounts can share a set of customers (please refer to Section Appendix A.3 in the Supplementary Materials for more detail), care has to be taken to analyze customer demographics. For each customer in a household, we compute fractional weights based on the size of the portfolios to which they are linked. The computational method is detailed in Section Appendix A.5 in the Supplementary Materials.

Some floating point values may be imprecise in the tables because we only give the results to two decimal places.

Age. As can be seen in Exhibit 11, people between the ages of 45 and 100 have a heightened tendency to make panic sales, both across the entire sample and during crisis periods. Younger investors are less likely to make panic sales by a wide margin.

Marital status. Exhibit 12 shows that investors who are married or divorced are more likely than other groups to freak out across the entire sample period.

Age group	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Missing	3,557.12	759.15	71,174.87	0.56 ⁺	0.55 ⁺
age <21	136.99	18.14	2,110.69	0.72 ⁺	0.44 ⁺
21 ≤ age <25	170.99	25.70	2,819.12	0.68 ⁺	0.47 ⁺
25 ≤ age <30	466.03	56.38	6,732.05	0.77 ⁺	0.43 ⁺
30 ≤ age <35	796.74	103.96	10,414.13	0.85 ⁺	0.52 ⁺
35 ≤ age <40	869.40	134.75	11,447.19	0.85 ⁺	0.61 ⁺
40 ≤ age <45	1,269.09	241.84	14,257.73	0.99	0.88
45 ≤ age <50	2,226.95	454.52	20,873.71	1.19 ⁺	1.12 ⁺
50 ≤ age <55	3,051.63	663.69	27,346.63	1.25 ⁺	1.25 ⁺
55 ≤ age <60	3,475.76	756.39	30,437.99	1.27 ⁺	1.28 ⁺
60 ≤ age <65	3,381.66	736.02	29,437.30	1.28 ⁺	1.29 ⁺
65 ≤ age <70	3,013.68	730.55	26,510.99	1.27 ⁺	1.42 ⁺
70 ≤ age <75	2,042.53	483.52	18,938.85	1.20 ⁺	1.32 ⁺
75 ≤ age <80	1,164.50	298.71	11,811.12	1.10 ⁺	1.31 ⁺
80 ≤ age <85	638.12	177.44	7,442.69	0.96	1.23 ⁺
85 ≤ age <90	359.30	95.05	4,728.17	0.85 ⁺	1.04
90 ≤ age <95	164.56	46.08	2,238.36	0.82	1.06
95 ≤ age <100	45.86	15.08	653.64	0.78	1.19
100 ≤ age <infy	21.08	4.03	234.76	1.00	0.89
Total	26,852	5,801	299,610		

Exhibit 11: Distribution of investors by age groups. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of other groups. ⁺ indicates significance at the 1% rejection level.

Category	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Separated	14.88	1.00	191.57	0.87	0.27
Minor	96.80	15.02	1,475.72	0.73 ⁺	0.53
Widowed	333.90	72.97	4,821.70	0.77 ⁺	0.78
Missing	7,868.55	1,734.12	113,814.44	0.77 ⁺	0.79 ⁺
Single	4,496.66	901.05	47,592.89	1.05 ⁺	0.98
Divorced	1,187.70	249.79	11,464.65	1.16 ⁺	1.13
Married	12,853.50	2,827.05	120,249.03	1.19 ⁺	1.21 ⁺
Total	26,852	5,801	299,610		

Exhibit 12: Distribution of investors by marital status. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of other groups. ⁺ indicates significance at the 1% rejection level.

Gender. We note that many investors in our dataset do not provide gender information. Among those who volunteered this information, males compose 56.2% of the sample. Previous behavioral finance studies have typically recorded a disproportionate proportion of males (Barber and Odean, 2001). Our analysis shows that males are slightly more likely than females to freak out (i.e. panic sell during periods of high financial stress) but are less likely to panic sell in general (Exhibit 13).

Number of dependents. Among those with known information about dependents, investors with no dependents are least likely to panic sell (Exhibit 14). There seems to be a positive correlation between the likelihood of panic selling and the number of dependents.

Gender	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Female	378.92	104.19	5,943.39	0.71 ⁺	0.91
Missing	25,822.43	5,525.99	286,025.14	1.01 ⁺	1.00
Male	650.66	170.82	7,641.47	0.95	1.15
Total	26,852	5,801	299,610		

Exhibit 13: Distribution of investors by gender. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of the other groups. ⁺ indicates significance at the 1% rejection level.

Number of Dep.	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Missing	3,532.45	754.99	70,801.04	0.56 ⁺	0.55 ⁺
0	14,808.74	3,186.52	156,436.02	1.06 ⁺	1.05 ⁺
1	3,090.31	670.57	26,994.69	1.28 ⁺	1.28 ⁺
2	3,055.74	710.32	27,584.11	1.24 ⁺	1.33 ⁺
3	1,514.31	319.54	11,957.38	1.41 ⁺	1.38 ⁺
4	587.07	110.88	4,117.34	1.59 ⁺	1.39 ⁺
≥5	263.39	48.19	1,719.41	1.71 ⁺	1.45
Total	26,852	5,801	299,610		

Exhibit 14: Distribution of investors by number of dependents. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of the other groups. ⁺ indicates significance at the 1% rejection level.

Self-declared investing experience. Exhibit 15 shows that the likelihood of panic sales and freak-outs is most pronounced when the investor has self-declared good or excellent investing experience. Interestingly, those for whom we lack this information, and those who declared themselves to have no investment experience, are less likely to panic sell or freak out.

