TRACKING RETAIL INVESTOR ACTIVITY

EKKEHART BOEHMER, CHARLES M. JONES, and XIAOYAN ZHANG*

October 30, 2017

ABSTRACT

We provide an easy way to use recent, publicly available U.S. equity transactions data to identify retail purchases and sales. Based on retail order imbalances, we find that retail investors are informed at horizons up to 12 weeks. Individual stocks with net buying by retail investors outperform stocks with negative imbalances; the magnitude is approximately 10 basis points over the following week, or 5% annualized. Retail investors are better informed in smaller stocks with lower share prices. They do not, however, exhibit any market timing ability.

*Boehmer is with Singapore Management University, LKC School of Business. Jones is with Columbia Business School. Zhang is with Purdue University, Krannert School of Management, and Tsinghua University, PBC School of Finance. We thank Kevin Crotty, Larry Glosten, Frank Hatheway, Eric Kelley, Meijun Qian, Tarun Ramadorai, Paul Tetlock, and seminar participants at Brigham Young University, the Frankfurt School, Humboldt University zu Berlin, INSEAD, Purdue University, Rice University, Singapore Management University, Southern Methodist University, Tsinghua University, University of Utah, National Singapore University, Nanyang Technology University, and conference audiences at the 2017 SFS Cavalcade, 2017 German Finance Association, 2017 ABFER, and the 2017 CICF for helpful comments.

Are retail equity investors informed? Do they make systematic, costly mistakes in their trading decisions? The answers to these questions are important to other market participants looking for useful signals about future price moves, to behavioral finance researchers, and to policymakers who need to decide whether these investors should be protected from themselves.

Many researchers have concluded that retail equity investors are generally uninformed and make systematic mistakes. However, some more recent evidence, including Kaniel, Saar, and Titman (2008), Kelley and Tetlock (2013), and Barrot, Kaniel and Sraer (2016), suggests otherwise. These studies show that retail investors' trading can predict future stock returns. Unfortunately, most existing studies of retail order flow are based on proprietary datasets with relatively small subsets of overall retail order flow. For example, Barber and Odean (2000) and Barber, Odean and Zhu (2009) analyze data from a single U.S. retail brokerage firm, Kelley and Tetlock (2013) have data from a single U.S. wholesaler, and Barrot, Kaniel and Sraer (2016) have data from one French brokerage firm. Kaniel, Saar, and Titman (2008) and Boehmer, Jones, and Zhang (2008) use proprietary account-type data from the NYSE during the early 2000's. During that sample period, only a couple of large brokerages sent their retail order flow to the NYSE. As a result, the NYSE's market share of overall retail order activity is quite small.

In existing work, many researchers use trade size as a proxy for retail order flow. Before the spread of computer algorithms that "slice and dice" large institutional parent orders into a sequence of small child orders, small trades were much more likely to come from retail customers, while institutions were likely behind the larger reported trades. For example, Lee and Radhakrishna (2000) use a $20,000 cutoff point to separate smaller individual trades from larger institutional trades. More recently, Campbell, Ramadorai, and Schwartz (2009) effectively allow these cutoff points to vary through a regression approach that is calibrated to observed quarterly changes in institutional ownership, but they maintain the same basic assumption that small trades are more likely to arise from individual trading. However, once algorithms become an important feature of institutional order executions in the early 2000's, this trade-size partition becomes far less useful as a proxy for retail order flow. In fact, the tendency for algorithms to slice orders into smaller and smaller pieces has progressed so far that we find that during our recent sample period retail order flow actually has a slightly larger average trade size compared to other flow.

More generally, researchers need an easily implementable method to isolate retail order flow given the current automated and algorithm-driven market structure. We introduce such a measure in this paper. As one of our main contributions, we show that our measure can identify a broad swath of marketable retail order flow. Our measure builds on the fact that in the U.S. retail order flow, but not institutional order flow, receives price improvement. Because most price improvements are quite small in magnitude, fractions of a penny per share, we can identify price-improved orders from the TAQ data, a publicly available data set that contains transactions data for the U.S. equity market. We do this by identifying trades that execute at share prices with fractional pennies. Most such price-improved orders take place off-exchange and are reported to a Trade Reporting Facility (TRF). From this TRF data, we identify transactions as retail-initiated buys if the transaction price is slightly below the round penny, and retail-initiated sells if the transaction price is slightly above the round penny. This approach isolates marketable orders from retail investors, because institutional trades do not receive this type of fractional penny price improvement.[1] We discuss our approach in more detail below. Overall, we believe that our method

[1] In contrast, institutional trades often occur at the midpoint of the prevailing bid and ask prices. If the bid-ask spread is an odd number of cents, the resulting midpoint trade price ends in a half-penny. Many of these midpoint trades take place on crossing networks and are reported to the TRF. Thus, trades at or near a half-penny may very well be from institutions and are not assigned to the retail category.

of retail trade identification is conservative and clean, and we confirm the accuracy of our approach using a small sample of NASDAQ TRF audit trail data.

