A Tale of Two Zoos:

Machine Learning Insights on Fifteen Million Retail Investors

Pulak Ghosh*, Huahao Lu†, Hong Zhang‡, Jian Zhang§

This Draft: May 30, 2026

Abstract

Both stock characteristics (the "factor zoo") and behavioral biases (the "bias zoo") impact retail investors' returns. But which factors/biases are more important? We address this question by utilizing machine-learning tools to analyze 15 million retail investor accounts in India. We observe that Neural Networks outperform other algorithms in uniquely predicting both good and bad out-of-sample performance, with (under)diversification, portfolio turnover, and momentum being the leading factors to influence investors' total returns. Neural Network-identified smart investors also help improve market efficiency. Our results provide insights into the role of retail investors and the plausible asset-pricing framework that explains their returns.

Keywords: Machine Learning, Anomalies, Behavioral Biases, Retail Investors, Market Efficiency
JEL Codes: D90, G40,

* Department of Decision Sciences and Centre of Public Policy, Indian Institute of Management Bangalore, Email: pulak.ghosh@iimb.ac.in
† PBC School of Finance, Tsinghua University, Email: luhh.21@pbcsf.tsinghua.edu.cn
‡ Lee Kong Chian School of Business, Singapore Management University, 50 Stamford Road, Singapore 178899; E-mail: hongzhang@smu.edu.sg.
§ The University of Hong Kong, Faculty of Business and Economics, Email: zhangj1@hku.hk

I. Introduction

Over the last few decades, the academic community has demonstrated that trading strategies based on a multitude of firm characteristics can yield striking returns (i.e., the "factor zoo"; see McLean and Pontiff 2016; Harvey, Liu, and Zhu 2016; Hou, Xue, and Zhang 2020; Jensen, Kelly, and Pedersen, 2023 for recent evidence). A related but distinct literature investigates the psychological heuristics and biases to which retail investors are susceptible (i.e., the "bias zoo"; see Barber and Odean 2013; Hirshleifer 2015; and Barberis 2018 for literature reviews). An important implication from these two streams of literature is that, when retail investors participate in stock markets, both firm characteristics and psychological heuristics—which we refer to as characteristics-based and behavioral "anomalies" for brevity—can affect their portfolio returns. For instance, an investor may unintentionally hold stocks with characteristics that predict strong returns in some periods, allowing her portfolio to perform well. In other periods, she might hold stocks with unfavorable characteristics or make behavioral mistakes that erode the value of her portfolio.

However, both strands of literature face a similar challenge, which Fama (1998) describes as the "lack of discipline" in proposed behavioral biases and Cochrane (2011) characterizes as a "multidimensional challenge" stemming from the excessive characteristics that exhibit return predictability. To the extent that both types of anomalies influence retail investors, this issue could become even more pronounced when examining their returns. Important questions arise as a result: Which type of anomalies contributes more to retail investors' returns? Moreover, do all reported anomalies affect retail returns or should we expect a much shorter list? Addressing these questions may provide valuable insights into a more parsimonious framework for tackling the multidimensional challenge. Meanwhile, retail investors worldwide rely on stock markets to build wealth, save for retirement, and achieve various financial goals. Identifying their most significant biases and misaligned exposures to characteristics thus also carries important normative implications for researchers and policymakers.

Our paper aims to shed light on these questions by employing various machine learning tools to a unique and large proprietary account-level dataset containing the daily trading activities of *all* retail investors on the National Stock Exchange of India (NSE). As the most populous country in the world, India provides an ideal testing ground to understand retail investors, with the NSE being

its largest stock exchange in India and the seventh largest stock market worldwide by Dec 2023.[1] From this dataset, we identify 15.4 million valid retail accounts and 1.523 billion investor-month return observations for the testing period of 2012-2020. For our baseline analysis, we construct 23 main holding-weighted *stock characteristics* and 13 leading proxies of *behavioral biases* for each account (including more characteristics will not change our main conclusions).

To investigate how these characteristics and biases affect retail performance, we employ a list of models, ranging from the traditional OLS to machine-learning models such as LASSO, Ridge, Random Forest, the *Feedforward Neural Network* (*FNN*) and an enhanced *Residual Neural Network* (*ResNN*).[2] Our primary objective is twofold. First, we investigate whether any of these models can help predict retail investors' monthly portfolio returns based on the multitude of behavioral biases and stock characteristics.[3] Since retail investors often make behavioral mistakes, it may not be surprising for statistical models to detect retail investors who consistently exhibit poor performance. In contrast, superior performance is difficult to predict even for institutional investors (e.g., Carhart 1997). Hence, we are particularly interested in whether some models can also identify retail investors with good performance. Second, if a model demonstrates reliable out-of-sample predictive power for both positive and negative returns, we use it to assess the relative importance of behavioral biases vis-à-vis firm characteristics as well as that of individual anomalies. In other words, we want to employ the most effective model to shed light on a more parsimonious framework of factors that impact retail investors.

Our return prediction analysis is employed as follows. In line with the literature (e.g., Kaniel, Lin, Pelger, and Van Nieuwerburgh 2023), we divide our return prediction period from 2012 to 2020 into three equal-length subperiods. We then train a model on two subsets of the data and use the trained model to predict returns on the remaining subset. This approach ensures that we can

[1] https://www.cnbc.com/2023/12/12/india-overtakes-hong-kong-to-become-worlds-seventh-largest-stock-market.html

[2] Of the two Neural Network models, FNN is more traditional, while ResNN reflects more recent development. Its key feature, "residual connections", is widely adopted in recent architectures such as BERT and ChatGPT. Later sections will delve into further details, demonstrating ResNN's superior capability for our purposes.

[3] We use ranks to normalize the distribution of all imputs so their importance can be more easily inferred. This method is widely used for machine learning models (e.g., Kelly, Pruitt, and Su, 2019; Freyberger, Neuhierl, and Weber, 2020). To calculate an investor's total investment returns, we calculate the daily return generated by her existing portfolios at the beginning of a given date and then compound her daily returns into monthly returns.

test the out-of-sample predictive power of the model over the entire prediction period.[4] During the out-of-sample predicting period, we categorize retail investors into five quintiles according to the predicted returns of a particular model; all quintiles are rebalanced monthly. The *High* and *Low* groups comprise the top and bottom 20% of predicted winners and losers among investors, respectively. We then calculate the out-of-sample returns of the high and low groups, along with their return difference. Since small stocks are difficult to trade in emerging markets (Liu et al. 2019) and may overstate the predicting power of machine learning models (Avramov, Cheng, and Metzker 2023; Cong et al., 2021), we exclude 30% of small stocks and value weight investors' stocks and portfolios in estimating the performance of investor quintiles. We further use the local three-factor (Fama and French 1992, 1993) or four-factor models (Carhart, 1997) to adjust these investor portfolio returns.

We observe that the two Neural Network models outperform other models in predicting total returns for retail investors. Both FNN and ResNN identify investors who can generate significantly positive out-of-sample returns. Despite being retail investors, the top 20% of predicted winners can generate a monthly return of 1.3% and 1.7%. The economic magnitude remains approximately the same when risk-adjusted (e.g., 1.4% and 1.5% adjusted by four factors). In contrast, all other models fail to predict winners. Given how difficult it is to predict four-factor adjusted mutual fund returns, the superior performance of retail investors strongly suggests that Neural Networks capture crucial characteristics of retail investors.

On the loser side, although many models can select *Low* groups that deliver significant or marginally significant returns, the two neural network models still perform the best. FNN and ResNN-predicted *Low* groups deliver a significant negative monthly return of -2.0% and -2.8%, respectively, allowing their *High* groups to outperform the *Low* groups by as much as 3.3% and 4.5% per month. Collectively, the two neural network models (particularly *ResNN*) outperform others in identifying winners and losers among retail investors.

Since neural networks exhibit superior predictive capabilities for investor returns, we next utilize them to explore the relative importance of behavioral biases and firm characteristics. To

[4] Our sample starts from 2010. We use the first two years of information to calculate the initial values of stock characteristics and behavioral biases. Hence, our return predicting test starts in 2012. We adopt the Kaniel et al., (2023) approach because it has the advantage of testing the model on every sample in the dataset, enhancing the robustness of model comparisons by mitigating the influence of specific periods. Additionally, within the training data, we randomly set aside 30% of the samples for validation purposes.

achieve this goal, we follow the literature (e.g., Horel and Giesecke 2020; Sadhwani et al. 2020; Kaniel et al. 2023) and employ FNN-based variable gradient analyses to estimate the relative importance of each variable in predicting returns. We then sum up the importance of all predictors falling into the two categories and scale our estimation so that the total importance of the two groups equals 100%. Behavioral biases and firm characteristics exert approximately 63.5% and 36.5% relative importance, respectively, confirming a higher influence of behavioral biases despite the smaller number of predictors in this category.[5]

Variable gradient analysis further identifies diversification, portfolio turnover, and momentum as the three leading anomalies that influence overall retail returns. The first anomaly arises because many retail investors fail to recognize and thus benefit from diversification (e.g., Barber and Odean 2000; Benartzi and Thaler 2001; Lusardi and Mitchell 2011). Next, overconfidence often causes investors to trade too aggressively, allowing their high portfolio turnover to reduce their welfare (Odean 1998; Barber and Odean 2000). The third denotes perhaps the most famous anomalies in the literature.[6] Hence, our results prompt under-diversification and overtrade as two leading behavioral mistakes and momentum as the leading characteristic to impact retail investors' returns.

We further observe that the top 3, 5, and 10 anomalies jointly span approximately 29%, 38%, and 53% of the total importance of all (36) anomalies, suggesting that the relevance of anomalies is quite skewed, with the top ones exhibiting disproportionally higher explanatory power. On the one hand, explaining retail investors' total returns also seems to require a higher dimensionality than implied by common factor models (e.g., Fama and French 1993, 2015; Carhart, 1997).

Next, we notice that investors' total returns can stem from two distinct sources: holding an existing portfolio for a specified period, such as a month (hereafter, *holding returns* when there is no confusion), and initiating new trades to buy and sell stocks during the month (hereafter, *trading*

[5] We also observe that when firm characteristics are used alone, FNN-predicted *High* and *Low* groups fail to deliver significantly positive or negative out-of-sample returns. Nor can FNN predict a significant *High*-minus-*Low* return spread. In contrast, using behavioral biases alone can predict a significant *High*-minus-*Low* return spread. This difference provides additional evidence leaning toward the relative importance of behavioral biases in affecting returns.
[6] Momentum can be related to investors' biases, such as the disposition effect (e.g., Grinblatt and Han, 2005). Since our goal is to identify the direct impact of characteristics and biases in a horse race, we include these two (and any other) effects as separate predictors. The behavioral category includes two of the top three factors affecting investment returns, aligning with its relatively higher total explanatory power, as observed above.

returns).[7] Behavioral theories propose that the motivations for new trading may differ from those for continuation. For instance, the disposition effect (e.g., Shefrin and Statman, 1985) suggests that unrealized capital gains motivate investors to trade (i.e., sell), whereas loss aversion incentivizes investors to retain losing assets, thus influencing holding returns. Another example is the salience theory (Bordalo, Gennaioli, and Shleifer 2012, 2013, 2020), which suggests that salient information, such as extreme stock prices, may also attract investors' attention to initiate new trades. Hence, our next question becomes how anomalies affect these two sources of returns.

To answer this question, we adjust our objective to predict trading or holding returns using Neural Networks. Again, we first validate the predicting power by observing the predicted *High* (*Low*) group to generate a significantly positive *High*-minus-*Low* total return spread. We then employ the variable gradient analysis to assess the relative importance of each variable in predicting each source of returns. We observe that behavioral anomalies play an even more striking role in predicting trading returns, whereas the relative importance of firm characteristics increases in explaining holding returns. Indeed, behavioral biases predominantly account for 95% of the predicting power for trading returns, compared to 52.2% when predicting holding returns. These observations are reasonable because characteristics-related returns likely contribute more to less rebalanced portfolios, whereas new trading is often initiated by behavioral reasons.

Portfolio turnover, the disposition effect, and diversification emerge as the three leading anomalies in predicting trading returns. In other words, the disposition effect replaces momentum, rendering all the top 3 trading return predictors being behavioral.[8] They jointly span an astonishing 87% of the total importance of all anomalies, suggesting that a highly parsimonious set of behavioral biases plays a dominant role in explaining the short-term performance of newly initiated trades, which differs from the case of total returns. Interestingly, even though the importance of

[7] For any given month, we calculate holding returns as the portfolio returns generated by the beginning-of-the-month holdings during the month. We then calculate trading returns as the cumulative daily returns generated by the newly initiated trades during the month (i.e., we compound these daily returns until the end of the month).

[8] Since Shefrin and Statman (1985), the development of this literature has been extensive, though the causes and consequences of the disposition effect are still under debate (see, among others, Grinblatt and Han, 2005; Barberis and Xiong, 2009, 2012; Calvet, Campbell, and Sodini, 2009; Ivkovic and Weisbenner, 2009; Kaustia, 2010; Ben-David and Hirshleifer, 2012; Henderson, 2012; Li and Yang, 2013; Frydman et al., 2014; An, 2016; Chang, Solomon, and Westerfield, 2016; Fischbacher, Hoffmann, and Schudy, 2017; Frydman and Wang, 2020). DellaVigna (2009; 2018), Hirshleifer (2015), and Barberis (2018) provide recent surveys.

behavioral biases drops in predicting holding-based returns, (under)diversification, portfolio turnover, and momentum remain the top three factors.

One notable feature of our analysis is the ability of the two neural network models to identify retail investors who can deliver superior performance, suggesting that not all retail investors are uninformed and biased. Indeed, an emerging body of research recognizes that at least some retail investors may trade smartly (Kaniel et al., 2012; Kelley and Tetlock, 2017; Boehmer et al., 2021; Coval, Hirshleifer, and Shumway, 2021). To further validate this conjecture, we investigate how trading activities of the *High* and *Low* groups of retail investors (hereafter, winners and losers) impact market efficiency. To this end, we follow Brogaard et al. (2021) to decompose stock return variance and focus on components related to noise and firm-specific private information.

Across all models, the trading intensity of losers—computed as their total trading volume in a stock scaled by the stock's overall trading volume—significantly amplifies noise. The two neural network models further indicate that the trading intensity of winners enhances firm-specific private information while reducing noise. Additional tests on variance ratios (Lo and MacKinlay, 1988; Campbell, Lo, and MacKinlay, 1997) and market delay (Hou and Moskowitz, 2005; Griffin, Kelly, and Nardari, 2010) confirm the role of these winners in enhancing market efficiency. Since institutional investors are often less informed in emerging markets (Chui, Subrahmanyam, and Titman, 2022), it is plausible that some Indian retail investors may capitalize on this gap by gathering information and disseminating it through trading. Our findings thus underscore both the contribution of these savvy retail investors to market efficiency and the ability of machine learning tools, in particular neural networks, to identify these key players.

We finally conduct a battery of additional analyses to shed further light on the economics and robustness of our main findings. We first examine whether market conditions, such as volatility, may affect the performance difference between the top and bottom quintiles of investors. Since behavioral biases are the main driving force of the difference, a significant impact of market conditions could suggest an intricate interaction between market conditions and behavioral biases that warrants further examination. However, we observe insignificant coefficients when we link the High-minus-Low spread to the return, volatility, and skewness of the local market.

Similarly, we also ask whether the aggregate sentiment of retail investors (proxied by the aggregate order imbalance of all investors in our sample) could impact the return difference

between top and bottom quintile investors. Although sentiment negatively affects the performance difference based on FNN, this effect becomes insignificant when the more advanced ResNN is used. As such, the relative performance of retail investors as well as the underlying behavioral and characteristics-based anomalies could be largely independent of market conditions or sentiment.[9]

Next, we investigate the persistence of performance and related welfare implications. Our previous analysis focuses on one-month-ahead performance. If top-quintile investors outperform bottom-quintile ones only in the next month but subsequently underperform, one might be wary about the implications of our baseline results. To address this concern, we use neural networks to predict the total returns of retail investors over a horizon of up to twelve months. Our analysis reveals that, after portfolio formation, the high-minus-low spread remains significant for two to five months (albeit with declining magnitude), depending on the factor model applied. Although the spread becomes insignificant in later months, it does not reverse. These results suggest that neural networks are effective in identifying anomalies with important welfare implications when retail investors engage in the stock market.

Lastly, we assess the robustness of the predicting power of neural networks to an alternative list of predictors, alternative thresholds of small stocks that we exclude, and when compared to alternative models. We first follow the cauterizations of Jensen, Kelly, and Pedersen (2023) to expand accounting variables to 67 firm characteristics as our firm-side predictors. Even in this case, neural networks still exhibit a similar power in predicting out-of-sample performance, and behavioral biases still hold similar relative importance. Next, our previous analyses exclude 30% of small stocks. We show that the predictive power of both FNN and ResNN remains robust to different removal thresholds (e.g., 20% or 40%). Finally, we present more formal statistics demonstrating that neural networks significantly outperform other models in predicting the returns of retail investors. Collectively, these observations suggest that neural networks provide a powerful and robust tool for uncovering the economic basis of retail investors.

Our results are related to several strands of literature. A growing body of literature demonstrates that machine-learning models can help predict asset prices across various market

[9] We have also verified that, although we divide our sample period into three subperiods when training models, the relative performance is highly robust to different subperiods.

sectors, ranging from equity to option pricing in the US and global markets.[10] Karolyi and Van Nieuwerburgh (2020) and Kelly and Xiu (2023) provide recent reviews. Our analysis is closely related to recent studies applying machine-learning models to predict the performance of institutional investors, such as mutual funds (e.g., Li and Rossi, 2020; DeMiguel, Gil-Bazo, Nogales, and Santos, 2023; Kaniel, Lin, Pelger, and Van Nieuwerburgh 2023) and hedge funds (Wu, Chen, Yang, and Tindall, 2021). We contribute by using a battery of machine-learning tools to scrutinize the performance of a large sample of retail investors. This extension is important, as the economic rationale guiding retail investors' investments can differ from that of institutional investors. Indeed, we observe that behavioral biases typically dominate in retail decisions.

In a closely related paper, Balasubramaniam, Campbell, Ramadorai, and Ranish (2023) use a large sample of Indian retail accounts to shed light on investor attributes that can give rise to investor clientele effects for stock characteristics. Unlike their focus on investor holdings, we aim to identify the most important investor bias and stock characteristics that can directly impact the returns and welfare of retail investors. To achieve this goal, we adopt a two-step methodology of first exploring a list of machine learning tools and then using the most reliable ones to systematically assess the return impact of individual anomalies. This goal and methodology also differ from the vast existing studies using account-level data to document how particular forms of behavioral biases influence retail investors' decisions (Barber and Odean 2013; Hirshleifer 2015; and Barberis 2018 provide comprehensive reviews).