Category	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Missing	5,281.71	1,124.01	89,774.51	0.66 ⁺	0.65 ⁺
None	2,044.11	395.43	24,317.28	0.94 ⁺	0.84 ⁺
Decline to report	853.97	163.33	9,531.72	1.00	0.89
Limited	8,972.14	1,869.94	98,277.98	1.02	0.98
Good	7,216.77	1,631.21	61,775.62	1.30 ⁺	1.36 ⁺
Excellent	2,483.30	617.08	15,932.89	1.74 ⁺	2.00 ⁺
Total	26,852	5,801	299,610		

Exhibit 15: Distribution of investors by investment experience. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of the other groups. ⁺ indicates significance at the 1% rejection level.

Self-declared investing knowledge. Similar to investing experience, we find that investors who describe their investment knowledge as good or excellent panic sell or freak out in higher proportions compared to their baselines (Exhibit 16).

Occupational Group. The occupational groups with the three highest risks of panic selling are self-employed, owners, and real estate, and the three occupational groups with the least risk of panic selling are paralegal, minor, and social worker (Exhibit 17).

Category	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Missing	8,757.51	1,581.76	131,408.17	0.74 ⁺	0.62 ⁺
Decline to report	1,083.23	220.11	11,902.07	1.02	0.96
Limited	7,282.50	1,650.57	77,048.98	1.05 ⁺	1.11 ⁺
None	1,750.25	401.47	16,543.69	1.18 ⁺	1.25 ⁺
Good	6,144.44	1,480.24	51,353.35	1.34 ⁺	1.49 ⁺
Excellent	1,834.07	466.85	11,353.75	1.80 ⁺	2.12 ⁺
Total	26,852	5,801	299,610		

Exhibit 16: Distribution of investors by investment knowledge. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of the other groups. ⁺ indicates significant at the 1% rejection level.

Category	(A) Liquidation (Full sample)	(B) Liquidation (Crisis periods)	All investors	Rel. prop for (A)	Rel. prop for (B)
Paralegal	36.35	5.45	576.67	0.70	0.49
Minor	96.81	15.02	1,476.89	0.73 ⁺	0.53
Social worker	16.95	2.99	285.85	0.66	0.54
Missing	5,934.96	1,259.68	96,257.00	0.69 ⁺	0.68 ⁺
Government	30.57	4.06	275.68	1.24	0.76
Police-military	133.37	18.39	1,244.83	1.20	0.76
Artist	173.55	33.68	2,222.40	0.87	0.78
Student	71.22	11.28	741.06	1.07	0.79
Medical	394.31	80.73	5,279.71	0.83 ⁺	0.79
Education	554.67	117.43	6,544.05	0.95	0.93
Skilled labor	1,020.02	171.87	9,574.47	1.19 ⁺	0.93
Scientist	114.69	29.41	1,595.88	0.80	0.95
Secretary	232.82	54.37	2,926.68	0.89	0.96
Unemployed	981.34	197.95	10,505.23	1.04	0.97
Homemaker	765.36	158.54	8,095.60	1.05	1.01
S-skilled office	350.59	77.65	3,898.63	1.00	1.03
Computer	645.39	130.43	6,481.64	1.11 ⁺	1.04
Attorney	426.05	89.00	4,417.51	1.08	1.04
Engineer	288.28	59.57	2,951.42	1.09	1.04
Clergy	32.76	6.06	294.84	1.24	1.06
Cpa	301.43	63.07	3,065.35	1.10	1.06
White-collar	1,162.82	245.28	11,625.54	1.12 ⁺	1.09
Physician	496.35	109.52	5,109.17	1.08	1.11
Retired	3,817.83	908.75	42,210.01	1.01	1.11 ⁺
Pilot	73.48	13.73	624.20	1.31	1.14
Manager	1,449.41	307.67	13,569.83	1.19 ⁺	1.17 ⁺
Marketing	1,142.77	219.33	9,365.87	1.36 ⁺	1.21 ⁺
Executive	1,412.26	293.41	10,756.85	1.46 ⁺	1.41 ⁺
Financial	733.94	163.30	5,903.92	1.39 ⁺	1.43 ⁺
Professional	1,679.33	405.09	14,271.18	1.31 ⁺	1.47 ⁺
Consultant	471.01	126.73	4,427.47	1.19 ⁺	1.48 ⁺
Owner	534.88	118.20	3,963.37	1.51 ⁺	1.54 ⁺
Self employed	967.02	225.22	6,914.92	1.56 ⁺	1.68 ⁺
Real estate	309.39	78.13	2,149.44	1.61 ⁺	1.88 ⁺
Disabled			6.85		
Total	26,852	5,801	299,610		

Exhibit 17: Distribution of investors by occupation groups, as classified by the broker. (A) shows the weights of investors who made panic sales across the entire sample period. (B) shows the weights of investors who freaked out. A proportion less than (greater than) 1 indicates that members of the group are less likely (more likely) to liquidate compared to members of the other groups. ⁺ indicates significance at the 1% rejection level.

6. Prediction of individual panic sells

Using logistic regression and deep neural network techniques, we attempt to predict panic sales for every individual in the next month in advance, given demographic attributes, past

trading patterns, portfolio history, and recent market conditions. In our logistic regression model, we seek a generalized linear model in which the separating hyperplane is linear with respect to the input feature space. This allows an easy interpretation of the coefficients in terms of odd ratios, but the class of functions that it can model accurately is restricted. In our machine learning model, we push the limits of prediction by training neural network models of 5 hidden layers and 15 hidden layers of 60 neurons to find similarities between panic-selling events. However, doing so necessarily sacrifices easy interpretation.³ Despite the drawbacks of each method, we hope to show that there exists significant information in the dataset that would allow us to predict panic selling.