We analyze retail marketable order flow from the U.S. equity market for six years between January 2010 and December 2015. We find that retail investors are contrarian on average, and they are informed: the cross-section of retail order imbalances in a given week predicts the cross-section of returns over the next several weeks, consistent with Kelley and Tetlock (2013) and Barrot, Kaniel, and Sraer (2016), but inconsistent with the findings of many others. We also examine whether aggregate retail order imbalances predict future market returns, and fail to find any predictive relation. Thus, it appears that our identified retail investors have some stock-picking ability but no market-timing skill.

The article is organized as follows. We describe the data and our identification method in Section I. Section II contains our main empirical results. We provide more discussion, robustness checks, and plausibility checks in Section III. Section IV concludes.

I. Data

In the U.S., most equity trades initiated by retail investors do not take place on one of the dozen or so registered exchanges. Instead, most marketable retail orders are executed either by wholesalers or via internalization. If an order is internalized, it is filled from the broker's own inventory. If an order is executed by a wholesaler, the wholesaler typically pays the retail customer's broker a small amount, typically a fraction of a penny per share, a practice that is called "payment for order flow." Orders executed in this way are usually reported to a FINRA Trade Reporting Facility (TRF), which provides broker-dealers with a mechanism to report transactions that take place off-exchange, and these TRF executions are then included in the "consolidated tape" of all reported transactions with exchange code "D".

Often, orders that are internalized or sold in this way are given a small amount of price improvement relative to the National Best Bid or Offer (NBBO). If a retail customer wants to sell, for example, the internalizing or wholesaling counterparty often agrees to pay slightly more than the National Best Bid. This price improvement is typically only a small fraction of a cent. Common price improvement amounts are 0.01 cents, 0.1 cent, and 0.2 cents. Notably, these types of price improvements are not a feature of institutional order executions, as institutional orders are almost never internalized or sold to wholesalers. Instead, their orders are sent to exchanges and dark pools, and Regulation NMS prohibits these orders from having subpenny limit prices. Thus, institutional transaction prices are usually in round pennies. The only exception applies to midpoint trades. Reg NMS has been interpreted to allow executions at the midpoint between the best bid and best offer. More specifically, institutions are heavy users of crossing networks and midpoint peg orders that generate transactions at this midpoint price. Since the quoted spread is now typically one cent per share, this means that there are many institutional transactions reported at a half-penny. There is also a dark pool that for a time allowed some negotiation around the midquote and thus printed trades at 0.4, 0.5 or 0.6 cents.

Based on these institutional arrangements, it is fairly straightforward to identify transactions initiated by retail customers. Transactions with a retail seller tend to be reported on a TRF at prices that are just above a round penny due to the small amount of price improvement, while transactions with a retail buyer tend to be reported on a TRF at prices just below a round penny. To be precise, for all trades reported to a FINRA TRF (exchange code 'D' in TAQ), let P_{it} be the transaction price in stock i at time t, and let $Z_{it} \equiv 100 * \text{mod}(P_{it}, 0.01)$ be the fraction of a penny associated with that transaction price. Z_{it} can take on any value in the unit interval [0,1). If Z_{it} is in the interval (0,0.4), we identify it as a retail seller-initiated transaction. If Z_{it} is in the interval (0.6,1), then the transaction is coded as retail buyer-initiated. To be conservative, transactions at a round penny ($Z_{it} = 0$) or near the half-penny ($0.4 \leq Z_{it} \leq 0.6$) are not assigned to the retail category.

We are confident that our data captures most of the marketable orders from retail investors. This can be discerned from SEC Rule 606 filings, where U.S. brokerage firms are required to provide regular summary statistics on their order routing practices for non-directed orders. A directed order instructs the broker to execute an order on a given exchange or trading venue; a non-directed order gives the broker discretion on execution venue. The vast majority of retail orders are non-directed. For example, Charles Schwab reports that 98.6% of their security orders during the second quarter of 2016 were non-directed orders. The corresponding figure for TD Ameritrade is 99%. According to the Rule 606 filings by these two retail brokerage firms, more than 90% of these orders receive price improvements.