In doing so, we also contribute to the literature on behavioral biases and that on asset pricing anomalies. One important goal of the former is to use psychological insights to explain many anomalies in individuals' financial decision-making. Although this effort provides profound insight into how individual investors make decisions, the multitude of proposed behavioral biases gives rise to a "lack of discipline" concern (Fama, 1998). To address this issue, a few recent papers use survey-based methods to nail down the relative importance of behavioral biases (e.g., Choi

[10] See, among others, Freyberger, Neuhierl, and Weber (2020), Gu, Kelly, and Xiu (2020), Bryzgalova, Pelger, and Zhu (2020), and Chen, Pelger, and Zhu (2023) for stock returns and characteristics, Jensen, et al. (2022) for trading-cost-adjusted portfolio optimization, Leippold, Wang, and Zhou (2022) for the Chinese equity market, Li et al (2023) on the spillover effect of the global supply chain, Bianchi, Büchner, and Tamoni (2021) for bond risk premium, Easley, López de Prado, O'Hara, and Zhang (2021) for market microstructure, Filippou et al. (2023) for currencies, Bali, Beckmeyer, Mörke, and Weigert (2023) for option pricing, and Van Binsbergen, Han, and Lopez-Lira, A. (2023) for the conditional biases in earnings expectations. Avramov, Cheng, and Metzker (2023) report that the predicted returns could drop substantially in magnitude when small firms are excluded.

and Robertson, 2020; Liu, Peng, Xiong, and Xiong, 2022). A similar "multidimensional challenge" (Cochrane 2011) exists in characteristics-based anomalies with excessive return predictability. We are related to recent studies using machine learning tools to address this issue (e.g., Giglio, Liao, and Xiu, 2021; Feng, Giglio, and Xiu, 2020; Lopez-Lira and Roussanov 2020). Our novelty is to use retail accounts to synchronize both psychology- and characteristics-based anomalies and shed light on a more parsimonious conceptual framework of asset pricing and investor behavior.

Our results also cast doubts on the conventional wisdom of retail investors acting as noise traders to weaken market efficiency (De Long et al., 1989; Shleifer and Summers, 1990; Shiller, 2003). Growing evidence suggests that institutional investors, who are often assumed rational in traditional financial theories, may exhibit biases (Frazzini 2006; Cici 2012; Daniel et al., 1998; Baker et al., 2010; Coval and Moskowitz 1999; Kacperczyk et al., 2016) and that some retail investors may also appear to trade smartly (Kaniel et al., 2012; Kelley and Tetlock 2017; Boehmer et al., 2021). We contribute by demonstrating that a proper tool, such as neural networks, can help identify smart retail investors who contribute to market efficiency.

Our analysis suggests that Residual Neural Networks (*ResNN*) could provide a helpful tool to financial analysis. A widely acknowledged challenge in deep learning is that deeper neural networks are more difficult to train due to issues such as vanishing gradients. *ResNN* addresses this difficulty by reformulating the output of a particular layer as a learning residual function plus the layer's input (He et al., 2015).[11] The key feature of *ResNN*—"residual connections" or the addition of the original input to the output of a deeper layer within a neural network—allows for deeper training and thus is widely used in Transformer models such as BERT and ChatGPT. Since our original inputs are characteristics, this approach resambles the use of characteristics adjustments in estimation. Our results suggest that *ResNN* could provide a powerful tool for financial tasks when characteristics- or benchmark-adjustments are desirable.

The remaining article is organized as follows. Section II describes the data and machine learning models. Section III provides baseline tests for predicting retail investors' returns. Section IV examines the importance of behavioral heuristics and firm characteristics. Section V provides additional tests and robustness checks, followed by a short conclusion with policy implications.

[11] Residual Neural Networks were originally developed to improve image recognition and won the *ImageNet* 2015 competition. On Mar 19, 2025, the seminar work of He et al., 2015 has garnered more than 260,564 Google citations.

II. Data, Main Variables, and Machine Learning Models

This section describes the data and explains how we construct our main variables. We then briefly describe the machine learning models used in our later analysis.

A. Data

We collected data from multiple sources. To characterize the impact on investors' trading behavior, we obtain a comprehensive database of all trading records on the NSE of India for the period 2010-2020. The NSE is the leading exchange in India and the world's 9th-largest stock exchange as of May 2021.[12] For each transaction, we can observe the anonymized permanent account number (PAN) of the individual[13], the transaction date, the ticker of the security, the number of shares purchased or sold, and the execution price. We require all transactions to be associated with stocks included in the Prowess Database (similar to CRSP in the US) maintained by the Centre for Monitoring Indian Economy (CMIE). Additionally, we retain only securities that are common shares of domestic stocks and exclude trading activities related to ETFs and foreign stocks.

The initial sample consists of the entire sample of 19.36 million retail accounts at NSE. For each retail investor, we further obtain sociodemographic data including gender, age, and, most importantly, geographic identifier (i.e., India PIN code), which allows us to identify the district of residence for each investor. We exclude accounts that have a negative balance, as such accounts could incur missing information or short selling. Our final sample includes 15.418 million valid individual accounts and approximately 1.52 billion investor-month portfolio-return observations.

We obtain stock returns and characteristics from the CMIE Prowess database maintained by CMIE, Center for Monitoring the India Economy. Previous studies on Indian firms have utilized this dataset, including works by Bertrand, Mehta, and Mullainathan (2002), Gopalan, Nanda, and Seru (2007), Lilienfeld-Toal, Mookherjee, and Visaria (2012) and Gopalan, Mukherjee, and Singh (2016). The detailed firm characteristics are summarized in Table 1. In addition, we employ the Fama-French three-factor model (Fama and French 1992, 1993) and Carhart's four-factor

[12] https://www.world-exchanges.org/our-work/statistics

[13] The PAN is a unique identifier issued to all taxpayers by the Income Tax Department of India. The trading data are at the individual level so that it is not a concern if a given individual investor may hold multiple accounts.

model (Carhart, 1997) to adjust investor returns. The data for these local factors are downloaded from Global Factor Data (Jensen, Kelly, and Pedersen, 2023).[14]

B. Main Variables

Our objective is to use machine-learning tools to predict investors' total returns. To construct the time series for an investor's total investment returns, we calculate the daily return generated by her existing portfolios at the beginning of a given date and then compound her daily returns into monthly returns. This approach is in line with the literature (e.g., Odean, 1998; Barber and Odean, 2000), which also allows us to further decompose the investor's monthly total returns into two sources: the part generated by the holding at the beginning of the month (i.e., holding-based returns) and the part generated by the newly initiated trading during the month (i.e., new trading-based returns). Empirically, the new trading-based return of a month is calculated as the difference between the monthly total return and the holding-based return. As we will see shortly, behavioral biases and firm characteristics play different roles in affecting the two sources of returns.

In constructing portfolio returns, we follow the literature on international stock returns—e.g., Liu et al.'s (2019) analysis of the Chinese stock market—and exclude 30% of small stocks. These small stocks are not only difficult to trade in emerging markets (Liu et al. 2019) but also may overstate the predicting power of machine learning models (Avramov, Cheng, and Metzker 2023; Cong et al., 2021). Due to the presence of some extreme values in the distribution of investors' monthly returns, we applied a winsorizing procedure at the 1st and 99th percentiles to mitigate the impact of outliers. Later sections will show that our results are robust to these data screening processes.

Behavioral Preditors: We resort to the recent behavioral and asset pricing literature to construct the list of predators. This data enables us to construct 13 investor characteristics at the monthly frequency to capture behavioral heuristics.

The Disposition Effect: Many studies have demonstrated the behavioral bias of investors to sell stocks that have gained profits while choosing to continue holding stocks that have incurred losses

[14] We thank the authors for maintaining a comprehensive global factor dataset and making is easily accessible on: https://jkpfactors.com/

(Shefrin and Statman, 1985; Odean, 1998). Specifically, we estimate the disposition effect of each investor through the following specifications:

$$Sell_{i,j,t} = \alpha + \beta_i Gain_{i,j,t-1} + \epsilon_{i,j,t},$$

where $Sell_{i,j,t}$ is a dummy variable that equals one if the investor i sells stock j at date t (including partial sales) and zero otherwise. $Gain_{i,j,t-1}$is also a dummy variable that equals one if the stock exhibits a positive unrealized gain at the time $t-1$ and zero otherwise. The gain for a particular stock is calculated based on the difference between the $t-1$ price and the average purchasing price using a first-in-first-out methodology. The disposition effect of the investor in any given month is the the parameter β_i estimated from a 500-trading-day rolling window (i.e., approximately two years) right before the beginning of the month.

Diversification: We measure each investor's degree of diversification based on the number of stocks held in their portfolio. Specifically, we calculate the daily count of stocks in the investor's portfolio and subsequently take the monthly average.

Turnover: We employ investor turnover as a proxy for their trading activity. Prior research has consistently shown that increased trading frequency is often associated with inferior performance (Odean 1998; Barber and Odean 2000). We calculate the daily turnover as the ratio of the trading amount to the total value of the investor's portfolio, followed by monthly averaging.

Local Bias: Investors often prefer companies located in close geographic proximity (e.g., Ivkovic and Weisbenner, 2005; Massa and Simonov, 2006). We utilized geographical location data for company headquarters and matched it with investors' registered addresses based on postal codes. Employing the Google Maps API, we obtained latitude and longitude coordinates for each company's headquarters and the investors' registered addresses. Subsequently, we calculated the pairwise distance (in kilometers) from each investor to every company using the Haversine formula, designed for computing surface distances between any two points on a sphere. We then performed a weighted summation of distances for companies held in each investor's portfolio, considering the weights associated with each holding.

Extrapolation: Since superior performance in the recent past is salient to retail investors, we proxy for each stock's salience score by its excess return over the market return in the preceding three

months. We then aggregated these excess returns at the investor level by value-weighting them based on the investor's portfolio investment value for the respective stocks.

Lottery Preference: We employed three variables to represent the lottery-like characteristics of stocks: the relative size of prices (using open, close, high, and low prices), idiosyncratic volatility, and idiosyncratic skewness, following the definition outlined in Kumar (2009). The idiosyncratic volatility and idiosyncratic skewness of an investor in any given month are estimated based on the daily CAPM Model from a 120-day rolling window (i.e., approximately six months) right before the beginning of the month.

Past Performance: We employed the investor's portfolio returns over the preceding three months as a proxy for their past investment ability.

Portfolio Value: To capture the potential wealth effect, we also include the total market value of the stocks held by investors in the previous month as an investor characteristic.

Firm Characteristics as Preditors: In our baseline analysis, we constructed 23 of the most important asset pricing anomalies as the holding-weighted stock characteristics based on investors' portfolios. Since these stock characteristics are constructed following the literature (e.g., Jensen, Kelly, and Pedersen, 2023), we do not explain them in detail.

Rank Normalization: After constructing the two sets of anomalies, we also apply rank normalization to ensure equal power of these proxies. More specifically, in each month, we rank each characteristics-based or behavioral anomaly in the cross-section of all retail investors (between 0 and 1). These ranks will be used as the preditors of investor returns.

Table 1 tabulates these variables and provides their detailed definitions. The Online Appendix (Table IN1) presents the summary statistics of our main variables. All portfolio-level variables have a reasonable distribution. Based on these summary statistics, it is reasonable to examine further how behavioral biases and firm characteristics affect retail investors' investment returns. We will undertake this task in the next section.

C. Machine-Learning Models

We employ a list of machine learning models, including Lasso, Ridge, Random Forests, and Neural Networks, to examine how behavioral biases and firm characteristics affect retail investors' investment returns. Below we describe their main algorithms.

C.1 Lasso and Ridge

When the number of predictors in a model is substantial, simple linear models may struggle to effectively fit the data, potentially leading to overfitting issues. Lasso and Ridge are both grounded in the linear assumption. However, unlike simple linear models, the objective function of these two models incorporates regularization. Specifically, these models no longer seek to minimize the error between fitted and observed values solely—they also assign penalties for the magnitude of linear model parameters. Lasso penalizes the first moment of model parameters, denoted as "l1" parameter penalization, whereas Ridge penalizes the second moment, known as "l2" parameter penalization. Specifically, a model with regularization can be expressed in the following form:

$$\min_{\beta \in R^p} \left\{ \frac{1}{N} \left|\left| y - X\beta \right|\right|_2^2 + \lambda \left|\left| \beta \right|\right|_1 + \gamma \left|\left| \beta \right|\right|_2^2 \right\},$$

where β is the model parameters, λ and γ are regularization coefficients, and in the context of the Lasso model, γ=0, while for the Ridge model, λ=0.

C.2 Random Forests

Random Forest is an ensemble learning method that operates by constructing a multitude of decision trees during the training phase. Decision trees are commonplace in machine learning, offering a non-linear modeling approach, which differs from traditional linear models. Notably, decision trees are non-parametric models. A tree is constructed by iteratively splitting the dataset into subsets, forming successive child nodes. The splits are based on predictor variables that most effectively discriminate among potential outcomes.

Random Forests employ an ensemble strategy by averaging multiple deep decision trees, each trained on different segments of the same training set. This approach aims to mitigate variance and offers a robust modeling technique.

C.3 Neural Network

Neural networks are currently highly popular models in various application domains, having achieved tremendous success in fields such as natural language processing and computer vision. According to the *Universal Approximation Theorem* (Kurt et al. 1989), neural networks can approximate any function between input x and output y. For our estimation, we employed a multi-

layer perceptron (MLP) network, also known as a feed-forward network (FNN), a standard and widely applicable neural network model in the financial market.

A multi-layer perceptron network consists of an input layer, an output layer, and one or more hidden layers. In each layer of the multi-layer perceptron, the input undergoes a linear transformation followed by an element-wise non-linear transformation (activation function). For the l-th layer of the MLP, its computational process can be expressed as follows:

$$X^l = g\left(W^{(l)T}X^{(l-1)} + b^{(l)}\right),$$

where $X^{(l-1)} \in R^{D^{l-1}}$ is the input to the l-th layer of the network, $W^{(l)} \in R^{D^{l-1} \times D^l}$ and $b^{(l)} \in R^{D^l}$ are the learnable parameters for the l-th layer, and $g(*)$ is the non-linear activation function. Notably, the output layer does not utilize a non-linear activation function. Instead, it directly aggregates the output from the previous layer through a linear mapping to form predictions for future returns, i.e.,

$$X^{output} = W^{(output)T}X^{(-1)} + b^{(output)}.$$

As for the choice of the activation function, we employ the most common rectified linear unit function (ReLU):

$$g(z) = ReLU(z) = \max(z, 0).$$

C.4 Residual Learning

In addition to the standard feed-forward network, we also employ a more advanced neural network: Residual Neural Network (ResNN). ResNN is a deep learning model in which the weight layers learn residual functions concerning the layer inputs. It is characterized by skip connections, termed "residual connections," which perform identity mappings and are combined with the layer outputs through addition. This architecture facilitates the training of deep learning models with tens or hundreds of layers, leading to improved accuracy as the depth of the network increases. Notably, the concept of identity skip connections, or residual connections, extends beyond Residual Networks and finds application in various other models such as Transformer models (e.g., BERT and GPT models like ChatGPT).

Following the seminal work of He et al. (2015), the computation for each layer can be expressed in the following form under the paradigm of residual connections:

$$X^l = g\left(W^{(l)T} X^{(l-1)} + b^{(l)}\right) + X^{(l-1)},$$

where $X^{(l-1)}$ denotes the original input, which is added back (e.g., through concatenation) to the output of the layer, X^l.

This design allows *ResNN* to be trained deeper and more easily optimized due to three beneficial features when compared to traditional neural networks. First, instead of finding the optimal function or true information for optimization, which often leads to overparameterization when the optimization process is complex, each layer of neural in the residual learning framework only needs to figure out the additional information—compared to the inputs—that helps to improve maximization. In other words, the goal now becomes to augment the initial data by providing additional information. This design helps simplify the task for each layer of neural, and it also gives each neural better information to achieve its simplified task.

Second, residual learning often combines several layers of neural into a block to facilitate residual connections. In this case, $X^{(l-1)}$ and X^l in the above equation become the original input and the output of the block. The benefit of this design is that it allows information to flow both within and aside the block, which shortens the gradient path of the blocks to speed up the training. This feature can significantly mitigate the issue of vanishing gradients in deep learning. Lastly, residual learning often adopts the modularity principle, which means building the neural network based on blocks with similar structures. Modularity allows for more blocks and deeper learning (see, e.g., Sun and Guyon 2023 for a recent survey).[15]

Among the three features, the first could be especially applicable to the financial market. This is because financial information is often expressed vis-à-vis a benchmark. For instance, mutual funds are often benchmarked against an index. In this case, the excess return of a fund provides the most important information about its operation. For another example, investment returns are typically adjusted by risk factors, allowing risk-adjusted returns to reveal important properties of the investment strategy. Roughly speaking, this feature of *ResNN* also allows machine learning to be benchmarked against some economically important inputs. As we will see shortly, this feature can help further improve the performance of neural networks.

[15] https://arxiv.org/abs/2310.01154

D. Data sampling and Optimization

We employed the cross-validation method to train and assess the performance of our model. Following the approach outlined by Kaniel et al. (2023), we uniformly divided the entire dataset into three parts. In each iteration, we trained the model on two of the folds and tested its performance on the remaining fold. This approach offers the advantage of testing the model on every sample in the dataset, enhancing the robustness of model comparisons by mitigating the influence of specific periods. Additionally, within the training data, we randomly set aside 30% of the samples for validation.

After partitioning the data, we employed a gradient-based approach to train neural networks. There are various neural network training strategies, and a common solution is to utilize the Adam optimizer. To enhance the optimization speed and performance of the model, the Adam optimizer randomly selects a subset of samples (batch) from the training data for gradient updates in each iteration.

A key parameter of the Adam optimizer is the learning rate, which dictates the step size for updates along the gradient direction. Since a well-chosen learning rate involves a trade-off between convergence speed and avoiding overshooting, it is essential to dynamically adjust the learning rate based on the state of the training process. Therefore, we implement a learning rate scheduler during training. A learning rate scheduler is a predefined framework that modifies the learning rate between epochs or iterations as the training advances. In particular, we employ a learning rate decay strategy, gradually reducing the learning rate as the training progresses.

Neural networks often exhibit strong expressive power and the ability to fit arbitrary functions, but they are also susceptible to overfitting. Overfitting occurs when a neural network performs well on the training data but poorly on unseen testing data. It occurs when a neural network memorizes the noises and details of the training data excessively while neglecting the overall distribution of the data, resulting in a decrease in the model's generalization ability.

To mitigate overfitting, we employ EarlyStopping and Dropout. EarlyStopping is a regularization technique in model training. If the model's performance on the validation dataset does not improve consistently, training is halted to prevent the model from excessively fitting the training data. Dropout involves ignoring the output of certain hidden layer nodes during training,

setting these nodes' output values to zero. This approach reduces interactions between hidden layer nodes, thereby minimizing overfitting in neural networks (Hinton et al., 2012).

For the specific parameters of the model, we employed a three-layer fully connected neural network with 32, 16, and 8 neurons in each layer, respectively. The learning rate is set at 0.001, and the maximum training epochs are set at 150. Additionally, we implement an Early Stopping mechanism with a patience setting of 3, meaning that the training would be terminated if there was no improvement in the model's performance for more than three epochs.

III. Predicting Total Returns for Retail Investors

We now use all the aforementioned models to predict retail investors' total investment returns.

A. The Portfolio Analysis Approach

Our baseline tests involve a machine-learning-based portfolio analysis. We first use all the models to predict retail investors' total investment returns. We then sort retail investors into five quintiles according to predicted returns, with the *High* and *Low* groups consisting of 20% of predicted winners and losers among investors, respectively. Finally, we report the value-weighted out-of-sample returns of the high and low groups as well as their return difference. we also use the locally estimated three-factor and four-factor models to adjust these returns.

B. The Performance of Model Selected Investors

Table 2 tabulates the predicted returns of investor quintiles. Columns 1-3 present average monthly returns and alpha adjusted through local FF-3 and Carhart-4 models. Columns 4-6 depict results for the high group, while columns 7-9 detail outcomes for the high minus low return.