In the rest of this section, we will refer to the occurrence of panic selling in the next month as a positive data point and its absence as a negative data point.

6.1. Construction of training and testing datasets

We created a dataset for machine learning using demographic attributes, portfolio states, and market states. Among the demographic attributes used are age, marital status, number of dependents, self-declared investment experience, self-declared investment knowledge, and occupational group. We assume equal weights for all the customers in a household when assigning scores to the one-hot categories. For example, if a household has three customers with ages 50, 50, and 70, the category ‘Age:50’ will have a score of $2/3$, while the category ‘Age:70’ will have a score of $1/3$.

For portfolio states, we consider the changes in portfolio balance, the distribution of the portfolio (in cash, equities, options, and penny stocks), the nominal and net values of trades, and, the number of trades as functions of time. We incorporated lags of six months to allow the models to easily pick up time-series signals. We use the month-to-month change, the volatility of prices, and the trading volume of the S&P 500 as indicators of market conditions. Market information was downloaded from Yahoo! Finance. We considered lags of 12 months for these market variables.

A summary of the variables is shown in Exhibit 18. For variables that are unbounded from either side (e.g., $\in \mathbb{Z}_+$), we shifted the midpoint value to zero and then scaled them to be within $[-1, 1]$. In total, the inputs into our models are vectors of length 507.

To benchmark the predictive power of the models, we perform a random 60–40 training–testing split. This ratio is maintained for each of the two classes so that the test set is representative of the entire sample. To prevent cross-contamination between the training and testing sets, which would falsely inflate the performance of the models, we ensure that all of a household’s data points are either in the testing set or the training set. We use the training set of investors for both rounds of training, but evaluate the performance of

³Of course, the notion of ‘interpretability’ is itself up for debate. See Doshi-Velez and Kim (2017) for discussions on the topic.

Description	Variable type
<i>Demographics</i>	
Age	Discrete, 83 groups
Marital status	Discrete, 7 groups
Number of dependents	8 groups
Investment experience	Discrete, 6 groups
Investment knowledge	Discrete, 6 groups
Occupation group	Discrete, 35 groups
<i>Portfolio Factors</i>	
Risky assets balance	\mathbb{R}_+
Penny stocks balance	\mathbb{R}_+
Options balance	\mathbb{R}_+
Cash balance	\mathbb{R}_+
Portfolio balance	\mathbb{R}_+
Number of risky assets	\mathbb{Z}_+
Number of penny stocks	\mathbb{Z}_+
Pct of cash in portfolio	[0,1]
Pct of risky asset (value) in portfolio	[0,1]
Pct of penny stocks in risky asset (value)	[0,1]
Pct of options in risky asset (value)	[0,1]
Pct of penny stocks in risky asset (count)	[0,1]
Net value of trades	\mathbb{R}
Nominal value of trades	\mathbb{R}_+
Nominal value of intraday trades	\mathbb{R}_+
Pct of trades involved in intraday trades (value)	[0,1]
Number of trades	\mathbb{Z}_+
Number of intraday trades	\mathbb{Z}_+
Pct of trades involved in intraday trades (abs num)	[0,1]
Net value of trades as percentage of portfolio balance	[0,1]
Nominal value of trades as percentage of portfolio balance	[0,1]
Is the investor in or out of the market	{0, 1}
We consider 6-month running lags for all the portfolio factors except the last	
<i>Market State</i>	
Month-to-month change in the S&P 500	\mathbb{R}
Month-to-month change in the volume traded of GSPC	\mathbb{R}
Volatility of the volume traded of GSPC within the month	\mathbb{R}_+
Volatility of the price over the past 20 days	\mathbb{R}_+
Volatility of the price over the past 60 days	\mathbb{R}_+
Volatility of the price over the past 180 days	\mathbb{R}_+
We consider 12-month running lags for market factors	

Exhibit 18: List of raw variables used to construct the machine learning data set

the models only on the test set to detect overfitting of data points. We do not require a validation set because we do not perform any parameter optimization or model selection.

6.2. Evaluation

Panic sales are rare events. In all, we obtain 25,418,786 data points, of which only 33,226, or 0.131%, are panic sales (the number of panic sales is less than the number reported in the previous section because we wish to create a lagged series, which forces us to drop some data points). This extremely unbalanced dataset poses a significant problem for any binary classification algorithm. For example, if an algorithm made the prediction of ‘not a panic sell’ for any input, it would achieve an accuracy of 99.869%, an eye-popping but practically irrelevant number. To get a better sense of the performance of the models, we compute accuracy rates separately for both the negative and the positive examples. In addition, we display the receiver operating characteristic (ROC) and precision-recall (PR) curves and report the areas under them. An explanation of these metrics is given in Section Appendix A.7.2 in the Supplementary Materials. Because the ROC and PR curves serve to answer different questions, we have included the results for both to allow our readers to decide if the models are useful for their applications.