We observe that the two Neural Network models outperform other models in predicting retail investors' total returns. In particular, both FNN and ResNN identify investors who can generate significantly positive out-of-sample returns. Indeed, column (4) reports that the top 20% of retail investors of the two models can generate a monthly return of 1.5% and 1.2%, which remains highly significant with a similar economic magnitude when risk-adjusted (e.g., 1.7% and 1.4% adjusted by four factors, as reported in column 6).

In contrast, all other models fail to predict winners. Given how difficult it is to predict four-factor adjusted superior mutual fund performance in the US (e.g., Carhart 1997), the superior retail

performance strongly suggests that Neural Network models capture important properties of retail investors.

On the loser side, ResNN performs the best. Column (1) reports that the ResNN-predicted *Low* group delivers a significantly negative monthly return of -3.1%, allowing the *High* group to outperform the *Low* group by 4.4% per month in column (7). FNN-predicted *Low* group delivers a slightly lower and marginally significant return of -2.5%. Although the low group's risk-adjusted performance becomes insignificant, its *High* group outperforms the *Low* group by 4.0% per month and remains highly significant after the risk adjustments.

Among other models, Lasso and Ridge can select retail investors that deliver marginal negative returns, whereas OLS and Random Forest do not exhibit significant predicting power on the negative return side. As a result, the *High* group of Lasso and Ridge can significantly outperform the *Low* group by 1.6% and 2.5%, as reported in column (7). Although the *Low* groups selected by OLS fail to deliver significant returns, its high-minus-low spread remains significant at 2.5%.

Collectively, we observe that the two Neural Network models outperform other models in predicting total returns for retail investors. In particular, both FNN and ResNN can identify retail investors who can consistently deliver positive returns. This predictive power is striking, given how difficult it is for professional investors—such as mutual funds—to deliver out-of-sample performance. Of course, the difficulty in predicting risky adjusted superior mutual fund performance is typically based on traditional OLS methods (e.g., Carhart 1997), whereas machine-learning tools are typically more powerful to predict the performance of mutual funds (e.g., Li and Rossi, 2020; DeMiguel, Gil-Bazo, Nogales, and Santos, 2023; Kaniel, Lin, Pelger, and Van Nieuwerburgh 2023) and hedge funds (Wu, Chen, Yang, and Tindall, 2021).

In our setup, Neural Networks outperform OLS and other machine-learning models. This evidence strongly suggests that Neural Networks capture crucial characteristics of retail investors that contribute to their returns. As a result, Neural Networks provide a reliable tool to further analyze how behavioral biases and firm characteristics affect retail investors' investment returns.

IV. The Analysis of Predictors

We now employ Neural Network models to investigate the relative importance of behavioral biases and firm characteristics in affecting retail investors' investment returns.

A. The Stand-alone Power of Behavioral Biases and Firm Characteristics

Our previous analysis uses both behavioral biases and firm characteristics as predictors of the Neural Network models. However, we can use each set of predictors alone, which can shed light on the relative importance of these predictors in predicting retail investors' investment returns.

The results are tabulated in Table 3. The first row reports the results when the FNN is trained only based on characteristics-based predictors. In this case, the Neural Network model fails to select the *High* and *Low* groups that can deliver significantly positive and negative out-of-sample returns. Nor can firm characteristics alone predict a significant *High*-minus-*Low* return spread.

The second row reports the results when behavioral biases are used alone by FNN. Unlike the first row, using behavioral biases alone can predict a significant *High*-minus-*Low* return spread of 3.3% per month. Moreover, its power mostly arises from the negative return (loser) side, with the Low group delivering a -2.5% return. Both observations hint at the relative importance of behavioral biases in affecting returns.

The third row reports outcomes from simultaneously incorporating behavioral biases and firm characteristics by the FNN model. Although this result has also been reported in the previous table, the side-by-side comparison between this and the behavioral-only result can help reveal more properties of the FNN estimation. We first observe that the simultaneous use of both behavioral biases and firm characteristics enables FNN to predict a significant High-minus-Low return spread. This return spread (4.0%) is larger than the case when the FNN algorithm is trained only by behavioral biases.

However, the High-minus-Low return spread of the third row is primarily driven by the positive returns generated by the High group, not by the Low group. Compared to the second-row result, i.e., FNN can select Low-group investors to deliver significantly negative returns when only behavioral biases are used, we find that including more predictors (i.e., firm characteristics) diminishes the model's ability to identify the Low-group investors. Conceptually, including more preditors should not impede an optimization algorithm because its optimization should include the cases with the more restrictive set of predictors. However, empirically, the increased complexity of parameter space may subject neural network optimization to common issues such as overparameterization and vanishing gradients, leading to a partial loss of predictive capacity (He

et al., 2015). In particular, the potential efficiency loss concentrates on the impact of behavioral bias on the Low-group investors.

The above issue motivates us to adopt the residual learning framework proposed by He et al. (2015) to facilitate better training of neural networks. Specifically, when passing through a hidden layer (we have three in total), we concatenate the initial set of economically important predictors (i.e., the behavioral biases) with the output from the layer as the combined inputs for the subsequent layer. As mentioned earlier and demonstrated by He et al. (2015), such an approach can significantly enhance the optimization efficiency of neural networks by mitigating common issues, such as overparameterization and vanishing gradients, across the hidden layers.[16]

The last row affirms that our designed Residual Neural Network can effectively predict future returns for both loser and winner groups. For the loser group, it yields a substantial -3.1% excess return and a -2.6% alpha after adjusting with the Carhart four-factor model. Conversely, for the winner group, it generates a notable 1.2% excess return and a 1.4% four-factor adjusted alpha. In other words, ResNN can identify investors with both good and bad performance. Its out-of-sample High-minus-Low return spread is approximately 10% higher in relative terms than that of FNN.

B. Behavioral vs. Firm Characteristics in Alternative Objectives

Our previous analysis leans toward the relative importance of behavioral biases when predicting the total returns of retail investors. However, investors' total returns may originate from two different sources: from holding an existing portfolio for a given period of, for instance, a month (i.e., holding returns) and from newly initiated trading during the month (i.e., trading returns). Behavioral theories suggest that the motivations to initiate trading may differ from those of continuation. For instance, the well-documented disposition effect (e.g., Shefrin and Statman, 1985) suggests that unrealized capital gains motivate investors to trade (i.e., sell), whereas unrealized capital losses incentivize investors to hold onto losing assets and thus affect holding returns. For another example, salient information, such as extreme stock prices, may also attract investors' attention to initiate new trades according to the salience theory (Bordalo, Gennaioli, and

[16] Our ResNN septically preserve the important predictors on Low-group investors. On the High-group side, we also note that, when used alone, neither firm characteristics nor behavioral biases would allow the FNN to successfully select the High group of investors to deliver superior out-of-sample returns. Hence, the interactions between investor behavioral biases and holding-based firm characteristics provide the source for FNN to predict good returns for investors. Such potential interaction effects remain effective in the residual network framework.

Shleifer 2012; 2013; 2020), even though traditional financial theory suggests investors should pay more attention to returns rather than the level of prices. These discussions suggest that the relative importance of behavioral biases may or may not hold for both components of returns.

We therefore ask how behavioral biases and firm characteristics may affect the two sources of returns. Before we can scrutinize this question, we need to investigate whether neural network-selected investor groups can help predict them. Hence, we alter the training goals of neural networks, asking them to predict trading and holding-based returns.

The results of this test are tabulated in Table 4. Panel A tabulates the VW out-of-sample monthly total returns generated by the Low and High groups and the High-minus-Low spread. In columns (1) to (3) and (4) to (6), the training goal is to predict the trading and holding-based returns of investors, respectively. In addition to training goals, we also differentiate the impact of different predictors. Similar to the previous table, the first three rows demonstrate the prediction power of FNN when different sets of predictors (i.e., firm characteristics or behavioral biases) are used. The last row reports the results for the residual neural network.

Across columns (1) to (3), firm characteristics alone fail to predict good, poor trading-based returns or a significant *High*-minus-*Low* spread. In contrast, behavioral biases alone allow neural networks to predict a significant trading-based return spread, though the prediction power arises mainly from the poor performance of Low-group investors (but not the good performance of High-group investors). Hence, behavioral biases seem to provide the most important ground for generating such poor performance.

Unlike trading returns, columns (4) to (6) suggest that firm characteristics alone allow neural networks to identify good performers and a marginally significant High-minus-Low spread. Behavioral biases alone can still generate a High-minus-Low spread. The joint use of both sets further improves the magnitude of the High-minus-Low spread. As a result, both firm characteristics and behavioral biases appear useful for neural networks to identify retail investors that deliver superior or poor performance.

Panels B and C present the risk-adjusted performance of the Low and High groups and the High-minus-Low spread. Their layout is the same as Panel A. We observe that our results remain highly robust under risk adjustments.

Between this and the previous table, we observed that behavioral biases may play a more important role in affecting trading returns, though firm characteristics also appear informative for holding-based returns. The caveat of these results is that they cannot quantify the importance of the two categories of predictors, not to mention that of individual predictors. To address this problem, we next turn to variable gradient analysis to quantify the impact of factors. Following the literature (e.g., Kaniel, Lin, Pelger, and Van Nieuwerburgh 2023), we mainly use the traditional FNN model to investigate and demonstrate this standalone predicting power.

C. Variable Gradient Analysis

We follow the methodologies proposed by Sadhwani et al. (2020) and Horel and Giesecke (2020) and conduct variable gradient analysis to demonstrate the relative importance of each variable when behavioral biases and stock characteristics are both used. More explicitly,

$$Importance(x) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{i=1}^{N_t}\left(\frac{\partial R_{i,t+1}^{pred}}{\partial x_{i,t}}\right)^2,$$

where T represents the number of periods in the data, and N_t denotes the total number of investors in the t-th period. The partial derivative measures the gradient of the model's predicted output with respect to each variable.

It is noteworthy to discuss two features of the estimation. First, the gradient can be positive or negative, so its square term is used to gauge its importance. For linear regression models, the partial derivative is simply the regression coefficient. Intuitively, a larger partial derivative implies a bigger influence of a variable on the model's output, indicating greater importance in predicting future returns. Second, the traditional neural network algorithm is more appropriate for this gradient analysis. Indeed, although the residual neural network can help improve model performance, concatenating the output from a layer with the initial set of economically important predictors (i.e., the behavioral biases) may introduce a selection bias on the importance of these variables. As a result, we compute the *Importance(x)* for each predictor when we use FNN to predict the total return of investors based on both firm characteristics and behavioral biases.

The variable gradient analysis enables us to address two economic questions. First, between behavioral heuristics and firm characteristics, which one contributes more to the investment returns

of retail investors? Second, which individual factors exert the most substantial influence? We use graphic plots to intuitively display the answers to both questions.

To address the first question, we sum up the importance of all predictors falling into one of the following two categories: behavioral heuristics or firm characteristics. We then scale our estimation so that the total importance of the two categories equals 100%. Figure 2 illustrates the Relative Importance of Behavioral Bias vs. Firm Characteristics.

We observe that the relative importance of Stock features gradually decreases as the prediction target shifts from Holding Return to Trading Return. Specifically, for Holding Return, the relative importance of characteristics and behavioral heuristics is 47.8% and 52.2%, respectively. In predicting total returns, we observe that the joint explanatory power of behavioral biases slightly exceeds that of firm characteristics, where the relative importance is 36.5% for Stock features and 63.5% for investor features. However, behavioral bias predictors dominate when it comes to predicting trading returns, accounting for nearly 95% of the total predictive power. These observations confirm the importance of behavioral biases for retail investors, particularly in trading and its associated short-term returns.

We now turn to the second question. The gradient analysis enables us to directly observe the importance of each variable across the two categories of predictors in the FNN estimation. We plot the distribution of variable importance in Figure 3, which ranks predictors based on their importance.

The FNN identifies diversification, portfolio turnover, and momentum as the top three leading factors to influence the total returns of retail investors. The first two variables are related to the behavioral biases of under-diversification and overconfidence. Under-diversification is among the most common features of retail investors (e.g., Barber and Odean 2000; Benartzi and Thaler 2001), as many investors may not fully understand the benefits of diversification (Lusardi and Mitchell 2011). Overconfidence often causes investors to trade too aggressively, allowing their high portfolio turnover to reduce their welfare (Odean 1998; Barber and Odean 2000). The third denotes

perhaps the most famous anomalies in the literature. It is interesting to observe that behavioral biases occupy two out of the top three factors affecting investment returns.[17]

Furthermore, behavioral biases still play a relatively more important role in predicting trading returns than holding returns. We observed that portfolio turnover, the disposition effect, and the degree of portfolio diversification emerge as the three most important factors in predicting new trading returns, followed by the opening and closing price of stocks. In other words, the disposition effect emerges as one of the leading predictors for trading, in addition to turnover and diversification.

Interestingly, (under)diversification, portfolio turnover, and momentum are still the three leading factors to affect holding-based returns. Unlike total returns, however, momentum becomes more prominent and surpasses turnover in terms of importance. The observation is intuitive: given the return predictability of momentum, its influence should be stronger for holdings.

The Online Appendix (Figure IN3) further plots the joint importance of the top three anomalies in predicting total returns, trading returns, and holding returns over the years. Two major observations emerge. First, the relative importance of these top anomalies remains quite stable over time, suggesting that they capture important and persistent economic grounds that influence investors' total returns.

Second, the top three anomalies exhibit a disproportionally high degree of relative importance. For instance, as far as total returns are concerned, the top three anomalies account for about 29% of the total importance of all 39 anomalies (which is normalized to 1 in the plot). For trading returns, the top three anomalies (all behavioral) span an astonishing 87% of total importance. These observations suggest that not all anomalies are equally important. On the contrary, anomalies exhibit a highly skewed distribution in their relative importance in influencing returns.

Figure 4 intuitively demonstrates the skewed distribution of anomaly importance. More specifically, Panels A, B, and C plot the cumulative distribution of the average relative importance of all behavioral and characteristics-based anomalies in predicting total returns, trading returns, and holding returns, respectively. The sequence of anomalies from left to right is determined by

[17] Note that observed portfolio diversification and turnover may also be related to alternative sources of biases, such as local bias, local information, and financial literacy. Since our purpose is to compare behavioral biases and stock characteristics, we do not further nail down the economic sources of behavioral bias.

their relative importance (i.e., the first anomaly is the most important one in predicting returns). Take total returns as an example. We already know that the top three anomalies span approximately 29% of relative importance. Panel A further shows that the cumulative relative importance of the top 5 and top 10 anomalies amounts to 38% and 53%, respectively. In other words, the top ten anomalies are more important than the remaining 29 anomalies in impacting total returns.

The distribution of anomaly importance is even more skewed for trading returns. There, the cumulative importance of the top 3, 5, and 10 anomalies amounts to 87%, 91%, and 96%, respectively. In other words, three behavioral anomalies suffice to capture the most important impacts on the short-term performance of newly initiated trades, suggesting that a parsimonious framework for behavioral biases could be especially important for these newly initiated trades. In contrast, explaining retail investors' total and holding returns likely requires a higher-dimensional factor structure despite the skewed distribution of the relative importance of anomalies.

D. Implications on Market Efficiency

Thus far, our results suggest that machine learning tools such as neural networks can properly predict retail investors' returns based on a relatively parsimonious set of factors. One notable feature of our results is that the two neural network models can identify not only poor-performing retail investors but also those who can deliver superior performance. The latter has important economic implications. For instance, in traditional theories, retail investors are often perceived as noise traders, whose trading brings risk to rational investors and thus weakens market efficiency (De Long et al., 1989; Shleifer and Summers, 1990; Shiller, 2003). In contrast, our results are in line with an emerging body of studies recognizing that at least some retail investors may trade smartly (Kaniel et al., 2012; Kelley and Tetlock, 2017; Boehmer et al., 2021; Coval, Hirshleifer, and Shumway, 2021). This section aims to provide further insights into this literature by undertaking additional tests to examine how the trading activities of the *High* and *Low* groups of retail investors—i.e., winners and losers—impact market efficiency.

More specifically, we investigate the role of retail investors in shaping variance components related to firm-specific private information and noise. In the literature, Brogaard et al. (2021) propose a VAR framework to decompose the variance of stock returns into four components based on their economic origins: firm-specific private information, firm-specific public information, market-wide information, and noise. We follow their approach to estimate the VAR system yearly

based on daily returns during a year. We are particularly interested in the two variance components related to firm-specific private information and noise because they directly speak to the role played by various types of investors. When there is no confusion, we label them the *Firm-specific Private Information Share* and *Noise Share* in later analysis.

To validate the role of retail investors, we first define trading intensity as the total trading volume in a stock made by a specific group of investors (such as winners and losers), scaled by the stock's overall trading volume. Since the variance decomposition is conducted yearly, we also construct the measure of trading intensity yearly by averaging monthly observations over the year. Lastly, we link *Firm-specific Private Information Share* and *Noise Share* to the contemporaneous trading intensity of winners and losers in panel specifications, which also control for the time and stock fixed effects.

The results are tabulated in Table 5, with Panels A and B reporting the regression coefficients for noise and firm-specific private information, respectively. Each column represents the results of a particular model. Across all columns in Panel A, the trading intensity of losers is significantly associated with noise. This pervasive pattern aligns with the economic prototype of "noise traders" for retail investors, whose trading adds noise to stock returns.

However, the observations for winners are drastically different. While most models report an insignificant role of winners, the two neural networks stand out as uniquely reporting significant negative coefficients. In other words, the trading of winners identified by neural networks can help enhance market efficiency by reducing noise contained in stock prices, suggesting that these investors may have collected valuable information and traded accordingly. Among the remaining models, *Random Forest*-identified winners also marginally reduce noise—but to a lesser extent in both statistical power and economic magnitude.

Panel B further confirms the role of neural network-identified winners in influencing the firm-specific private information component of return variance. The trading intensity shows significant positive coefficients, indicating that their trading enhances the proportion of firm-specific private information embedded in stock prices. The informational role of smart retail investors is plausible in emerging markets such as India, where institutional investors often contribute positively to liquidity but are less informed compared to their counterparts in other countries (Chui, Subrahmanyam, and Titman, 2022). This gap creates an opportunity for sophisticated retail

investors who can gather firm-specific private information from diverse channels. Their trading not only enables them to attain higher returns but also incorporates such information into stock prices, thereby enhancing market efficiency.

Our Online Appendix provides additional tests on market efficiency measures, including variance ratios (Lo and MacKinlay, 1988; Campbell, Lo, and MacKinlay, 1997) in Table IN2 and market delay (Hou and Moskowitz, 2005; Griffin, Kelly, and Nardari, 2010) in Table IN3. The primary message of these results is that trading by neural network-identified winners reduces variance ratios and market delay, thereby enhancing market efficiency. Among the remaining models, only *Random Forest* can identify winners to reduce variance ratios and market delay—but the influence is marginal. As such, the two neural networks remain the most powerful tool for identifying retail investors who are plausibly informed.

Collectively, these findings validate the ability of neural networks to identify smart retail investors capable of achieving superior performance. Perhaps even more notably, they carry important normative implications for market efficiency, demonstrating that at least some retail investors are informed and play a vital role in enhancing price informativeness in the market.

V. Additional Analyses

This section provides additional analysis to shed light on the economic interpretation and robustness of our existing results.

A. The Impact of Market Conditions

We first examine whether market conditions, such as volatility, may affect the performance difference between the top and bottom quintiles of investors. Since behavioral biases are the main driving force of the difference, a significant impact of market conditions would suggest an interaction effect between market conditions and behavioral biases. We test the significance of this interaction by linking the High-minus-Low spread to the aggregate order flows, return, volatility, and skewness of the local market.

The results are reported in Table 6, where Models (1)-(4) and (5)-(8) use FNN- and Res-NN generated High-minus-Low return spread as the dependent variable. Columns (1) and (5) tabulate the time series regression when the four factors are included. Its constant represents the four-factor-adjusted performance as reported in Table 2. These two benchmark models show that the High-

minus-Low spread is generally unrelated to market returns. Instead, momentum and HML are the two factors that correlate with the spread.

The next two columns include separately the volatility and skewness of the market returns. We observe insignificant regression coefficients, suggesting that the performance spread between the top and bottom quintiles of investors is not affected by market volatility or skewness. Collectively, the first three moments of the market have little impact on performance difference, suggesting that the relative advantage and disadvantage of investors in generating performance—or the economic grounds that lead to this difference, such as the leading behavioral heuristics estimated by neural networks—do not interact with market conditions.

Lastly, columns (4) and (8) include investor sentiment as an additional aggregate variable. Unlike most existing studies, we directly observe the orders of *all* retail investors in the market. Hence, we proxy for sentiment by the aggregate order imbalance of all retail investors in our sample, calculated as the buy-minus-sell orders scaled by the buy-plus-sell orders. We observe that sentiment negatively affects FNN's High-minus-Low spread but has little impact on the spread constructed by ResNN. As a result, our tests report an inconsistent role of investor sentiment in our setup. Since the selection made by ResNN is superior to that by FNN, we can interpret the result as the more advanced neural network tool could better identify the subgroups of investors who are less affected by investor sentiment.

The Online Appendix (Table IN4) provides another robustness check. When training models, we divide our sample period into three subperiods. A natural concern is whether the High-minus-Low return spread only concentrates on one or two subperiods. However, when we use dummy variables to indicate whether the specific month belongs to the first or second subperiods, we observe that subperiods do not impact the return spread.

Collectively, our tests suggest that the economic source captured by our neural network models, such as behavioral heuristics, may affect investors' relative performance regardless of market conditions or testing periods. When the more advanced neural network model is used, relative performance is also independent of aggregate investor sentiment.

B. The Persistence of Performance

Next, we investigate the persistence of performance and related welfare implications. Our previous analysis focuses on one-month ahead performance. If top-quintile investors outperform bottom-quintile ones only in the next month but subsequently underperform, our baseline findings may have little impact on investor welfare. To mitigate this concern, we use neural networks to predict the total returns of retail investors for up to six months.

Table 7 tabulates the results, with Panels A and B reporting the respective estimations from FNN and ResNN. In each panel, column "T+k" (k = 1, 2,…, 6) reports the monthly High-minus-Low return spread or its risk-adjusted performance k months after constructing the High and Low quintiles. Column "T+1" presents the results for one-month ahead performance, which are the main results that we reported in previous tables.

We observe that the outperformance remains significant (albeit with declining magnitude) for about three months and then dissipates to insignificant in the fourth to sixth months. Importantly, we do not observe a reversal, suggesting that neural networks may capture long-term economic determinants of retail investors' performance. As a result, the factors we identify may carry significant normative implications to impact the welfare of retail investors when they participate in the stock market.

C. Expanded List of Firm Characteristics

In the main analysis, we use 23 holding-weighted stock characteristics and 13 proxies of behavioral heuristics to predict investor returns, and we observe that behavioral anomalies play a more important role. Although our list already includes all categories of stock characteristics (e.g., Jensen, Kelly, and Pedersen, 2023) and the most important characteristics therein, one potential conjecture could be that including more firm characteristics can also increase their relative power. To investigate this conjecture, we follow the cauterizations of Jensen, Kelly, and Pedersen (2023) to incorporate more available accounting variables to create a list of 67 firm characteristics. The expanded list is tabulated in our online Appendix (Table IN5).

We re-estimate our baseline analysis and tabulate the results in Panel A of Table 8. We observe that neural networks exhibit a very similar power in predicting out-of-sample performance. For instance, the four-factor adjusted High-minus-Low return spread amounts to 0.035 and 0.043 for FNN and ResNN. Compared to the estimates of Table (0.031 and 0.041), expanding the list of stock characteristics does not allow the two models to substantially improve the out-of-sample

performance of selected investors. Indeed, the Online Appendix (Table IN5) also reports the return differences between estimates from the expanded firm characteristics and the original sample. Across all return estimates, whether risk-adjusted or not, and regardless of the models we use, the difference remains insignificant. This observation suggests that our original list of firm characteristics already captures the most economic grounds to understand the investments of retail investors.[18]

Figure 5 further plots the relative importance of behavioral biases vs. firm characteristics. These numbers are also reported in Panel B. Behavioral biases still hold similar relative importance when compared to firm characteristics. For instance, behavioral biases and firm characteristics now exert approximately 56.6% and 43.4% relative importance, respectively, for investors' total returns. Hence, the significant expansion of firm characteristics does not transfer into a substantial increase in relative importance.

D. Removing More Microcap Stocks

In the main analysis, we exclude 30% of small stocks because these stocks are difficult to trade in emerging markets (Liu et al. 2019) and may overstate the predicting power of machine learning models (Avramov, Cheng, and Metzker 2019 and Cong et al., 2020). Table 9 further examines whether our main conclusions are robust to different removal thresholds (e.g., 20% or 40%). Columns 1-3 tabulate the raw or risk-adjusted *High*-minus-*Low* return spread when 20% of small stocks are removed. Columns 4-6 report similar results when 40% of small stocks are removed. Neural Network models (FNN and ResNN) continue to generate significantly out-of-sample returns, which dominate all other alternative machine-learning or OLS models. These patterns suggest that our conclusions derived from Neural Network models are robust to the less or more strict control of small stocks.

It is also worth noting that ResNN outperforms FNN in both cases. For instance, the two models deliver, respectively, 4.4% and 3.4% value-weighted *High*-minus-*Low* return spread when 40% microcap stocks are removed. Hence, ResNN outperforms FNN by almost 30% in generating out-of-sample returns. Indeed, ResNN outperforms FNN across all empirical specifications, suggesting that ResNN may provide a superior tool for return predictive analysis.

[18] This observation is reasonable as many characteristics in the expanded list may share similar economic foundations. For this reason, our main test focuses on the 23 holding-weighted stock characteristics because of its parsimony.

E. Model Comparisons in Predicted Performance

Lastly, we more formally report the differences in predicted performance across various models. Since the two neural network models outperform other models, we first investigate the difference between FNN and other non-neural network models. We then move on to tabulate the difference between the two neural network models. Our analysis focuses on the High-minus-Low return spread as the overall performance measure of each model. To assess the robustness of our results, we also systematically remove 20%, 30%, and 40% of small stocks and use the local three or four factors to adjust the High-minus-Low return spread.

The results are tabulated in Table 10. The first four rows compare FNN to other non-neural network models. Regardless of the sample we use or the factor models we use to adjust returns, the High-minus-Low return spread of FNN significantly outperforms other non-neural network models across all empirical specifications. These observations reveal the striking predictive power of neural networks in general.

Between the two neural network models, ResNN significantly outperforms FNN across all empirical specifications, as reported in the last row. ResNN generates between 9 and 13 basis points (bps) more in monthly four-factor adjusted alpha than FNN. These results again confirm the advantage of residual neural networks for analyzing large financial data like ours.

Conclusions

This paper employs various machine learning models to analyze the returns for millions of retail investors in India. We observe that Neural Network outperforms other models, including traditional linear OLS models, in predicting investor returns. In particular, the more recently developed Residual Neural Network (*ResNN*) exhibits superior power in identifying both good and bad out-of-sample performance. Such a predicting power suggests that Neural Network models comprehend important information about investors that contributes to their returns.

We further conduct variable gradient analysis, which indicates that behavioral biases, in general, play a more important role than holding-weighted firm characteristics to affect retail investors' investment returns. We identify diversification, portfolio turnover, and momentum as the leading factors influencing investors' total returns. Turnover, the disposition effect, and diversification emerge as the three most important factors in predicting new trading-generated

returns. The explanatory power of firm characteristics and behavioral heuristics exhibits a skewed distribution. A highly parsimonious anomaly structure is feasible to explain new trading-generated returns, whereas a higher dimensional structure seems needed to explain total returns.

Our results call for further research, potentially utilizing state-of-the-art machine learning tools, to comprehensively understand the framework of how behavioral biases and firm characteristics can jointly influence the investments of retail investors.

References

Acharya, V.V., Steffen, S., (2020). The risk of being a fallen angel and the corporate dash for cash in the midst of COVID. The Review of Corporate Finance Studies, 9(3), pp.430-471.

An, L., (2016). Asset pricing when traders sell extreme winners and losers. Review of Financial Studies 29, pp.823--861.

An, L., Engelberg, J., Henriksson, M., Wang, B. and Williams, J., (2019). The portfolio-driven disposition effect. PBCSF-NIFR Research Paper.

Avramov, D., Cheng, S., and Metzker, L. (2023). Machine learning vs. economic restrictions: Evidence from stock return predictability. Management Science, 69(5):2587–2619.

Avramov, D., Cheng, S., Metzker, L., and Voigt, S. (2023). Integrating factor models. The Journal of Finance, 78(3):1593–1646.

Baker, M., Litov, L., Wachter, J. A., & Wurgler, J. (2010). Can mutual fund managers pick stocks? Evidence from their trades prior to earnings announcements. Journal of Financial and Quantitative Analysis, 45(5), 1111–1131.

Balasubramaniam, V., Campbell, J. Y., Ramadorai, T., & Ranish, B. (2023). Who owns what? A factor model for direct stockholding. The Journal of Finance, 78(3), 1545-1591.

Bali, T. G., Beckmeyer, H., Morke, M., and Weigert, F. (2023). Option Return Predictability with Machine Learning and Big Data. The Review of Financial Studies, 36(9):3548–3602.

Barber, B. M. and Odean, T. (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 Odean, T. (2013). Chapter 22 - the behavior of individual investors. volume 2 of Handbook of the Economics of Finance, 1533–1570. Elsevier.

Barberis, N., & Xiong, W. (2009). What drives the disposition effect? An analysis of a long ‐ standing preference ‐ based explanation. the Journal of Finance, 64(2), 751-784.

Barberis, N. (2018). Chapter 2 - psychology-based models of asset prices and trading volume., Handbook of Behavioral Economics - Foundations and Applications 1, 79–175. North-Holland.

Ben-David, I., & Hirshleifer, D. (2012). Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect. The Review of Financial Studies, 25(8), 2485-2532.

Benartzi, S. and Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. The American Economic Review, 91(1):79–98.

Bertrand, M., Mehta, P., & Mullainathan, S. (2002). Ferreting out tunneling: An application to Indian business groups. The Quarterly Journal of Economics, 117(1), 121-148.

Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. The Review of Financial Studies, 34(2), 1046-1089.

Boehmer, E., Jones, C. M., Zhang, X., & Zhang, X. (2021). Tracking retail investor activity. The Journal of Finance, 76(5), 2249–2305.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and asset prices. American Economic Review, 103(3):623–28.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. Quarterly Journal of Economics, 135(3):1399–1442.

Brogaard, J., Nguyen, T. H., Putnins, T. J., & Wu, E. (2022). What moves stock prices? The roles of news, noise, and information. The Review of Financial Studies, 35(9), 4341–4386.

Bryzgalova, S., Huang, J., & Julliard, C. (2023). Bayesian solutions for the factor zoo: We just ran two quadrillion models. The Journal of Finance, 78(1), 487-557.

Bryzgalova, S., Pelger, M., & Zhu, J. (2020). Forest through the trees: Building cross-sections of stock returns. Available at SSRN 3493458.

Calvet, L. E., Campbell, J. Y., & Sodini, P. (2009). Measuring the financial sophistication of households. American Economic Review, 99(2), 393-398.

Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). The econometrics of financial markets. Princeton, NJ: Princeton University Press.

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

Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. The Journal of Finance, 71(1), 267-302.

Chen, L., Pelger, M., & Zhu, J. (2024). Deep learning in asset pricing. Management Science, 70(2), 714-750.

Choi, J. J., & Robertson, A. Z. (2020). What matters to individual investors? Evidence from the horse's mouth. The Journal of Finance, 75(4), 1965-2020.
Chui, A., Subrahmanyam, A., & Titman, S. (2022). Momentum, reversals, and investor clientele. Review of Finance, 26, 217–255.
Cici, G. (2012). The prevalence of the disposition effect in mutual funds' trades. Journal of Financial and Quantitative Analysis, 47(4), 795–820.
Cong, L., Tang, K., Wang, J. & Zhang, Y., (2021), AlphaPortfolio: Direct Construction Through Deep Reinforcement Learning and Interpretable AI, working paper.
Coval, J. D., Hirshleifer, D., & Shumway, T. (2021). Can individual investors beat the market? The Review of Asset Pricing Studies, 11(3), 552–579.
Coval, J. D., & Moskowitz, T. J. (1999). Home bias at home: Local equity preference in domestic portfolios. The Journal of Finance, 54(6), 2045–2073.
Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. Journal of Finance, 53(6), 1839–1885.
DellaVigna, S. (2009). Psychology and economics: Evidence from the field. Journal of Economic literature, 47(2), 315-372.
DellaVigna, S. (2018). Structural behavioral economics. In Handbook of Behavioral Economics: Applications and Foundations 1 (Vol. 1, pp. 613-723). North-Holland.
Cochrane, J. H. (2011). Presidential address: Discount rates. The Journal of Finance, 66(4):1047–1108.
DeMiguel, V., Gil-Bazo, J., Nogales, F. J., and Santos, A. A. (2023). Machine learning and fund characteristics help to select mutual funds with positive alpha. Journal of Financial Economics, 150(3):103737.
Easley, D., Lopez de Prado, M., O'Hara, M., and Zhang, Z. (2021). Microstructure in the Machine Age. The Review of Financial Studies, 34(7):3316–3363.
Fama, E. F., and K. R. French. (1992). The cross-section of expected stock returns. Journal of Finance 47:427-65.
Fama, E. F., and K. R. French. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33:3-56.
Fama, E. F. and French, K. R. (1998). Value versus growth: The international evidence. The Journal of Finance, 53(6):1975–1999.
Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. Journal of financial economics, 116(1), 1-22.
Feng, G., Giglio, S., & Xiu, D. (2020). Taming the factor zoo: A test of new factors. The Journal of Finance, *75*(3), 1327-1370.
Filippou, I., Rapach, D., Taylor, M. P., & Zhou, G. (2023). Out-of-sample exchange rate prediction: A machine learning perspective. Available at SSRN 3455713.
Fischbacher, U., Hoffmann, G., & Schudy, S. (2017). The causal effect of stop-loss and take-gain orders on the disposition effect. The Review of Financial Studies, 30(6), 2110-2129.
Frazzini, A. (2006). The disposition effect and underreaction to news. Journal of Finance, 61(4), 2017–2046
Freyberger, J., Neuhierl, A., & Weber, M. (2020). Dissecting characteristics nonparametrically. The Review of Financial Studies, 33(5), 2326-2377.
Frydman, C., & Rangel, A. (2014). Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. Journal of economic behavior & organization, 107, 541-552.
Frydman, C., & Wang, B. (2020). The impact of salience on investor behavior: Evidence from a natural experiment. The Journal of Finance, 75(1), 229-276.
Giglio, S., Liao, Y., & Xiu, D. (2021). Thousands of Alpha Tests, The Review of Financial Studies, 34(7), 3456–3496.
Gopalan, R., Nanda, V., & Seru, A. (2007). Affiliated firms and financial support: Evidence from Indian business groups. Journal of Financial Economics, 86(3), 759-795.
Griffin, J. M., Kelly, P. J., & Nardari, F. (2010). Do market efficiency measures yield correct inferences? A comparison of developed and emerging markets. The Review of Financial Studies, 23(8), 3225–3277.
Grinblatt, M., & Han, B. (2005). Prospect theory, mental accounting, and momentum. Journal of financial economics, 78(2), 311-339.
Gu, S., Kelly, B., and Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. The Review of Financial Studies, 33(5):2223–2273.
Harvey, C. R., Liu, Y., and Zhu, H. (2015). . . . and the Cross-Section of Expected Returns. The Review of Financial Studies, 29(1):5–68.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778.
Henderson, V. (2012). Prospect theory, liquidation, and the disposition effect. Management Science, 58(2), 445-460.
Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. ArXiv, abs/1207.0580.
Hirshleifer, D. (2015). Behavioral finance. Annual Review of Financial Economics, 7(1):133–159.
Horel, E. and Giesecke, K. (2020). Significance tests for neural networks. J. Mach. Learn. Res., 21(1).
Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. Neural Networks, 2(5):359–366.
Hou, K., & Moskowitz, T. J. (2005). Market frictions, price delay, and the cross-section of expected returns. Review of Financial Studies, 18(3), 981–1027.
Hou, K., Xue, C., and Zhang, L. (2020). Replicating anomalies. Review of Financial Studies, 33(5):2019–2133.
Ivoovic, Z. and Weisbenner, S. (2005). Local does as local is: Information content of the geography of individual investors' common stock investments. The Journal of Finance, 60(1):267–306.
Ivković, Z., & Weisbenner, S. (2009). Individual investor mutual fund flows. Journal of Financial Economics, 92(2), 223-237.
Jensen, T. I., Kelly, B. T., Malamud, S., & Pedersen, L. H. (2022). Machine learning and the implementable efficient frontier. Swiss Finance Institute Research Paper, (22-63).
Jensen, T. I., Kelly, B., and Pedersen, L. H. (2023). Is there a replication crisis in finance? The Journal of Finance, 78(5):2465–2518.
Kacperczyk, M., Van Nieuwerburgh, S., & Veldkamp, L. (2016). A rational theory of mutual funds' attention allocation. Econometrica, 84(2), 571–626.
Kaniel, R., Liu, S., Saar, G., & Titman, S. (2012). Individual investor trading and return patterns around earnings announcements. The Journal of Finance, 67(2), 639–680.
Kaniel, R., Lin, Z., Pelger, M., and Van Nieuwerburgh, S. (2023). Machine-learning the skill of mutual fund managers. Journal of Financial Economics, 150(1):94–138.
Karolyi, G. A., & Van Nieuwerburgh, S. (2020). New methods for the cross-section of returns. The Review of Financial Studies, 33(5), 1879-1890.
Kaustia, M. (2010). Prospect theory and the disposition effect. Journal of Financial and Quantitative Analysis, 45(3), 791-812.
Kelly, B. T., Pruitt, S., & Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. Journal of Financial Economics, 134(3), 501-524.
Kelly, B., & Xiu, D. (2023). Financial machine learning. Foundations and Trends® in Finance, 13(3-4), 205-363.
Kumar, A. (2009). Who gambles in the stock market? The Journal of Finance, 64(4):1889–1933.
Leippold, M., Wang, Q., and Zhou, W. (2022). Machine learning in the chinese stock market. Journal of Financial Economics, 145(2):64–82.
Li, B., & Rossi, A. G. (2020). Selecting mutual funds from the stocks they hold: A machine learning approach. Available at SSRN 3737667.
Li, B., A. Rossi, X. Yan and L. Zheng (2023). Real-Time Machine Learning in the Cross-Section of Stock Returns, working paper.
Li, Y., & Yang, L. (2013). Prospect theory, the disposition effect, and asset prices. Journal of Financial Economics, 107(3), 715-739.
Lilienfeld‐Toal, U. V., Mookherjee, D., & Visaria, S. (2012). The distributive impact of reforms in credit enforcement: Evidence from Indian debt recovery tribunals. Econometrica, 80(2), 497-558.
Liu, H., Peng, C., Xiong, W. A., and Xiong, W. (2022). Taming the bias zoo. Journal of Financial Economics, 143(2):716–741.
Liu, J., Stambaugh, R. F., and Yuan, Y. (2019). Size and value in china. Journal of Financial Economics, 134(1):48–69.
Lo, A. W., & MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. Review of Financial Studies, 1(1), 41–66.
Lopez-Lira, A., & Roussanov, N. L. (2020). Do Common Factors Really Explain the Cross-Section of Stock Returns?. Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.
Lusardi, A. and Mitchell, O. S. (2011). Financial literacy around the world: an overview. Journal of Pension Economics & Finance, 10(4):497–508.