6.3. Computation

For all models, we use cross-entropy loss and train the models to optimally use batch gradient descent (GD) with Adaptive Momentum (Kingma and Ba, 2014). It can be shown that minimizing the cross-entropy loss yields the maximum likelihood estimate of the parameters. For each batch of 150,000, we randomly draw half of the samples from each of the two classes (with replacement for the positive class and without replacement for the negative class) to prevent the classifiers from over-emphasizing either class, which would be the natural tendency of the classifier had we selected the training batch at random. We terminate the training when we determine that the accuracy and/or loss has been saturated. We note in passing that, given the appropriate training schedule, the solution converged on by GD for the logistic regression will be the global minimum solution with respect to the loss of the training set.

All the models were implemented on Tensorflow 1.6 with CUDA 9.0/CuDNN 7.0, and training was executed on a single Microsoft Azure NC12 instance, which contains 2 Nvidia Tesla K80 GPUs.

6.4. Results

The accuracy curves, ROC curves, and PR curves on the testing set are shown in Exhibits 19, 20, and 21, respectively. As can be seen from Exhibit 19, all the models have been trained to convergence. The neural networks converge after approximately 2,000 steps, whereas the logistic classifier converged after approximately 8,000 steps. There is no evidence that there

is any form of overfitting on the training set, despite the 15-layer neural network containing over 56,000 parameters.

The final accuracy rates, areas under the ROC curves (AUROCs), and areas under the PR curves (AUPRCs) on the testing set for all the models are reported in Exhibit 22. We can see that the neural networks outperform the logistic classifier on all metrics. Between the neural networks, the 15-layer network showed an improvement of 1.3 percentage points over the 5-layer network on the positive data, but a deterioration of 2.5 percentage points on the negative data. We can see that the 5-layer neural network marginally outperforms the 15-layer neural network on the AUROC and AUPRC metrics, but the differences may be simply due to randomness in training. The comparable performance of the neural networks shows that a 5-layer neural network has enough capacity to approximate the function and a larger network is unnecessary.

6.4.1. Interpreting the logistic classifier

We attempt to interpret the coefficients of the logistic classifier. We group the variables according to their classification type (demographic factor vs. market factor vs. portfolio factor) and report the top 10 most important variables according to the absolute value of the weights of the coefficients. This works in our analysis because we have monotonically transformed values to between -1 and 1 . Our results are shown in Exhibit 23.

Age dominates the list of the most important demographic variables. In general, being young or elderly decreases the risk of panic selling. Being disabled or a minor also lowers the likelihood of panic selling. Although not shown, declaring oneself a member of the clergy, an owner, or an executive increases the likelihood of panic selling. In addition, having self-declared excellent investment experience increases the odds of panic selling. These results substantially agree with the analysis by demographic slices in Section 5.5.

Among all the market factors, lagged series of the 20-day S&P 500 volatility, the 60-day S&P 500 volatility, and the volatility of the S&P 500 trading volume are the most important factors in predicting panic sales. The signs of the coefficients are mixed.

Our analysis of the coefficients for the portfolio factor shows that the likelihood of a panic sale increases with the percentage of daily trades made by the investor. Furthermore, an investor will be more likely to panic sell if options compose a larger proportion of the entire portfolio. The liquidation of the portfolio has been added as a variable to help the convergence of the model, and the model accurately deciphered that the chance of a panic sale is high when the portfolio has not been liquidated. This serves as a sanity check that our model is picking up the correct signals.




Exhibit 19: Accuracy curves over training steps. The training of the 15-layer and 5-layer neural networks were terminated at around the 2,650th and 3,150th step, respectively, as we deemed that they have converged. The logistic classifier was terminated at around the 8,000th step.

7. Conclusion

The analyses in this paper hinge on the heuristic we developed to identify panic sales. To test the robustness of our results, we performed additional runs with different parameters. We find that, although decreasing the thresholds will increase the number of panic sales identified across all time periods, there is still a disproportionate number of accounts which panic sell in periods of high financial stress (see Section Appendix A.6 of the Supplementary Materials).

Panic selling and freaking out are distinct behavioral patterns in finance that differ from




Exhibit 20: Receiver operating characteristic (ROC) curves of the trained models.

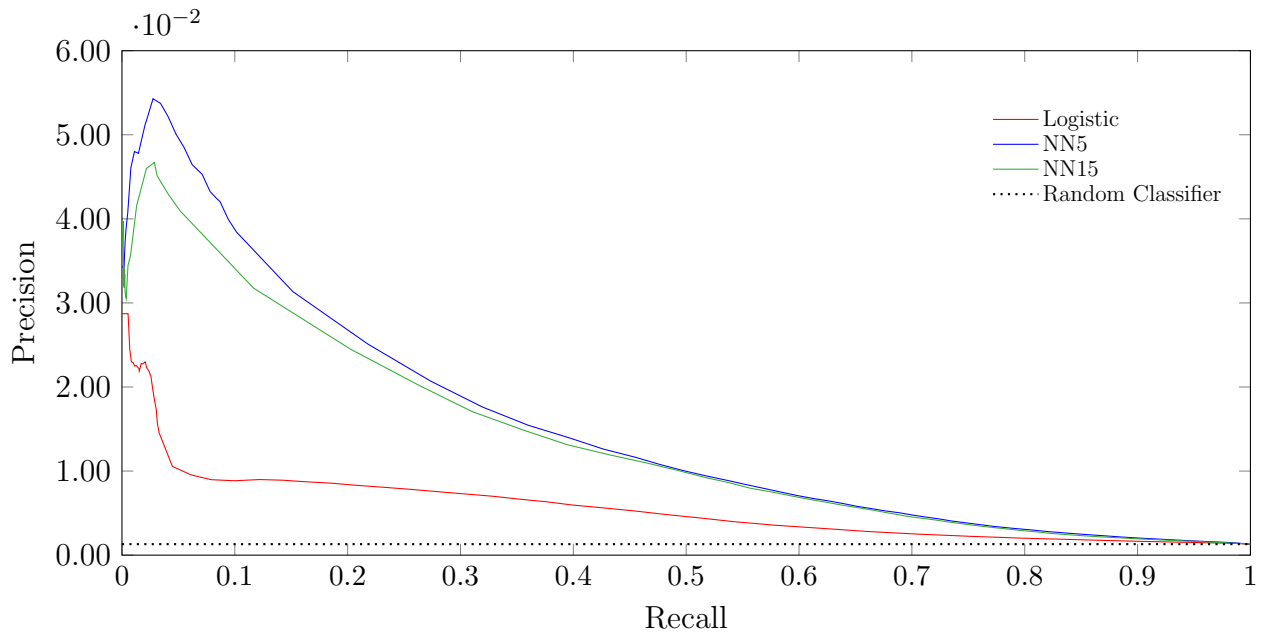


Exhibit 21: Precision-recall (PR) curves of the trained models.

other previously studied patterns. While the disposition effect claims that investors tend to hold on to their losers and keep their winners, we see that investors who made panic sales achieve only a slightly negative return after they liquidate. Furthermore, in contrast to overtrading, investors who made panic sales did so infrequently. We see that panic selling spikes in periods of crisis, suggesting a relationship between panic selling and market conditions. Our logistic model suggests that recent market volatility influences panic-selling behavior.

Model	Accuracy		AUROC	AUPRC ($\times 10^{-3}$)
	Positive samples	Negative samples		
Random Predictor	—	—	0.500	1.307
Logistic Regression	57.9%	78.8%	0.739	5.521
Neural Net (5 hidden layers)	69.5%	81.5%	0.821	15.184
Neural Net (15 hidden layers)	70.8%	79.0%	0.813	13.819

Exhibit 22: Performance of the models on the test set

Variable Name	Description	Coefficient
<i>Demographic factors</i>		
Age:97	Age of 97	-0.885
dmsa_curr_occup_tx SOCIALWORKER	Occupation: Social Worker	-0.759
Age:<21	Age less than 21	-0.642
Age:23	Age of 23	-0.620
Age:26	Age of 26	0.588
Age:94	Age of 94	-0.587
Age:99	Age of 99	-0.586
dmsa_curr_occup_tx DISABLED	Occupation: Disabled	-0.546
dmsa_curr_occup_tx MINOR	Occupation: Minor	-0.531
invst_exprc_cd E	Investment experience: excellent	0.501
<i>Market factors</i>		
60d_price_vol_lag_8	60 days volatility of S&P500 (8 months ago)	0.758
20d_price_vol_lag_4	20 days volatility of S&P500 (4 months ago)	-0.748
20d_price_vol	20 days volatility of S&P500	0.730
volume_vol	Volatility of volume traded in S&P500 across one month	0.723
60d_price_vol_lag_5	60 days volatility of S&P500 (5 months ago)	-0.701
20d_price_vol_lag_6	20 days volatility of S&P500 (6 months ago)	-0.690
20d_price_vol_lag_9	20 days volatility of S&P500 (9 months ago)	0.687
20d_price_vol_lag_11	20 days volatility of S&P500 (11 months ago)	0.680
volume_vol_lag_7	Volatility of volume traded in S&P500 across one month (7 months ago)	0.671
volume_vol_lag_5	Volatility of volume traded in S&P500 across one month (5 months ago)	0.620
<i>Portfolio factors</i>		
pct_intra_day_trades_lag_5	Percentage of intraday trades in a month (5 months ago), by counts	0.786
pct_val_options_lag_1	Percentage of portfolio that is options (1 month ago), by value	0.765
pct_val_options_lag_6	Percentage of portfolio that is options (6 months ago), by value	0.736
pct_intraday_val_lag_2	Percentage of intraday trades in a month (2 months ago), by value	0.716
inMarket	'1' if the portfolio has not been liquidated	0.708
pct_intraday_val_lag_1	Percentage of intraday trades in a month (1 month ago), by counts	0.699
pct_val_options_lag_3	Percentage of portfolio that is options (3 months ago), by value	0.687
pct_intraday_val	Percentage of intraday trades in this month, by value	0.664
pct_intra_day_trades	Percentage of intraday trades in this month, by counts	0.661
pct_val_options_lag_4	Percentage of portfolio that is options (4 months ago), by value	0.648

Exhibit 23: Most important variables in the logistic classifier.

Panic selling and freak-outs often have negative connotations. We show that this negativity may not always be warranted. Although panic selling in normal market conditions is indeed harmful to the median retail investor, freaking out in environments of sustained market decline prevents further losses and protects one's capital.

Panic sales are not random events. Specific types of investor, such as those with less than \$20,000 in portfolio value, tend to liquidate more frequently than others. Subtle patterns in portfolio history, past market movements, and demographic profile can be exploited by deep

neural networks to accurately predict whether an investor will panic sell in the near future.

Unfortunately, the problem of causation cannot be addressed with the data we have. Therefore, our study does not address *why* investors panic sell. This topic, however, would doubtless be an interesting direction for future research.