Massa, M. and Simonov, A. (2006). Hedging, familiarity and portfolio choice. Review of Financial Studies, 19(2):633–685.
Mclean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? The Journal of Finance, 71(1):5–32.
Odean, T. (1998). Are investors reluctant to realize their losses? The Journal of Finance, 53(5):1775–1798.
Sadhwani, A., Giesecke, K., and Sirignano, J. (2020). Deep Learning for Mortgage Risk*. Journal of Financial Econometrics, 19(2):313–368.
Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. The Journal of Finance, 40(3):777–790.
Van Binsbergen, J. H., Han, X., & Lopez-Lira, A. (2023). Man versus machine learning: The term structure of earnings expectations and conditional biases. The Review of Financial Studies, 36(6), 2361-2396.
Wu, W., Chen, J., Yang, Z., & Tindall, M. L. (2021). A cross-sectional machine learning approach for hedge fund return prediction and selection. Management Science, 67(7), 4577-4601.

Table 1: Variable Names and Explanations

This table tabulates the list of firm characteristics and behavioral biases that we use in our analysis. In Panel A, the first six categories represent firm characteristics and the last consists of behavioral biases. Panel B summarizes the literature on behavioral biases.

Panel A: Variable Names and Explanations

Name	Explanation	Name	Explanation
Profitability		***Investor Behavioral Bias***	
ROA	Return on assets	Distance	Local Bias
NSOLA	Net Sales Over Lagged Assets	Port_Value	Investor's Portfolio Value
COGS	Cost of Goods Sold over lagged assets	Extrapolation	Extrapolation
SaleGrow	Sales Growth	Disp	Disposition Effect
Past Returns		Month_Diver	Diversification
R1_0	Last month return	Investor_Tvr	Investor Turnover
R2_1	Return from t-2 to t-1	Past_Perform	Investor Past Performance
R12_7	Intermediate momentum	IVOL	Idiosyncratic Volatility (Proxy for Lottery Preference)
R12_2	Momentum	Low_Price	Low Price Rank (Proxy for Lottery Preference)
Investments		High_Price	High Price Rank (Proxy for Lottery Preference)
DPI2A	Change in property, plants, and equipment	Open_Price	Open Price Rank (Proxy for Lottery Preference)
NI	Net Share Issues	Close Price	Close Price Rank (Proxy for Lottery Preference)
Intangibles		Skew	Idiosyncratic Skewness (Proxy for Lottery Preference)
NIA	Net Intangible Asset		
Value			
TobinQ	Tobin's Q		
Div_Yield	Dividend Yield		
EPS	Earnings Per Share		
BVPS	Book Value Per Share		
PE	Price to Earnings		
PB	Price to Book Value		
Trading Frictions			
TA	Total Asset		
Size	Market Equity		
Turnover	Monthly Turnover		
TradVol	Monthly Trading Volume		
Leverage	Financial leverage		
NOE	Number of Employees Growth		

Panel B: Reference of the behavioral bias used in our analysis.

Bias	Proxy	Papers
The disposition effect	Regression coefficient	Shefrin and Statman (1985), Odean (1998), Ben-David and Hirshleifer (2012)
Lottery preference	Ivol Iskew Stock price	Kumar (2009), Harvey and Siddique (2000), Bordalo, Gennaioli, and Shleifer (2012; 2013; 2020)
Extrapolation	Excess return of holding stocks	Barber and Odean (2013)
Underdiversification	Number of stocks in an investor's portfolio	Barber and Odean (2000), Benartzi and Thaler (2001), Lusardi and Mitchell (2011)
Local bias	Average distance between an investor's location and the headquarters of the stocks the investor bought	Ivkovic and Weisbenner (2005), Massa and Simonov (2006)
Turnover	The frequency of trading for the investors	Odean (1998), Barber and Odean (2000)

Table 2. Model Comparison: Predict Total Return

This table reports the performance comparison of various models, including Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN). We first use each model to predict investor returns and then sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. Finally, we report the value-weighted out-of-sample returns of the high and low groups as well as their return difference. we also use the locally estimated three-factor and four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Low Group			High Group			High Minus Low		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
Linear	-0.018*	-0.017	-0.014	0.007	0.007	0.008	0.025***	0.024***	0.022***
	(-1.77)	(-1.63)	(-1.40)	(1.37)	(1.33)	(1.52)	(3.34)	(3.14)	(2.95)
Lasso	-0.008	-0.006	-0.004	0.008	0.008	0.010*	0.016**	0.014*	0.013*
	(-0.75)	(-0.53)	(-0.34)	(1.56)	(1.61)	(1.90)	(2.06)	(1.80)	(1.68)
Ridge	-0.018*	-0.017	-0.014	0.007	0.007	0.008	0.025***	0.024***	0.022***
	(-1.76)	(-1.61)	(-1.38)	(1.36)	(1.32)	(1.52)	(3.32)	(3.10)	(2.91)
Random Forest	-0.015	-0.014	-0.012	0.011	0.009	0.007	0.026	0.023	0.019
	(-1.24)	(-1.17)	(-1.03)	(1.08)	(0.90)	(0.70)	(1.47)	(1.59)	(1.63)
FNN	-0.025*	-0.018	-0.015	0.015***	0.015***	0.017***	0.040***	0.033***	0.031***
	(-1.74)	(-1.62)	(-1.43)	(2.66)	(2.70)	(3.08)	(3.38)	(3.24)	(3.11)
Residual Neural Network	-0.031**	-0.029**	-0.026**	0.012**	0.013**	0.014**	0.044***	0.042***	0.041***
	(-2.38)	(-2.17)	(-2.00)	(2.00)	(2.08)	(2.33)	(4.57)	(4.38)	(4.26)

Table 3. Information Set Comparison: Holding-Based Characteristics vs. Investor Behavioral Biases

This table reports the performance of Neural Network algorithms when different subsets of predictors (i.e., firm characteristics or investor behavioral biases) are used alone or jointly to predict returns. The first (second) line uses the standard Feedforward Neural Network (FNN) to predict investor returns based on firm characteristics (behavioral biases) only. The third line allows FNN to use both firm characteristics and investor behavioral biases to predict returns. The last line utilizes the Residual Neural Network (ResNN) using both sets of predictors. For each model-predictor combination (e.g., FNN using only firm characteristics), we sort retail investors into five quintiles according to their predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among investors, respectively. Finally, we report the value-weighted out-of-sample returns of the high and low groups as well as their return difference. We also use the locally estimated three-factor and four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Low Group			High Group			High Minus Low		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
Stock Characteristics	-0.003	-0.001	0.001	0.006	0.006	0.008	0.009	0.007	0.007
	(-0.23)	(-0.08)	(0.09)	(1.01)	(1.02)	(1.32)	(0.99)	(0.79)	(0.72)
Behavioral Biases	-0.025**	-0.024**	-0.022**	0.008	0.009	0.010*	0.033***	0.033***	0.032***
	(-2.62)	(-2.43)	(-2.25)	(1.46)	(1.55)	(1.80)	(5.63)	(5.41)	(5.29)
Stock Chars + Behavioral	-0.025*	-0.018	-0.015	0.015***	0.015***	0.017***	0.040***	0.033***	0.031***
	(-1.74)	(-1.62)	(-1.43)	(2.66)	(2.70)	(3.08)	(3.38)	(3.24)	(3.11)
Residual Neural Network	-0.031**	-0.029**	-0.026**	0.012**	0.013**	0.014**	0.044***	0.042***	0.041***
	(-2.38)	(-2.17)	(-2.00)	(2.00)	(2.08)	(2.33)	(4.57)	(4.38)	(4.26)

Table 4. Predictor Comparison when Predicting Holding and Trading-based Returns

This table reports the performance of Neural Network algorithms when different subsets of predictors (i.e., firm characteristics or investor behavioral biases) are used alone or jointly to predict returns. In each panel, the training objectives are to predict the two components of the total returns received by retail investors: trading returns (columns 1-3) and holding returns (columns 4-6). The first (second) line uses the standard Feedforward Neural Network (FNN) to predict investor returns based on firm characteristics (behavioral biases) only. The third line allows FNN to use both firm characteristics and investor behavioral biases to predict returns. The last line utilizes the Residual Neural Network (ResNN) using both sets of predictors. For each model-predictor-training goal combination (e.g., FNN using only firm characteristics to predict trading returns), we sort retail investors into five quintiles according to their predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among investors, respectively. Panel A reports the value-weighted out-of-sample *total returns* of the high and low groups as well as their return difference. Panels B and C further report the risk-adjusted total returns based on the locally estimated three-factor and four-factor models. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Excess Return

	(1)	(2)	(3)	(4)	(5)	(6)
	Trading Return			Holding Return		
	Low	High	High Minus Low	Low	High	High Minus Low
Stock Characteristics	-0.000	-0.006	-0.006	-0.005	0.012**	0.017*
	(-0.05)	(-0.71)	(-1.11)	(-0.38)	(1.99)	(1.96)
Behavioral Biases	-0.031***	0.003	0.034***	-0.022**	0.010*	0.031***
	(-4.01)	(0.42)	(17.75)	(-2.00)	(1.69)	(4.32)
Stock Chars + Behavioral	-0.034***	0.003	0.037***	-0.019	0.012**	0.031**
	(-4.10)	(0.43)	(13.81)	(-1.16)	(2.18)	(2.35)
Residual Neural Network	-0.028***	0.003	0.031***	-0.023	0.012**	0.035***
	(-3.67)	(0.39)	(14.69)	(-1.61)	(2.25)	(3.04)

Panel B: Fama-French Three Factor Adjusted Alpha

	(1)	(2)	(3)	(4)	(5)	(6)
	Trading Return			Holding Return		
	Low	High	High Minus Low	Low	High	High Minus Low
Stock Characteristics	0.001	-0.005	-0.006	-0.003	0.013**	0.016*
	(0.09)	(-0.60)	(-1.13)	(-0.26)	(2.12)	(1.85)
Behavioral Biases	-0.029***	0.004	0.034***	-0.020*	0.010*	0.031***
	(-3.73)	(0.58)	(17.00)	(-1.82)	(1.79)	(4.13)
Stock Chars + Behavioral	-0.032***	0.004	0.036***	-0.019	0.013**	0.032**
	(-3.82)	(0.59)	(13.23)	(-1.12)	(2.27)	(2.34)
Residual Neural Network	-0.027***	0.004	0.031***	-0.021	0.013**	0.034***
	(-3.40)	(0.56)	(14.20)	(-1.48)	(2.33)	(2.90)

Panel C: Carhart Four Factor Adjusted Alpha

	(1)	(2)	(3)	(4)	(5)	(6)
	Trading Return			Holding Return		
	Low	High	High Minus Low	Low	High	High Minus Low
Stock Characteristics	0.003	-0.004	-0.006	0.000	0.015**	0.015*
	(0.35)	(-0.41)	(-1.17)	(0.03)	(2.57)	(1.68)
Behavioral Biases	-0.028***	0.006	0.033***	-0.018	0.012**	0.030***
	(-3.55)	(0.84)	(16.79)	(-1.65)	(2.09)	(4.04)
Stock Chars + Behavioral	-0.031***	0.006	0.036***	-0.015	0.014**	0.029**
	(-3.65)	(0.84)	(13.12)	(-0.90)	(2.62)	(2.17)
Residual Neural Network	-0.025***	0.006	0.031***	-0.020	0.014***	0.034***
	(-3.22)	(0.82)	(14.07)	(-1.35)	(2.74)	(2.87)

Table 5: Market Efficiency Test

We follow Brogaard, Nguyen, Putnins, and Wu (2023) to decompose the return variance of individual stocks into four components: noise, private firm-specific information revealed through trading, firm-specific information revealed through public sources and market-wide information. Based on the predicted returns from the Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN), investors are divided into 5 groups. We calculate the group level trading intensity each month, which is defined as the total trading volume of each group of investors in each stock, scaled by the total trading volume of that stock. We employ a panel regression framework to estimate the model. We include both firm fixed effects and year fixed effects. Standard errors are clustered at both the firm and year levels. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Regression Results for Noise Information Share

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Linear	Lasso	Ridge	Random Forest	FNN	ResNN
Winner	0.005	0.017	-0.005	-0.008*	-0.020***	-0.019***
	(0.95)	(1.46)	(-1.53)	(-1.89)	(-4.20)	(-4.34)
Loser	0.013***	0.017***	0.015***	0.021***	0.019***	0.020***
	(3.15)	(3.67)	(4.26)	(5.85)	(5.68)	(5.81)
PE	0.002	0.002	0.002	0.002	0.003	0.005**
	(1.56)	(1.62)	(1.69)	(1.49)	(1.61)	(2.58)
Size	-0.016***	-0.018***	-0.018***	-0.015***	-0.015***	-0.016***
	(-4.00)	(-3.99)	(-4.09)	(-4.21)	(-3.94)	(-5.20)
Constant	0.495***	0.498***	0.498***	0.503***	0.498***	0.497***
	(272.52)	(249. 46)	(274.46)	(247.78)	(207.85)	(236.88)
Time FE	YES	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES	YES

Panel B: Regression Results for Private Information Share

VARIABLES	(1) Linear	(2) Lasso	(3) Ridge	(4) Random Forest	(5) FNN	(6) ResNN
Winner	0.024*	0.029	0.026	0.046**	0.062***	0.064***
	(1.86)	(1.23)	(1.41)	(2.48)	(5.16)	(5.67)
Loser	0.013	0.016	0.008	0.004	-0.014	-0.014
	(0.83)	(1.11)	(0.61)	(0.33)	(-1.53)	(-1.53)
PE	0.003**	0.003**	0.003**	0.003**	0.003**	0.005**
	(2.51)	(2.63)	(2.64)	(2.55)	(2.55)	(2.58)
Size	-0.035***	-0.035***	-0.034***	-0.036***	-0.036***	-0.034***
	(-6.74)	(-6.66)	(-6.35)	(-6.90)	(-6.90)	(-6.20)
Constant	0.168***	0.167***	0.164***	0.167***	0.167***	0.165***
	(165.75)	(133.95)	(200.82)	(184.43)	(184.30)	(236.88)
Time FE	YES	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES	YES

Table 6. The Impact of Market Conditions

This table reports the results when the High-minus-Low spread estimated by neural network models are regressed against the local factors, the volatility and skewness of the local market, as well as investor sentiment. Sentiment is measured as the buy-minus-sell orders scaled by the buy-plus-sell orders of all investors in our sample. The left and right four columns use FNN- and Res-NN generated High-minus-Low return spread as the dependent variable. Columns (1) and (5) tabulate the time series regression when the four factors are included. Columns (2) – (4) and columns (6) – (8) tabulate the results when these three market condition variables are included. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

VARIABLES	(1) FNN	(2) FNN	(3) FNN	(4) FNN	(5) ResNN	(6) ResNN	(7) ResNN	(8) ResNN
Constant	0.031***	0.020	0.021***	0.021***	0.041***	0.037**	0.032***	0.033***
	(3.11)	(0.96)	(2.69)	(2.86)	(4.26)	(2.02)	(3.91)	(4.24)
HML	-0.897***	-0.960***	-0.995***	-0.939***	-0.679***	-0.706***	-0.719***	-0.707***
	(-3.84)	(-4.06)	(-4.17)	(-4.07)	(-3.51)	(-3.57)	(-3.58)	(-3.57)
SMB	-0.006	-0.007	-0.007	-0.008*	-0.010**	-0.010**	-0.010**	-0.010**
	(-1.34)	(-1.52)	(-1.60)	(-1.67)	(-2.53)	(-2.59)	(-2.61)	(-2.57)
MKT	0.574	0.741	0.801	0.742	0.877	0.951	0.967	0.945
	(0.83)	(1.05)	(1.14)	(1.10)	(1.53)	(1.62)	(1.64)	(1.61)
MOM	0.922***	0.911***	0.919***	0.931***	0.456***	0.442**	0.453***	0.451***
	(4.63)	(4.54)	(4.63)	(4.85)	(2.77)	(2.60)	(2.73)	(2.71)
MKT_VOL		0.199				-0.434		
		(0.12)				(-0.28)		
MKT_SKEW			-0.010				-0.004	
			(-0.84)				(-0.38)	
Sentiment				-1.316**				0.130
				(-2.18)				(0.25)
Observations	95	95	95	95	95	95	95	95
R-squared	0.429	0.435	0.440	0.468	0.300	0.304	0.304	0.304

Table 7: Performance Persistence

This table reports the performance comparison of FNN and ResNN models on different time horizons. We first use each model to predict investor returns and then sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. Finally, we report the value-weighted out-of-sample returns of the high and low groups on next 12 months respectively. We also use the locally estimated three-factor and four-factor models to adjust these returns. Model (1) corresponds to the results of our main analysis. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: The performance of High-minus-Low Return Spread selected by FNN

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	T+1	T+2	T+3	T+4	T+5	T+6	T+7	T+8	T+9	T+10	T+11	T+12
Mean	0.040***	0.012**	0.011*	0.011*	0.010*	0.010	0.005	0.004	0.003	0.001	0.001	0.001
	(3.38)	(2.13)	(1.74)	(1.69)	(1.68)	(1.59)	(1.33)	(1.23)	(1.02)	(0.90)	(0.61)	(0.32)
FF-3	0.033***	0.012**	0.009*	0.010*	0.010*	0.008	0.004	0.003	0.002	0.001	0.000	-0.001
	(3.24)	(2.17)	(1.70)	(1.68)	(1.69)	(1.53)	(1.31)	(1.20)	(0.85)	(0.64)	(0.47)	(-0.15)
Carhart-4	0.031***	0.013**	0.009	0.009	0.008	0.006	0.004	0.001	0.001	0.000	-0.003	0.000
	(3.11)	(2.21)	(1.66)	(1.59)	(1.60)	(1.47)	(1.27)	(1.01)	(0.60)	(0.35)	(-0.24)	(0.06)

Panel B: The performance of High-minus-Low Return Spread selected by Res-NN

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	T+1	T+2	T+3	T+4	T+5	T+6	T+7	T+8	T+9	T+10	T+11	T+12
Mean	0.044***	0.021***	0.016**	0.014*	0.011*	0.010*	0.008	0.006	0.005	0.004	0.002	0.001
	(4.57)	(2.98)	(2.08)	(1.95)	(1.79)	(1.65)	(1.54)	(1.48)	(1.25)	(1.12)	(0.83)	(0.51)
FF-3	0.042***	0.020***	0.014**	0.012*	0.011*	0.009	0.006	0.006	0.004	0.002	0.000	-0.000
	(4.38)	(2.76)	(2.00)	(1.67)	(1.64)	(1.60)	(1.45)	(1.37)	(1.02)	(0.84)	(0.67)	(-0.01)
Carhart-4	0.041***	0.020***	0.012*	0.011	0.010	0.008	0.006	0.004	0.002	0.000	-0.000	0.000
	(4.26)	(2.74)	(1.96)	(1.60)	(1.55)	(1.53)	(1.41)	(1.18)	(0.60)	(0.39)	(-0.04)	(0.09)