References

- Amini, S., B. Gebka, R. Hudson, and K. Keasey (2013). A review of the international literature on the short term predictability of stock prices conditional on large prior price changes: Microstructure, behavioral and risk related explanations. *International Review of Financial Analysis* 26, 1–17.
- Baker, M. and J. Wurgler (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance* 61(4), 1645–1680.
- Bali, T. G., K. O. Demirtas, and H. Levy (2009). Is there an intertemporal relation between downside risk and expected returns? *Journal of Financial and Quantitative Analysis* 44(4), 883–909.
- Barber, B. M. and T. 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, B. M. and T. Odean (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics* 116(1), 261–292.
- Barber, B. M., T. Odean, and N. Zhu (2009). Systematic noise. *Journal of Financial Markets* 12(4), 547–569.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics* 121(3), 823–866.
- Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review* 99(1), 243–64.
- Bauman, W. S., C. M. Conover, and R. E. Miller (1999). Investor overreaction in international stock markets. *Journal of Portfolio Management* 25, 102–111.
- Benos, A. V. (1998). Aggressiveness and survival of overconfident traders. *Journal of Financial Markets* 1(3-4), 353–383.
- Bollerslev, T. and V. Todorov (2011). Tails, fears, and risk premia. *The Journal of Finance* 66(6), 2165–2211.
- Brown, K. C., W. Harlow, and S. M. Tinic (1993). The risk and required return of common stock following major price innovations. *Journal of Financial and Quantitative Analysis* 28(1), 101–116.
- Chawla, N. V., K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer (2002). Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16, 321–357.
- Chopra, N., J. Lakonishok, and J. R. Ritter (1992). Measuring abnormal performance: do stocks overreact? *Journal of Financial Economics* 31(2), 235–268.
- De Bondt, W. F. and R. Thaler (1985). Does the stock market overreact? *The Journal of Finance* 40(3), 793–805.

- De Bondt, W. F. and R. H. Thaler (1987). Further evidence on investor overreaction and stock market seasonality. *The Journal of Finance* 42(3), 557–581.
- Doshi-Velez, F. and B. Kim (2017, Feb). Towards a rigorous science of interpretable machine learning. *ArXiv e-prints*.
- Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics* 127(2), 645–700.
- Kaminski, K. M. and A. W. Lo (2014). When do stop-loss rules stop losses? *Journal of Financial Markets* 18, 234–254.
- Kaul, G. and M. Nimalendran (1990). Price reversals: Bid-ask errors or market overreaction? *Journal of Financial Economics* 28(1-2), 67–93.
- Kelly, B. and H. Jiang (2014). Tail risk and asset prices. *The Review of Financial Studies* 27(10), 2841–2871.
- Kingma, D. P. and J. Ba (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klement, J. (2013). Assessing stop-loss and re-entry strategies. Available at SSRN: <https://ssrn.com/abstract=2277722> or <http://dx.doi.org/10.2139/ssrn.2277722>.
- Lo, A. W. and A. Remorov (2017). Stop-loss strategies with serial correlation, regime switching, and transaction costs. *Journal of Financial Markets* 34, 1–15.
- Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance* 53(5), 1775–1798.
- Park, J. (1995). A market microstructure explanation for predictable variations in stock returns following large price changes. *Journal of Financial and Quantitative Analysis* 30(2), 241–256.
- Pindyck, R. S. and N. Wang (2013). The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy* 5(4), 306–39.
- Puckett, A. and X. S. Yan (2008). Short-term institutional herding and its impact on stock prices. Unpublished manuscript, University of Missouri.
- Rotblot, C. (2004). *The Danger of Getting Out of Stocks During Bear Markets*.
- Rozeff, M. S. and M. A. Zaman (1998). Overreaction and insider trading: Evidence from growth and value portfolios. *The Journal of Finance* 53(2), 701–716.
- Shefrin, H. and M. Statman (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance* 40(3), 777–790.
- Shi, L., L. Han, Y. Wang, D. Chen, Y. Piao, and C. Gou (2011). Market crowd trading conditioning and its measurement. Available at SSRN: <https://ssrn.com/abstract=1661515> or <http://dx.doi.org/10.2139/ssrn.1661515>.

- Statman, M., S. Thorley, and K. Vorkink (2006). Investor overconfidence and trading volume. *The Review of Financial Studies* 19(4), 1531–1565.
- Wang, J., G. Meric, Z. Liu, and I. Meric (2013). Investor overreaction to technical insolvency and bankruptcy risks in the 2008 stock market crash. *The Journal of Investing* 22(2), 8–14.
- Weber, M. and C. F. Camerer (1998). The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization* 33(2), 167–184.

Appendix A. Supplementary Materials

Appendix A.1. Account security holding and portfolio allocations data

The raw position files consist of monthly snapshots that record the quantities and month-end prices of each security held in the portfolio of all accounts within the sample that were open on the last day of the month. Securities are uniquely identified either by CUSIP ID or ticker symbol, and accounts are uniquely identified by an anonymized numeric key (‘acid key’ or ‘acid’). An internal asset class assignment for each security is also provided within the brokerage account files, which classifies each CUSIP/ticker as one of the following: equities, mutual funds, fixed income securities, cash or cash equivalents, or options. In addition, a separate identifier is provided distinguishing cash equities from ETFs within the equities category. Exhibit A2 outlines the data fields in the position files.

Key	Description	Format
month	Month of snapshot (all positions are those held at month-end)	YYYYMM (e.g. 201512)
settle_qty	Quantity of shares held in security	double
ticker_symbol	Ticker symbol	string (e.g. AAPL)
cusip_num	Security identification number registered with the US SEC	9-digit alpha-numeric (e.g. 17275R102)
issue_price	Exchange-listed close price on the last market day of the assigned month	double
product_grplv1	Top level security type identifier	string (e.g. ‘EQUITY’)
product_grplv2	Mid level security type identifier	string (e.g. ‘EQUITY’)
product_grplv3	Bottom level security type identifier	string (e.g. ‘EQUITY’)
acid_key	Unique account identification number	integer (e.g. 9374629673)

Exhibit A2: Summary of the data fields in the positions datafile.

Appendix A.2. Trading data

The raw trade files consist of annual records of all trades executed by the sampled accounts during the year. Each trade is timestamped by date, uniquely identified by acid and CUSIP/ticker, and includes the dollar principal (either positive or negative) expended on the trade (a buy or sell, respectively). The commission in dollars paid by the account for the trade is also recorded. The daily timestamped nature of the trading data is critical to our analysis because it enables computation of metrics based on intra-month trading decisions and returns and therefore exposes granular patterns of behavior that would not be visible at fixed-interval monthly or quarterly frequencies. Furthermore, whereas portfolio holdings data reflect both individual allocation decisions and changes in asset values, making it difficult to disentangle the effects of investor decisionmaking from the effects of changing prices, the trade data reflect the decision to allocate in a much more direct manner. For these reasons, the availability of trade data differentiates this paper from similar studies focusing on retail brokerage account or government stock holdings data. Exhibit ?? provides a summary of the data fields in the trade files.

Key	Description	Format
trade_date	Date of trade	YYYYMMDD (e.g. 20080317)
buy_sell	Indicator of buy or sell	string (e.g. 'B', 'S')
principal	Principal amount traded	double
quantity	Units of asset traded	integer
tcommission	Trade commission	double
cusip_nr	Security identification number registered with the US SEC	9-digit alpha-numeric (e.g. 17275R102)
ticker_symbol	Ticker symbol	string (e.g. AAPL)
product_grplv1	Top level security type identifier	string (e.g. 'EQUITY')
product_grplv2	Mid level security type identifier	string (e.g. 'EQUITY')
product_grplv3	Bottom level security type identifier	string (e.g. 'EQUITY')
acid_key	Unique account identification number	integer (e.g. 9374629673)

Exhibit A3: Summary of the data fields in the trades datafile.

Appendix A.3. Relationship between household, accounts, and customers

An investing account can be co-owned by multiple customers. The brokerage firm has associated a group of accounts into a household based on the relationships between the customers. An investing account can only belong to one household, whereas the map between investing accounts and customers can be many-to-many. Exhibit A4 illustrates the relationship between the accounts and customers for one of the households.

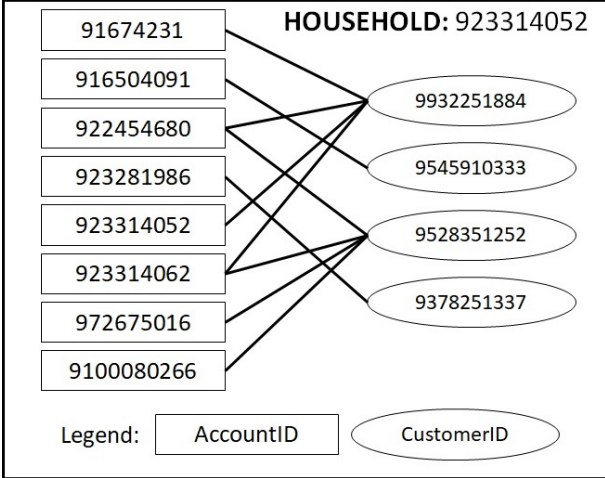


Exhibit A4: A graphical representation of how the households, accounts and customers are related.

Appendix A.4. Demographic data

The demographic files record the personal information on the account application forms of the accounts selected by the random sampling procedure and can be merged with the historical account data contained in the position and trade files using the anonymized key. Demographic fields include age, income, profession, investment knowledge ('knowledge'), investment experience ('experience'), and marital status. Knowledge and experience are

survey questions included with the other components of the application questionnaire, and can receive values of ‘Excellent,’ ‘Good,’ ‘Limited,’ ‘None,’ or ‘Decline to report.’ These fields reflect the account holder’s self-reported view of his or her familiarity with personal finance and financial decision-making, and therefore offer a novel way to measure the behavior and performance of investors as a function of their financial sophistication. Exhibit A5 provides a summary of the data fields in the demographics files.

Key	Description	Format
month_last_record	Month of last record	integer (YYYYMM)
cust_age	Age	integer
dmsa_martrl_stat_cd	Marital status	string
cust_depndt_qy	Number of dependents	integer
ps_gndr_cd	Gender	char
dmsa_curr_occup_tx	Occupation group	string
invst_knldg_cd	Investment knowledge	string
invst_exprc_cd	Investment experience	string
acctid_key	Account ID	integer

Exhibit A5: Summary of the data fields in the demographic attributes datafile.

Appendix A.5. Computing the demographic distribution

The computation of the distribution of demographic features in our dataset is complicated by the fact that a household can consist of multiple customers. Furthermore, some customers in a household can be associated with more accounts than others. In Exhibit A4, customer 9932251884 is associated with four accounts, whereas customer 9378251337 is only associated with one account. One can conceptualize that the former customer is more influential and should be assigned a higher weight.