Table 8. Expanding to Fifty Firm Characteristics

Panel A reports the performance comparison of various models, including Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN), when we construct a list of 50 firm characteristics. We first use each model to predict investor returns and then sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. Finally, we calculate the value-weighted out-of-sample returns of the high and low groups and report their return difference. We also use the locally estimated three-factor and four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. Panel B reports the relative importance of behavioral biases and firm characteristics based on the variable gradient analysis. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Model Comparison with Fifty Firm Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Low Group			High Group			High Minus Low		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
Linear	-0.019*	-0.018*	-0.015	0.009	0.008	0.008	0.028***	0.026***	0.023***
	(-1.76)	(-1.81)	(-1.52)	(1.43)	(1.41)	(1.61)	(3.51)	(3.35)	(3.06)
Lasso	-0.018*	-0.016	-0.014	0.010	0.010*	0.011*	0.028***	0.024***	0.022***
	(-1.85)	(-1.58)	(-1.34)	(1.54)	(1.71)	(1.79)	(3.46)	(3.40)	(3.28)
Ridge	-0.020*	-0.019*	-0.017	0.010	0.010	0.009	0.030***	0.029***	0.026***
	(-1.89)	(-1.81)	(-1.66)	(1.61)	(1.69)	(1.58)	(3.61)	(3.57)	(3.20)
Random Forest	-0.018*	-0.016	-0.016	0.013	0.011	0.010	0.031***	0.027***	0.026***
	(-1.74)	(-1.67)	(-1.53)	(1.27)	(1.09)	(1.07)	(2.97)	(2.63)	(2.69)
FNN	-0.027*	-0.022*	-0.019	0.016***	0.018***	0.016***	0.042***	0.040***	0.035***
	(-1.94)	(-1.91)	(-1.63)	(2.71)	(2.64)	(3.11)	(3.57)	(3.54)	(3.32)
Residual Neural Network	-0.034***	-0.030**	-0.028**	0.014**	0.014**	0.015**	0.048***	0.044***	0.043***
	(-2.42)	(-2.22)	(-2.05)	(2.13)	(2.20)	(2.28)	(4.65)	(4.51)	(4.30)

Panel B: Relative Importance of Behavioral Biases vs. Firm Characteristics

	(1)	(2)	(3)
VARIABLES	Total Return	Holding Return	Trading Return
Stock Characteristics	43.40%	54.80%	6.50%
Investor Behavioral Bias	56.60%	46.20%	93.50%

Table 9. Robustness Checks on Removing 20% or 40% of Small Stocks

This table reports the performance comparison of various models, including Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN). Different from our main analysis, we remove the bottom 40% of small stocks in this robustness check. We first use each model to predict investor returns and then sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. Finally, we report the value-weighted out-of-sample returns of the high and low groups as well as their return difference. we also use the locally estimated three-factor and four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	High Minus Low (Removing 20%)			High Minus Low (Removing 40%)		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
Linear	0.014	0.012	0.013	0.013	0.021	0.018
	(1.71)	(1.61)	(1.57)	(1.52)	(1.21)	(2.00)
Lasso	0.022*	0.017*	0.015	0.022*	0.024*	0.026
	(1.87)	(1.78)	(1.22)	(1.87)	(1.78)	(1.22)
Ridge	0.023**	0.019*	0.017*	0.021**	0.021**	0.020**
	(2.27)	(2.00)	(1.88)	(2.57)	(2.27)	(2.22)
Random Forest	0.023	0.017	0.016	0.024	0.012	0.012
	(1.15)	(1.04)	(1.00)	(1.02)	(1.31)	(1.15)
FNN	0.034***	0.032***	0.029***	0.034***	0.032***	0.028***
	(3.67)	(3.53)	(3.18)	(3.86)	(3.47)	(3.12)
Residual Neural Network	0.045***	0.040***	0.038***	0.044***	0.042***	0.041***
	(4.59)	(4.48)	(4.32)	(4.78)	(4.66)	(4.37)

Table 10. Model Comparisons Across Models

This table reports the difference in performance across various models, including Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN). In conducting the comparison, we remove 20%, 30%, and 40% of small stocks. Each model is used to predict investor returns and sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. The value-weighted out-of-sample High-minus-Low return spread is calculated for each model. The first four rows report the difference between the High-minus-Low return spread of FNN and that of Linear, Lasso, Ridge, and Random Forest models. The last row reports the difference between ResNN and FNN. we also use the locally estimated three-factor and four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	High Minus Low (Excluding 20% Small Stocks)			High Minus Low (Excluding 30% Small Stocks)			High Minus Low (Excluding 40% Small Stocks)		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
FNN - Linear	0.020***	0.020***	0.016***	0.015**	0.015***	0.016***	0.021***	0.020***	0.016***
	(2.98)	(2.86)	(2.80)	(2.37)	(2.77)	(2.84)	(3.03)	(2.94)	(2.89)
FNN - Lasso	0.012**	0.015***	0.014***	0.024***	0.019***	0.017***	0.012**	0.011**	0.008**
	(2.43)	(2.97)	(2.82)	(2.95)	(3.11)	(3.17)	(2.37)	(2.28)	(2.22)
FNN - Ridge	0.011**	0.013***	0.012**	0.015**	0.009**	0.010**	0.013**	0.011**	0.008**
	(2.37)	(2.80)	(2.48)	(2.37)	(2.50)	(2.41)	(2.43)	(2.38)	(2.27)
FNN - Random Forest	0.011**	0.015***	0.013***	0.014**	0.010**	0.013***	0.010**	0.020***	0.016***
	(2.24)	(2.98)	(2.85)	(2.33)	(2.48)	(2.85)	(2.02)	(2.91)	(2.83)
ResNN - FNN	0.011**	0.008**	0.009***	0.004*	0.009**	0.010***	0.010**	0.010**	0.013***
	(2.17)	(2.32)	(2.92)	(1.79)	(2.26)	(2.83)	(2.08)	(2.25)	(2.88)

Figure 1: The Cumulative Returns of High-minus-Low Return Spread from Various Models

This figure plots the cumulative returns generated by the High-minus-Low investor groups predicted by various models, including Linear, Lasso, Ridge, Random Forest, the standard Neural Network algorithms (FNN), and the enhanced Residual Neural Network (ResNN). We first use each model to predict investor returns and then sort retail investors into five quintiles according to predicted returns, with the High and Low groups consisting of 20% of predicted winners and losers among all investors, respectively. We next calculate the value-weighted out-of-sample returns of the high and low groups. Finally, we plot the cumulative returns of the High-minus-Low spread for the period from Jan 2012 to June 2020.

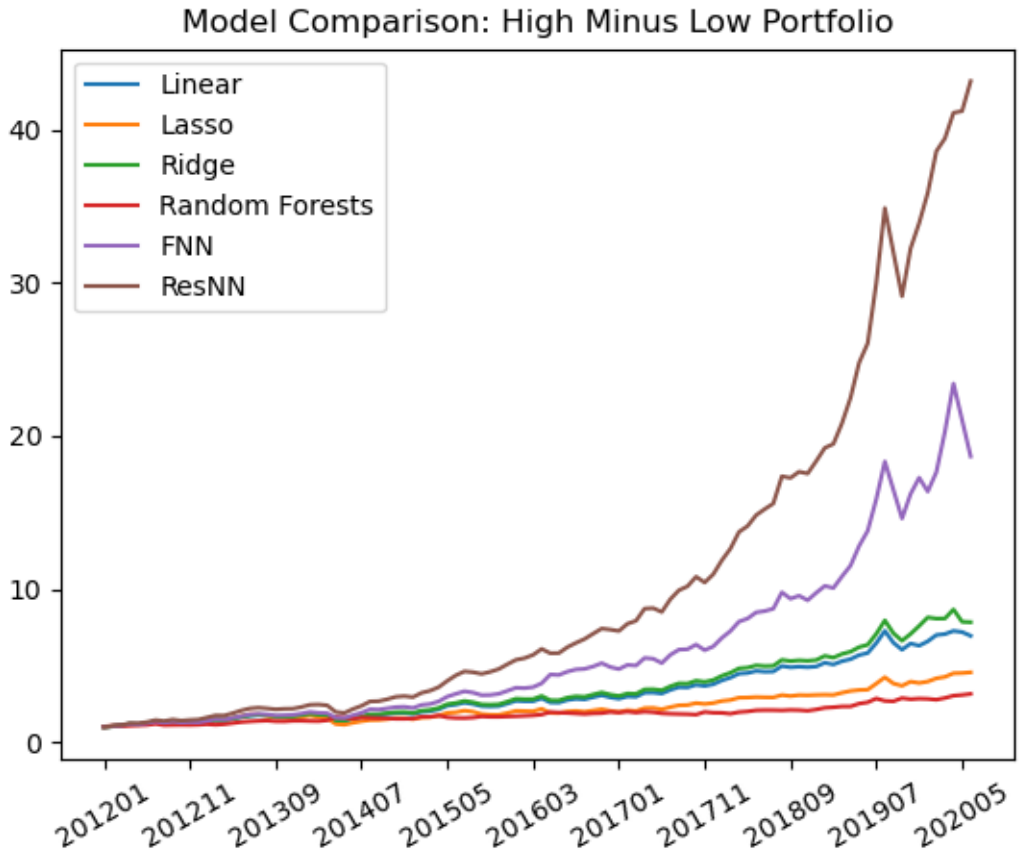


Figure 2: The Relative Importance of Behavioral Bias vs. Firm Characteristics

This figure plots the relative importance of behavioral biases vis-à-vis firm characteristics. Based on our earlier delineation, where predictors are categorized into Investor Behavioral Bias and Firm Characteristics, we define the variable importance measure of a group by computing the average of the importance measures within that group, which can also be expressed as the joint explanatory power of all predictors falling into each category. Without loss of generality, we normalize the variable importance to sum up to 1.

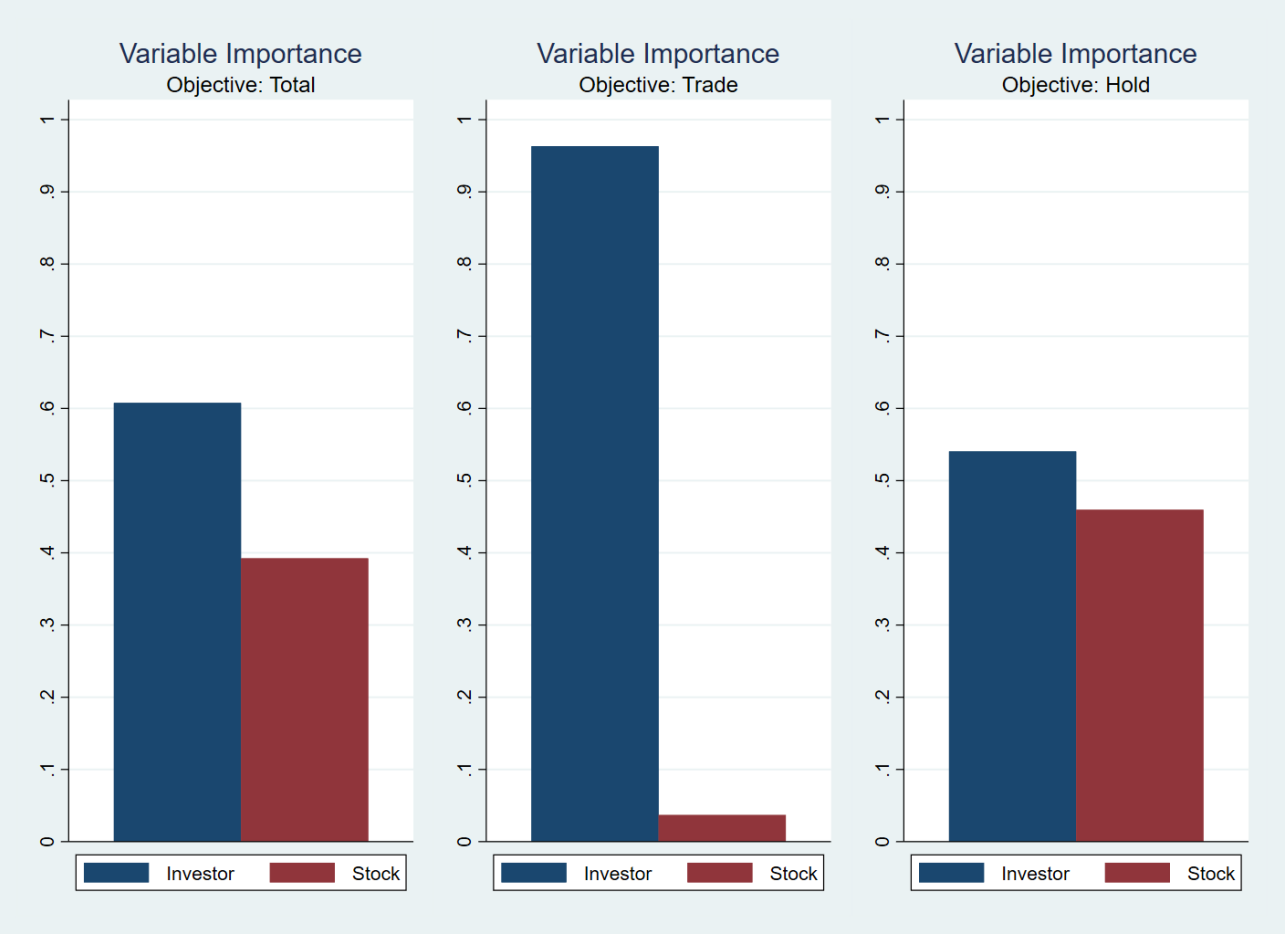


Figure 3: Top Variable Importance of Behavioral Bias vs. Stock Characteristics

This table reports the importance of each predictor when the Feedforward Neural Network (FNN) is used to predict retail investors' returns based on all predictors, including behavioral biases and stock characteristics. The marginal importance of a variable is computed from the variable gradient analysis: $Importance(x) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{i=1}^{N_t}\left(\frac{\partial R_{i,t+1}^{pred}}{\partial x_{i,t}}\right)^2$, where T represents the number of periods in the data, and N_t denotes the total number of investors in the t-th period. The partial derivative measures the gradient of the model's predicted output with respect to each variable. Intuitively, a larger partial derivative implies a greater influence of a variable on the model's output, indicating greater importance in predicting future returns. We computed the Importance(x) separately for the training objectives of total return, trading return, and holding return, as described in Table 4.

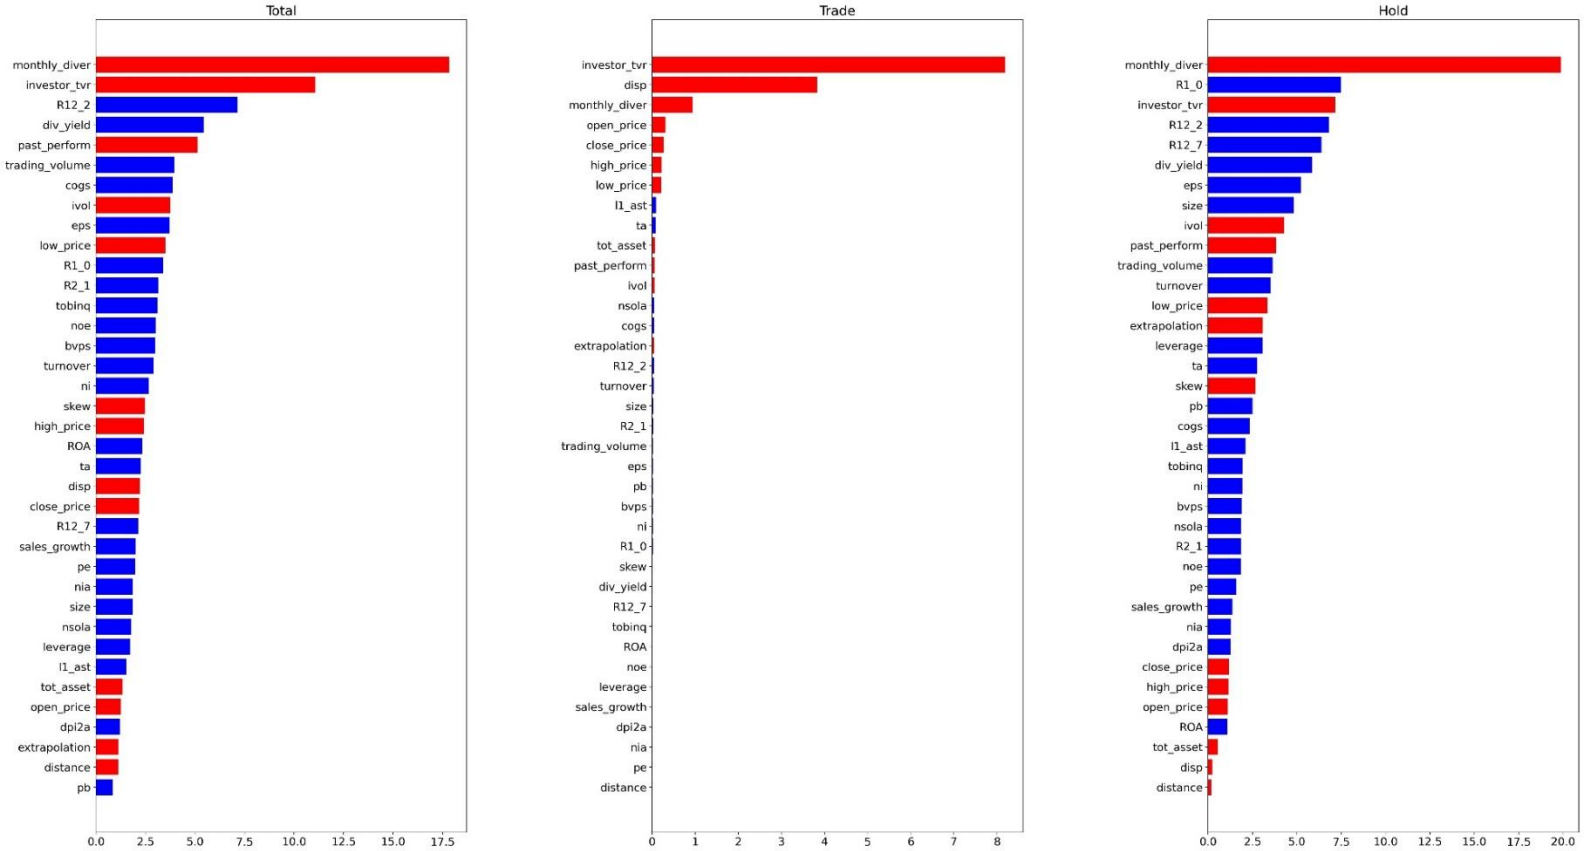


Figure 4: Cumulative Distribution of Relative Importance of Anomalies

This figure reports the cumulative distribution of the relative importance of the third-nine behavioral and characteristics-based anomalies. Panels A, B, and C plot the cumulative distribution when the relative importance of anomalies is determined when predicting total returns, trading returns, and holding returns, respectively. In each plot, the total summation of variable importance is normalized to be 1. The sequence of anomalies from left to right is determined by their relative importance (i.e., the first anomaly is the most important one in predicting returns). The quantile values of the top 3, 5, 10, and 20 anomalies are ploted as dotted lines.