There are many ways to aggregate this information. The typical method used by the brokerage firm is to consider either the minimum or the maximum of all customers for a single variable. For example, it will use the maximum age of all customers in a household when analyzing the age of a household, or consider the highest level of investing experience declared by all customers in a household to be the household’s investing experience. While this is useful for marketing purposes, where one is only interested in finding a target audience (e.g., if the household has someone who needs retirement products), it does not suffice for our study. Furthermore, this method will fail when one attempts to apply it to unordered information, such as occupational groups.

We choose to take into account the portfolio weights of each customer to analyze the demographic distribution of our dataset. To do so, we will first compute the weight of every account based on its average portfolio value over its lifetime. Let the portfolio weight of account i be p_i . For every account, we assume an equal weight among all its registered customers. Denote the set of customers in account i by c_i . Thus, the effective weight of customer j will then be $\sum_{i,j \in c_i} \frac{p_i}{|c_i|}$.

We demonstrate an example of the computational process using Exhibit A6. There, we have three investing accounts (in rectangles) and three customers (in ovals). First, we compute the average portfolio values across the entire time horizon to find that the portfolio weights of the accounts are $\frac{1}{3}$, $\frac{1}{2}$, and $\frac{1}{6}$, respectively. For each account, we then assign a weight from the account to the customers on an equal basis. The results of this step are in green. Finally, for each customer, we can compute the total weight. For customer 9932251884, the overall weight is then $\frac{1}{2} \times 0.5 + \frac{1}{6} \times 1 = 0.417$.




Exhibit A6: An example of how the demographic weights are computed. The rectangles represent investing accounts, and the ovals represent customers. The numbers in red are the portfolio weights for each investing account, and the green numbers are the weights to a customer from an investing account.

We also attempted a method in which all customers in a household were assigned equal weights. Although the resulting numbers differed slightly, the conclusions drawn were similar.

Appendix A.6. Changing the parameters for the identification of panic sales

Our method of determining a panic sale requires us to define two parameters, p_1 and p_2 , the monthly portfolio decline and the monthly portfolio net sell, which we set to 0.9 and 0.5, respectively. We conduct additional runs with different parameter pairs to determine how they affect the identification of panic selling. Because the amount of computation required is immense, however, costing more than 5,000 CPU-hours per run, we performed only two additional runs with the parameter settings shown in Exhibit A7. We did not vary p_3 and p_4 , the portfolio rebound and the cumulative net buy, because they do not affect the identification of panic sales.

Run	p_1	p_2	p_3	p_4
1	0.9	0.5	0.5	0.5
2	0.5	0.25	0.5	0.5
3	0.25	0.1	0.5	0.5

Exhibit A7: Summary of the parameters used in the various run

Exhibit A8 shows the results of our additional runs versus our baseline. As expected, decreasing the thresholds will increase the number of panic sales being captured. Although we still observe that the major spikes in reaction to major events remain across all runs, lowering the thresholds also amplifies noise in our data.




Exhibit A8: The number of panic sales over time for different parameter sweeps. The red, blue, and gold bars represent the results for the parameter sets $\{0.9, 0.5, 0.5, 0.5\}$, $\{0.5, 0.25, 0.5, 0.5\}$, and $\{0.25, 0.1, 0.5, 0.5\}$ respectively.

Appendix A.7. Explanation of machine learning models

Appendix A.7.1. Issue of imbalanced data

One of the biggest issues encountered in training our machine learning models is the extremely imbalanced dataset. Given that the negative class comprises of 99.87% of all data points, a naive classifier that always predicts zero will easily achieve an accuracy of 99.87%. Naively training the models based on the usual cross-entropy minimization will lead to this outcome.

To mitigate this problem, we oversampled the underrepresented class, which we achieved by creating training batches with equal weights. We also considered using SMOTE (Chawla et al., 2002), but we found that interpolating variables generated nonsensical data points; our data were constructed in such a way that too many constraints have to be fulfilled for this method to be applicable.

Appendix A.7.2. Metrics for evaluating models

As discussed, accuracy over the entire test set is not a valid measure for imbalanced data. Instead, we evaluated our models on the accuracy of both positive-labeled and negative-labeled data points to get a more useful idea of their real-world performances.

In addition, we reported two other metrics that characterize the performance of machine learning models: the area under the AUROC and the AUPRC.

We define the following measures:

$$\text{Sensitivity} = \text{Recall} = \frac{\#\text{True Positives}}{\#\text{True Positives} + \#\text{False Negatives}} \quad (\text{A.1})$$

$$\text{Specificity} = \frac{\#\text{True Negatives}}{\#\text{True Negatives} + \#\text{False Positives}} \quad (\text{A.2})$$

$$\text{Precision} = \frac{\#\text{True Positives}}{\#\text{True Positives} + \#\text{False Positives}} \quad (\text{A.3})$$

The ROC curve is created by plotting the true positive rate of a classifier, also known as its ‘sensitivity’ or ‘recall,’ against its false positive rate, or 1–specificity, at different thresholds. A naive classifier will have an ROC profile that is a diagonal from (0,0) to (1,1). In this case, the AUROC of the naive classifier will be 0.5. On the other hand, a perfect classifier will have an AUROC of 1. Mathematically, the AUROC is the probability that the score of a randomly selected positive example is higher than the score of a randomly selected negative example.

The ROC is not useful if one is interested in the rate that the models produce false alarms. In such cases, the PR curve is more useful. A naive classifier will have a precision that is equal to the proportion of positive data points in the entire sample for all thresholds.