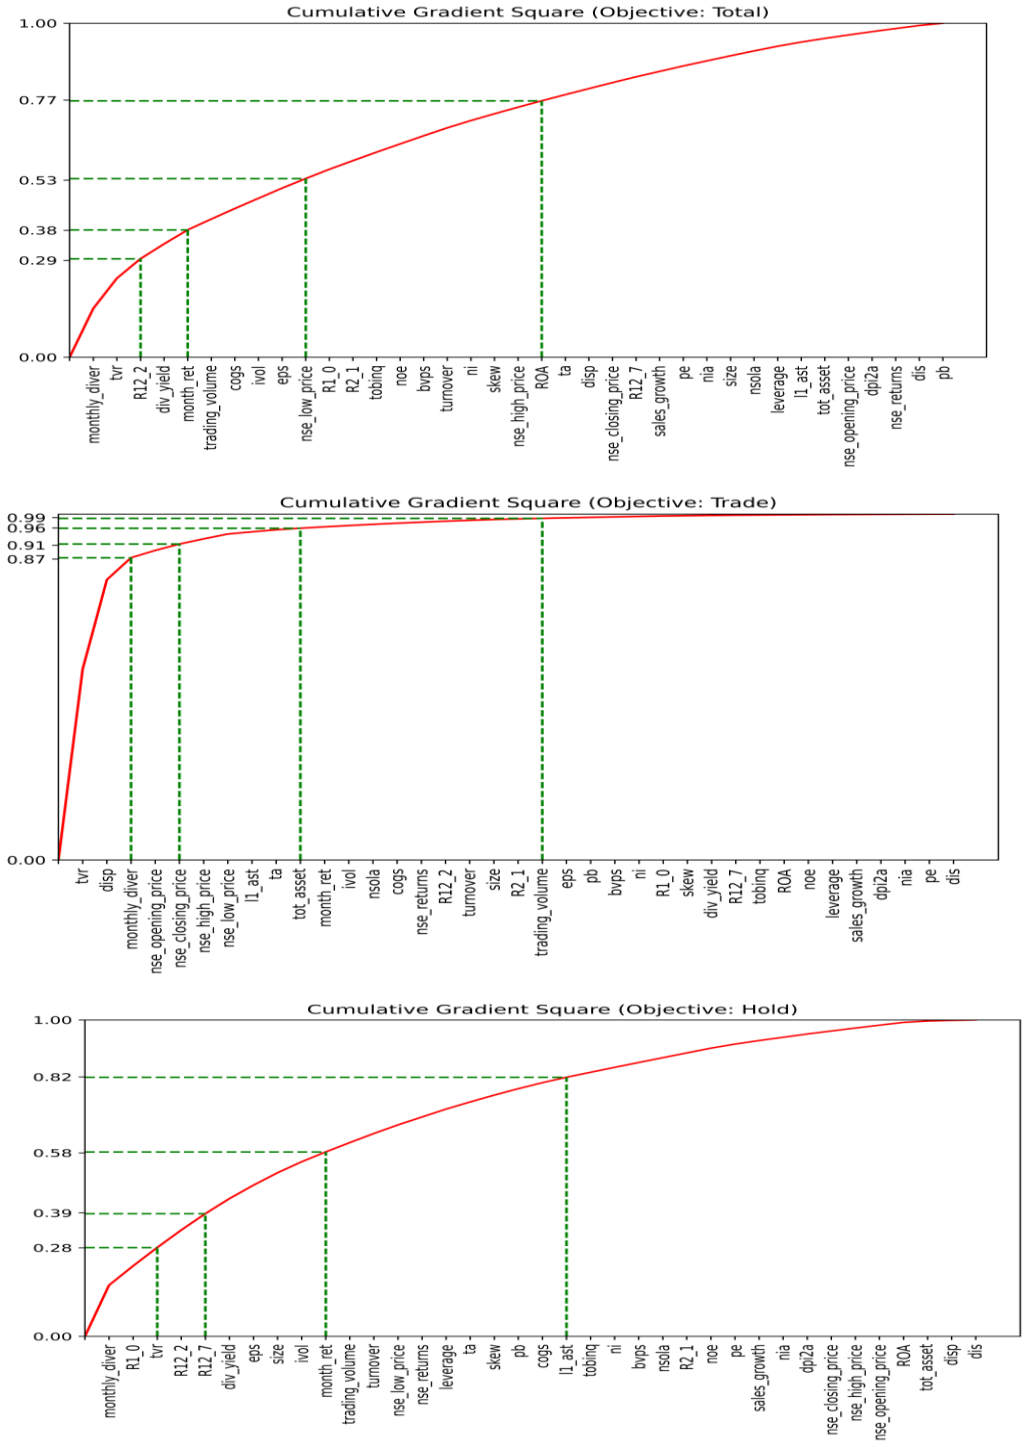
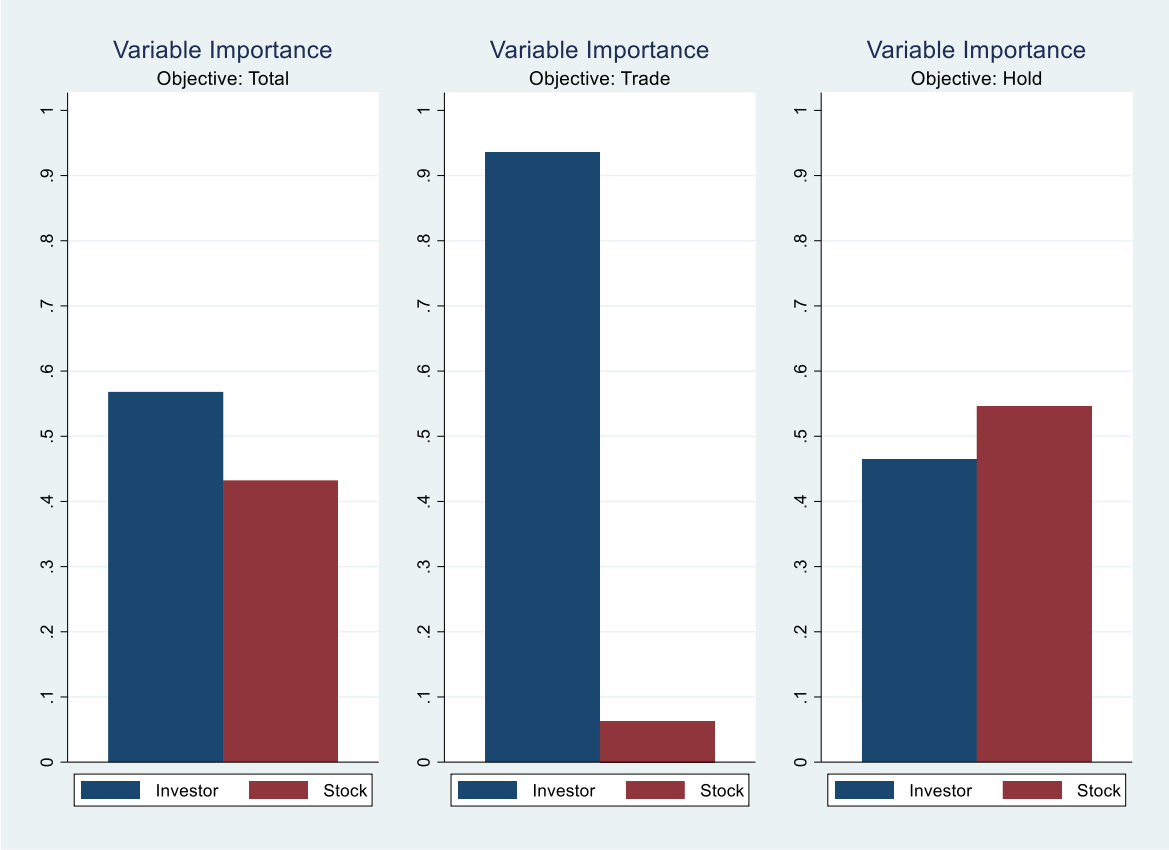


Figure 5: The Relative Importance of Behavioral Bias vs. Firm Characteristics (Expanded Firm Characteristics)

This figure plots the relative importance of behavioral biases vis-à-vis firm characteristics. Based on our earlier delineation, where predictors are categorized into Investor Behavioral Bias and Firm Characteristics, we define the variable importance measure of a group by computing the average of the importance measures within that group, which can also be expressed as the joint explanatory power of all predictors falling into each category. Without loss of generality, we normalize the variable importance to sum up to 1.



Online Appendix

Table IN1: Summary Statistics of Main Variables.

Panel A: Holding based firm characteristics

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Dpi2a	1.523e+09	-0.324	0.750	-1.342	-0.893	-0.339	0.185	0.684
R1_0	1.523e+09	0.0630	0.719	-0.930	-0.417	0.0980	0.566	0.994
R2_1	1.523e+09	0.0460	0.722	-0.949	-0.441	0.0780	0.553	0.984
R12_7	1.523e+09	0.0330	0.738	-0.987	-0.473	0.0560	0.568	0.994
R12_2	1.523e+09	0.0230	0.747	-1.019	-0.496	0.0450	0.575	0.999
Net Issue	1.523e+09	1.192	0.533	0.467	0.974	1.365	1.582	1.671
Nsola	1.523e+09	-0.371	0.742	-1.376	-0.952	-0.394	0.134	0.615
Cogs	1.523e+09	-0.404	0.738	-1.428	-0.966	-0.423	0.0710	0.580
ROA	1.523e+09	0.0530	0.702	-0.839	-0.441	0.0500	0.526	0.995
Sales Growth	1.523e+09	0.124	0.711	-0.870	-0.301	0.159	0.578	1.050
Nia	1.523e+09	0.725	0.744	-0.336	0.321	0.861	1.318	1.559
PB	1.523e+09	0.171	0.733	-0.856	-0.330	0.241	0.709	1.079
PE	1.523e+09	0.120	0.663	-0.760	-0.294	0.132	0.536	0.989
BVPS	1.523e+09	0.323	0.739	-0.746	-0.158	0.419	0.882	1.213
EPS	1.523e+09	0.253	0.802	-0.960	-0.295	0.367	0.867	1.239
Leverage	1.523e+09	0.0250	0.719	-0.917	-0.456	0.00400	0.533	1.030
TobinQ	1.523e+09	0.163	0.725	-0.840	-0.354	0.198	0.695	1.090
Div Yield	1.523e+09	0.497	0.771	-0.651	0.0500	0.640	1.085	1.404
TA	1.523e+09	1.190	0.516	0.471	0.985	1.363	1.557	1.650
Size	1.523e+09	1.141	0.514	0.411	0.898	1.283	1.529	1.657
Turnover	1.523e+09	1.192	0.506	0.498	0.973	1.349	1.563	1.669
Trading Vol	1.523e+09	1.227	0.499	0.560	1.025	1.389	1.585	1.677
NOE	1.523e+09	0.839	0.734	-0.198	0.442	1.012	1.418	1.625

Appendix Table IN1: Summary Statistics of Main Variables.

Panel B: Investor behavioral biases

Variable	N	Mean	SD	p10	p25	p50	p75	p90
Port. Value	1.523e+09	1.187	0.521	0.462	0.979	1.363	1.560	1.651
Diver	1.523e+09	9.669	17.013	1.000	2.000	5.000	11.000	22.000
Disp	1.523e+09	0.001	0.008	-0.004	-0.001	0.002	0.003	0.004
ivol	1.523e+09	-0.0570	1.018	-1.416	-0.944	-0.157	0.877	1.360
iskew	1.523e+09	-0.0150	1.031	-1.408	-0.924	-0.126	0.915	1.399
Distance	1.523e+09	877.144	498.822	258.547	528.240	858.926	1185.493	1484.389
Investor Tvr	1.523e+09	-0.128	0.909	-0.662	-0.627	-0.589	-0.534	1.498
Open Price	1.523e+09	0.317	0.773	-0.852	-0.181	0.426	0.910	1.258
High Price	1.523e+09	0.315	0.773	-0.854	-0.184	0.424	0.907	1.257
Low Price	1.523e+09	0.318	0.773	-0.849	-0.180	0.428	0.911	1.259
Close Price	1.523e+09	0.317	0.773	-0.851	-0.181	0.427	0.910	1.258
Extrapolat~n	1.523e+09	0.0380	0.174	-0.155	-0.0510	0.0430	0.135	0.228
Past Perform	1.523e+09	1.009	0.124	0.897	0.947	1.004	1.063	1.130

Table IN2: Variance Ratio Test

Based on the predicted returns from the Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN), investors are divided into 5 groups. We calculate the group level trading intensity each day, which is defined as the total trading volume of each group of investors in each stock, scaled by the total trading volume of that stock. We follow Lo and MacKinlay (1988) and Campbell, Lo, and MacKinlay (1997) in the construction of variance ratio (VR) statistics and of their heteroscedasticity consistent significance tests. To improve the small sample properties of the tests, in all reported VR analysis, we use overlapping observations. We first estimate autocorrelation coefficients of order five and ten for the daily return series for individual stocks and then calculate the absolute deviation of the variance ratios from one (|VR-1|) as VR statistics. We employ a panel regression framework to estimate the model. We include both firm fixed effects and date (or time) fixed effects. Standard errors are clustered at both the firm and date levels. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Regression Results for Variance Ratio (5 days)

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Linear	Lasso	Ridge	Random Forest	FNN	ResNN
Winner	-0.003	0.003	-0.003	-0.008*	-0.011**	-0.012***
	(-0.65)	(0.44)	(-0.67)	(-1.69)	(-2.32)	(-2.85)
Loser	-0.004	0.010	-0.004	-0.005	-0.009	-0.004
	(-0.52)	(1.14)	(-0.51)	(-1.11)	(-1.31)	(-0.81)
PE	0.005*	0.005*	0.005*	0.004*	0.006*	0.006*
	(1.66)	(1.66)	(1.66)	(1.66)	(1.78)	(1.72)
Size	-0.096***	-0.096***	-0.096***	-0.096***	-0.094***	-0.094***
	(-11.53)	(-11.53)	(-11.53)	(-11.52)	(-10.68)	(-10.71)
Constant	0.248***	0.248***	0.248***	0.248***	0.242***	0.243***
	(111.79)	(111.50)	(111.80)	(111.63)	(120.37)	(120.17)
Time FE	YES	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES	YES

Panel B: Regression Results for Variance Ratio (10 days)

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Linear	Lasso	Ridge	Random Forest	FNN	ResNN
Winner	-0.007	0.006	-0.007	-0.010*	-0.019***	-0.021***
	(-1.21)	(0.77)	(-1.15)	(-1.76)	(-3.26)	(-3.43)
Loser	-0.001	0.013	-0.001	-0.001	-0.004	0.002
	(-0.13)	(1.12)	(-0.19)	(-0.85)	(-0.31)	(0.17)
PE	0.006	0.006	0.006	0.005	0.006	0.006
	(1.45)	(1.46)	(1.45)	(1.38)	(1.41)	(1.44)
Size	-0.123***	-0.123***	-0.123***	-0.121***	-0.122***	-0.122***
	(-9.68)	(-9.68)	(-9.68)	(-9.52)	(-9.66)	(-9.66)
Constant	0.326***	0.326***	0.326***	0.327***	0.327***	0.326***
	(106.76)	(106.35)	(106.77)	(106.57)	(106.15)	(106.02)
Time FE	YES	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES	YES

Table IN3: Market Delay Test

Based on the predicted returns from the Linear, Lasso, Ridge, Random Forest, the standard Feedforward Neural Network (FNN), and the enhanced Residual Neural Network (ResNN), investors are divided into 5 groups. We calculate the group level trading intensity each day, which is defined as the total trading volume of each group of investors in each stock, scaled by the total trading volume of that stock. For each individual stock we estimate the restricted and the unrestricted models over the entire sample period. The unrestricted model uses four lags of daily market returns and is

$$r_{i,t} = \alpha_i + \beta_{0,i} r_{m,t} + \beta_{1,i} r_{m,t-1} + \beta_{2,i} r_{m,t-2} + \beta_{3,i} r_{m,t-3} + \beta_{4,i} r_{m,t-4} + \epsilon_{i,t}$$

Where $r_{i,t}$ is the daily individual stock return at time t and $r_{m,t}$ is the local market return. The restricted model constrains the coefficients on the lagged market returns to zero:

$$r_{i,t} = \alpha_i + \beta_{0,i} r_{m,t} + \epsilon_{i,t}$$

The adjusted R2s from these regressions are used to calculate delay as follows:

$$Delay = Adj.R^2_{unrestricted} - Adj.R^2_{restricted}$$

We employ a panel regression framework to estimate the model. We include both firm fixed effects and date (or time) fixed effects. Standard errors are clustered at both the firm and date levels. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Linear	Lasso	Ridge	Random Forest	FNN	ResNN
Winner	0.000	-0.003	0.000	-0.004*	-0.007**	-0.011***
	(0.12)	(-0.83)	(0.04)	(-1.96)	(-2.13)	(-3.08)
Loser	-0.002	-0.004	-0.001	0.006	0.000	0.007
	(-0.17)	(-1.22)	(-0.31)	(1.53)	(0.21)	(1.35)
PE	0.007***	0.007***	0.007***	0.007***	0.007***	0.007***
	(7.65)	(7.65)	(7.65)	(7.47)	(7.55)	(7.54)
Size	-0.039***	-0.039***	-0.039***	-0.040***	-0.037***	-0.037***
	(-5.61)	(-5.59)	(-5.61)	(-5.72)	(-5.59)	(-5.57)
Constant	-0.133***	-0.132***	-0.133***	-0.131***	-0.132***	-0.132***
	(-215.11)	(-214.77)	(-215.12)	(-214.75)	(-215.86)	(-214.19)
Time FE	YES	YES	YES	YES	YES	YES
Stock FE	YES	YES	YES	YES	YES	YES

Table IN4: The impact of subperiods

We further conducts robustness when we include dummy variables of our estimated subsamples. The dummies are associated with the training of our model. Recall that we divide our sample period into three subperiods when training models. We use dummy variables to indicate whether the specific month belongs to the first or second subperiods. Finally, we report the value-weighted out-of-sample returns of the high and low groups as well as their return difference. we also use the locally estimated four-factor models to adjust these returns. The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	NN	NN	NN	ResNN	ResNN	ResNN
Constant	0.031***	0.025***	0.022*	0.041***	0.034***	0.031***
	(3.11)	(2.80)	(1.77)	(4.23)	(4.52)	(2.96)
HML	-0.897***	-0.910***	-0.926***	-0.679***	-0.685***	-0.699***
	(-3.84)	(-3.88)	(-3.86)	(-3.51)	(-3.51)	(-3.51)
SMB	-0.006	-0.006	-0.007	-0.010**	-0.010**	-0.010**
	(-1.34)	(-1.39)	(-1.43)	(-2.53)	(-2.54)	(-2.56)
MKT	0.574	0.631	0.668	0.877	0.901	0.935
	(0.83)	(0.90)	(0.94)	(1.53)	(1.55)	(1.58)
MOM	0.922***	0.909***	0.910***	0.456***	0.451***	0.452***
	(4.63)	(4.54)	(4.52)	(2.77)	(2.71)	(2.70)
Period1		-0.011	-0.008		-0.005	-0.002
		(-0.73)	(-0.47)		(-0.37)	(-0.14)
Period2			0.006			0.006
			(0.35)			(0.40)
Observations	95	95	95	95	95	95
R-squared	0.429	0.432	0.433	0.300	0.301	0.303

Table IN5. Model Comparison on Exdended Firm Characteristics

This table reports the performance comparison between Table 8 and Table 2. Both tables use the same list of models and behavioral factors, with Table 8 utilizing a more extensive list of firm characteristics than Table 2. Panel A reports the return or performance differences. Panel B presents the expanded list of firm characterists, which we construct based on available information in the Indian market following the categorization of Jensen, Kelly, and Pedersen (2023). The OLS t-statistics are reported in parentheses. ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A. Return Difference between the Estimation of Expanded Characteristics and the Orignal Estimation of Table 2

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Low Group			High Group			High Minus Low		
	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4	Mean	FF-3	Carhart-4
Linear	-0.001	-0.000	0.001	0.002	0.001	0.001	0.003	0.002	0.001
	(-0.53)	(-0.21)	(0.15)	(0.36)	(0.30)	(0.11)	(1.09)	(0.65)	(0.41)
Lasso	-0.010	-0.005	-0.003	0.002	0.001	-0.000	0.012*	0.008	0.005
	(-1.32)	(-0.82)	(-0.38)	(0.53)	(0.48)	(0.09)	(1.73)	(1.21)	(0.73)
Ridge	-0.002	-0.001	0.000	0.003	0.002	0.001	0.005	0.002	0.001
	(-0.65)	(-0.27)	(-0.05)	(0.77)	(0.55)	(0.13)	(1.32)	(0.86)	(0.38)
Random Forest	-0.003	-0.003	-0.001	0.002	0.002	0.001	0.005	0.003	0.002
	(-0.92)	(-0.67)	(-0.31)	(0.59)	(0.46)	(0.22)	(1.44)	(0.89)	(0.62)
FNN	-0.002	-0.002	-0.001	0.001	0.002	0.001	0.002	0.001	0.001
	(-0.48)	(-0.46)	(-0.26)	(0.18)	(0.12)	(0.05)	(0.99)	(0.71)	(0.56)
Residual Neural Network	-0.003	-0.003	-0.002	0.002	0.002	0.001	0.004	0.003	0.002
	(-0.97)	(-0.71)	(-0.42)	(0.42)	(0.45)	(0.27)	(1.15)	(1.02)	(0.78)

Panel B. The Expanded List of Firm Characterists

Acronym	Definition	Acronym	Definition
at_gr1	Asset Growth 1yr	debtlt_gr3a	Long-Term Debt Change 3yr
sale_gr1	Sales Growth 1yr	ebitda_gr3a	Operating Profit before Depreciation Change 3yr
lt_gr1	Total Liabilities Growth 1yr	ebit_gr3a	Operating Profit after Depreciation Change 3yr
be_gr1	Book Equity Growth 1yr	dp_gr3a	Depreciation and Amortization Change 3yr
cogs_gr1	Cost of Goods Sold Growth 1yr	fcf_gr3a	Free Cash Flow Change 3yr
opex_gr1	Operating Expenses Growth 1yr	eqnetis_gr3a	Equity Net Issuance Change 3yr
at_gr3	Asset Growth 3yr	eqnpo_gr3a	Equity Net Payout Change 3yr
sale_gr3	Sales Growth 3yr	div_gr3a	Dividend Payout Ratio Change 3yr
lt_gr3	Total Liabilities Growth 3yr	eqbb_gr3a	Equity Buyback Change 3yr
be_gr3	Book Equity Growth 3yr	eqis_gr3a	Equity Issuance Change 3yr
cogs_gr3	Cost of Goods Sold Growth 3yr	eqpo_gr3a	Net Equity Payout Change 3yr
opex_gr3	Operating Expenses Growth 3yr	capx_gr3a	Capital Expenditures Change 3yr
gp_gr1a	Gross Profit Change 1yr	capx_at	Capital Expenditures scaled by Assets
cash_gr1a	Cash and Short-Term Investments Change 1y	rd_at	R&D scaled by Assets
ppeg_gr1a	Property, Plans and Equiptment Gross Change 1yr	xido_at	Extraordinary Items and Discontinued Operations scaled by Assets
intan_gr1a	Intangible Assets Change 1yr	gp_sale	Gross Profit Margin
debtst_gr1a	Short-Term Debt Change 1yr	ebitda_sale	Operating Profit Margin before Depreciation
debtlt_gr1a	Long-Term Debt Change 1yr	ebit_sale	Operating Profit Margin after Depreciation
ebitda_gr1a	Operating Profit before Depreciation Change 1yr	fcf_sale	Free Cash Flow Margin
ebit_gr1a	Operating Profit after Depreciation Change 1yr	gp_at	Gross Profit scaled by Assets
dp_gr1a	Depreciation and Amortization Change 1yr	ebitda_at	Operating Profit before Depreciation scaled by Assets
fcf_gr1a	Free Cash Flow Change 1yr	ebit_at	Operating Profit after Depreciation scaled by Assets
eqnetis_gr1a	Equity Net Issuance Change 1yr	fcf_be	Free Cash Flow scaled by BE
eqnpo_gr1a	Equity Net Payout Change 1yr	gp_ppen	Gross Profit scaled by PPEN
div_gr1a	Dividend Payout Ratio Change 1yr	ebitda_ppen	Operating Profit before Depreciation scaled by PPEN
eqbb_gr1a	Equity Buyback Change 1yr	fcf_ppen	Free Cash Flow scaled by PPEN
eqis_gr1a	Equity Issuance Change 1yr	R1_0	Last month return
eqpo_gr1a	Net Equity Payout Change 1yr	R2_1	Return from t-2 to t-1
capx_gr1a	Capital Expenditures Change 1yr	R12_7	Intermediate momentum
gp_gr3a	Gross Profit Change 3yr	R12_2	Momentum
cash_gr3a	Cash and Short-Term Investments Change 3y	Turnover	Monthly Turnover
ppeg_gr3a	Property, Plans and Equiptment Gross Change 3yr	TradVol	Monthly Trading Volume
intan_gr3a	Intangible Assets Change 3yr	Size	Market Equity
debtst_gr3a	Short-Term Debt Change 3yr		

Figure IN1: High vs. Low Group Returns from Selected Models.

Figure 1 shows the value-weighted out-of-sample returns of the high and low groups. We first use all the models to predict retail investors' total investment returns. We then sort retail investors into five quintiles according to predicted returns, with the *High* and *Low* groups consisting of 20% of predicted winners and losers among investors, respectively.

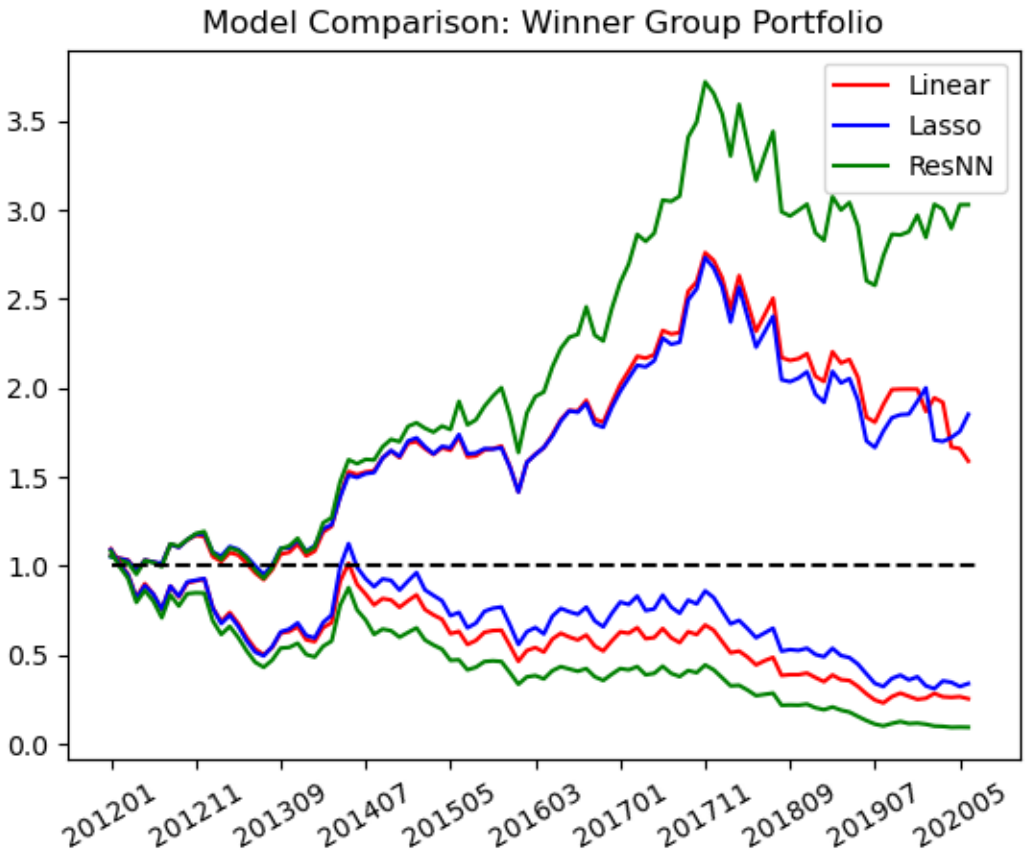


Figure IN2: Top Variable Importance of Behavioral Bias vs. Stock Characteristics

In Figure 2, we delve deeper into the analysis of the direction of influence each variable has on the prediction outcomes. Specifically, we define the directional impact of a variable as:

$$Direction(x) = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{i=1}^{N_t}\frac{\partial R_{i,t+1}^{pred}}{\partial x_{i,t}}$$

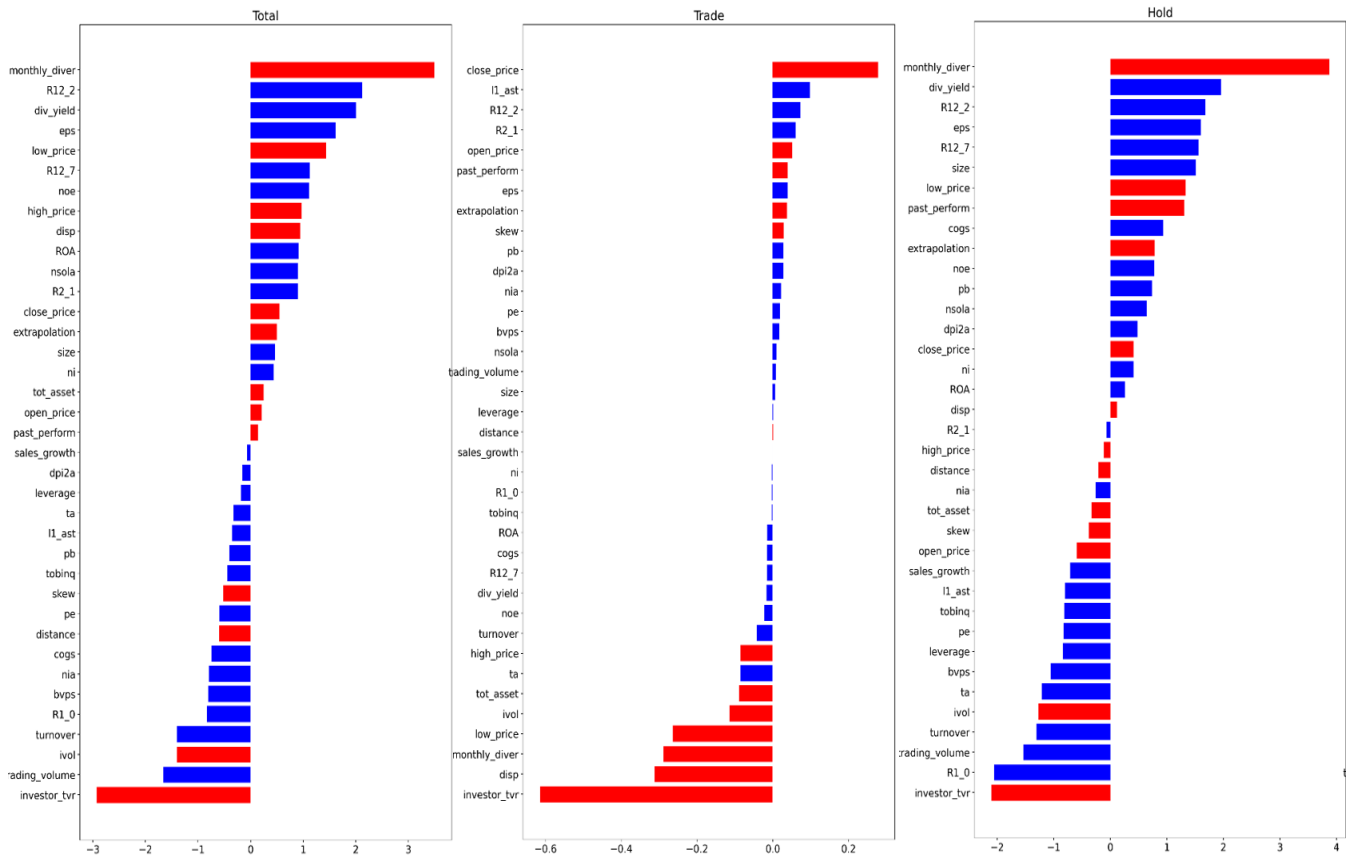
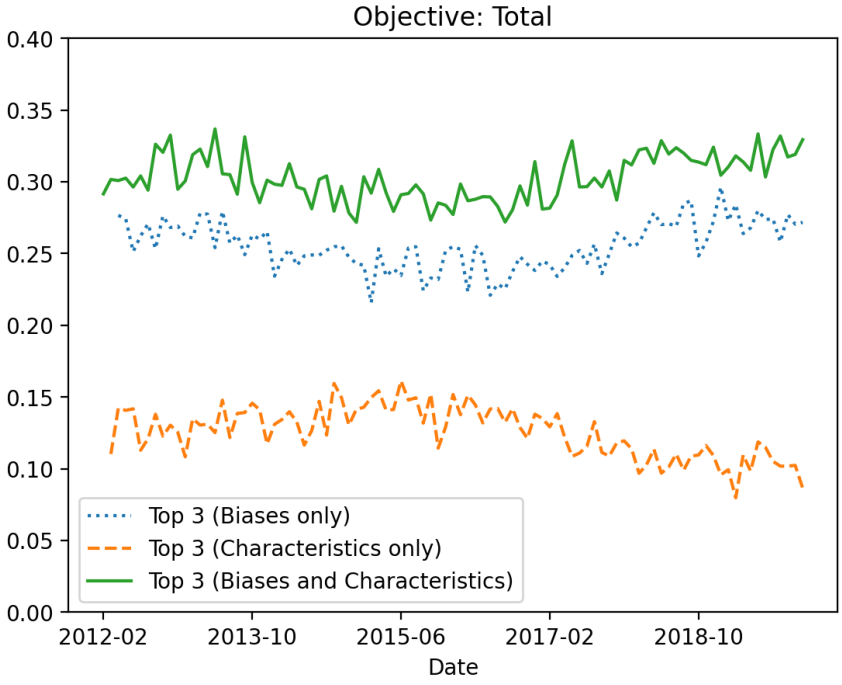
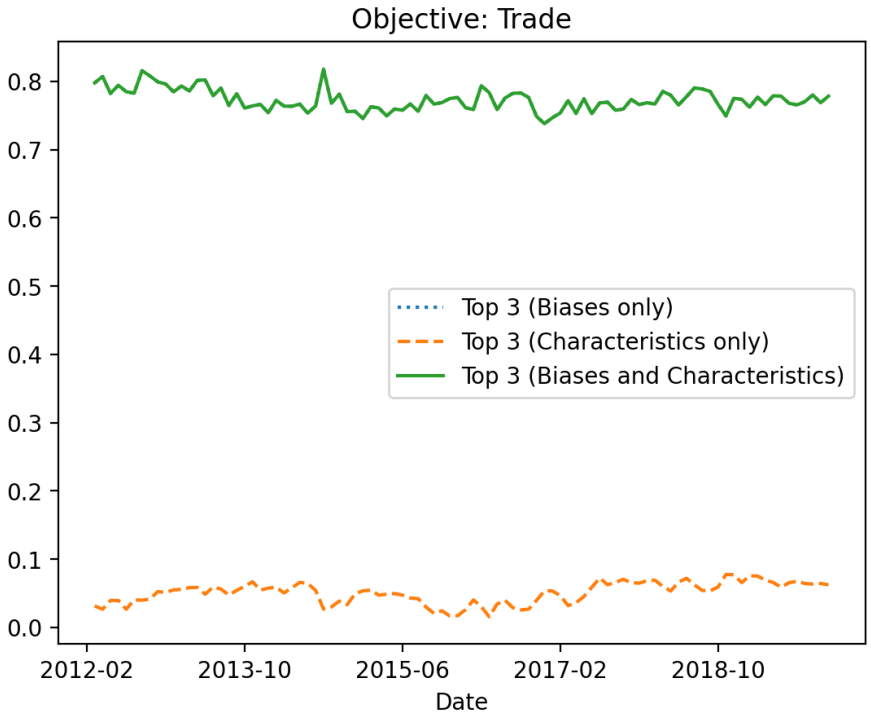
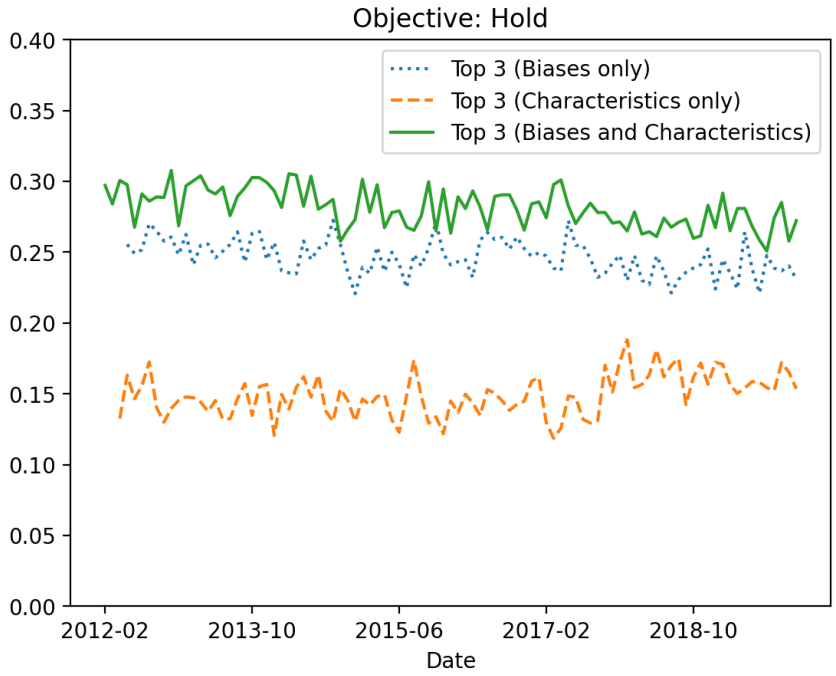


Figure IN3: The Relative Importance of Top 3 Anomalies over Time

This figure plots the time series of the joint relative importance of the top three anomoalies when the total importance of all 39 anomoalies is normalized to 1. Panel A plots the joint relative importance of the top three anomoalies to impact total returns over time. The solid line represents the joint importance of the top three anomoalies out of all 39 (solid line, including behavioral and characteristics-based), while the dashed and dotted lines plots the joint importance of the top three behavioral and characteristics-based anomoalies, respectively. Panels B and C plot the the joint relative importance top three anomoalies to impact trading and holding returns.




Objective: Trade
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
Top 3 (Biases only)
Top 3 (Characteristics only)
Top 3 (Biases and Characteristics)
2012-02
2013-10
2015-06
2017-02
2018-10
Date


Objective: Hold
0.40
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Top 3 (Biases only)
Top 3 (Characteristics only)
Top 3 (Biases and Characteristics)
2012-02
2013-10
2015-06
2017-02
2018-10
Date