As discussed above, Reg NMS requires that limit orders be priced at round pennies, so our approach by definition will only identify marketable retail orders. The 606 filings by brokerage firms are also partitioned into market and limit orders, which allows us to gauge the relative prevalence of these two types of orders. While we have not collected data systematically, from our spot checks, retail investors use market orders and limit orders in more or less equal numbers. For example, the Charles Schwab brokerage firm reports that for the second quarter of 2016, market orders account for 50.0% of its customers' non-directed orders in NYSE-listed securities, while limit orders account for 45.1%, and other orders account for the remainder. For securities listed on NASDAQ, limit orders are slightly more prevalent than market orders at Schwab, with market orders accounting for 44.0% and limit orders accounting for 50.7%. Note also that limit orders may be cancelled without executing, and limit orders may in fact be marketable, so it is likely that most overall retail trading activity arises from marketable orders. Thus, we believe that our approach picks up a majority of overall retail trading activity.

After collecting information on retail investor activity, we merge it with stock return data and accounting data from CRSP and Compustat, respectively. We only include common stocks with share code 10 or 11 (which excludes mainly ETFs, ADRs, and REITs) listed on the NYSE, NYSE MKT (formerly the Amex), or NASDAQ. We remove low-priced stocks by requiring the minimum stock price to be $1 at the previous month-end. Our sample period is from January 3, 2010 to December 31, 2015. On each day, we have an average of around 3,000 firms included in the sample.

Table I presents summary statistics of retail investors' activity. We pool observations across stocks and across days, and compute the mean, standard deviation, median, 25th and 75th percentile. Our sample includes over 4.6 million stock-day observations. For the number of shares traded per day (*vol*), the mean share volume is around 1.23 million, and the standard deviation is about 6.85 million shares. The average stock has 5,917 trades each day (*trd*), and comparing this to the average daily share volume implies that the average trade size over this sample period is about 200 shares. Our identified retail investor activity is only a small part of the overall trading volume. The average daily buy volume from retail investors (*indbvol*) is 42,481 shares, and the average daily sell volume from retail investors (*indsvol*) is 42,430 shares. Thus, we identify an average of 84,911 shares per stock-day traded by retail investors, about 6.91% of the average total shares traded each day. The average number of buy trades from retail investors (*indbtrd*) each day is 110, and the average sell trades from retail investors (*indstrd*) each day is 108. Thus, the total number of trades per stock-day from retail investors is 218, around 3.68% of the total number of trades. Over our sample period, there is slightly more buying than selling by retail investors.

Information on odd lot trades (trades of fewer than 100 shares) is reported on the TRF and on the consolidated tape beginning in December of 2013 (see O'Hara, Yao, and Ye, 2014). During the December 2013 through December 2015 sample period where odd lot data are available, the daily average of odd lot retail buy and sell volume (*oddindbvol* and *oddindsvol*, respectively) are 506 and 443 shares respectively, for a total of 949 shares traded by retail investors in odd lots per average stock-day. This is about one-third of all odd lot share volume at 3,027 shares. A similar pattern exists for the number of trades. Older papers studying odd lots generally find that these retail-dominated orders are virtually uninformed, so we study odd lots separately to determine whether the information content in odd lots executed by retail customers differs from that of retail round lots.

In Figure 1, we provide further statistics on the overall properties of our identified retail trades. Panel A characterizes trade sizes in dollars. For each retail trade, we compute the trade size in dollars by multiplying the number of executed shares by the transaction price. For each year of our sample, we compute the 25^{th} percentile, the median, and the 75^{th} percentile of retail trade size. The median retail trade size is mostly around \$8,000, and the interquartile range is mostly between \$2,000 and \$25,000. Panel B reports the distribution of subpenny prices. We separate all trades into 12 groups or bins. We separate out trades that take place at a round penny or a half penny; the other bins are each 0.1 cent wide. We pool the sample across days and stocks, and we report the number of shares reported in the different subpenny buckets. Not surprisingly, most of the share volume occurs at round and half pennies, with average stock-day share volumes around 25,000 and 7,000, respectively. The next most prevalent occurrence, averaging around 3,000 shares per day per stock, is a subpenny price within 0.1 cents of a round penny. Other subpenny bins are less prevalent, with most averaging around 1,000 shares per stock per day.

To measure retail investors' directional trades, we compute four order imbalance measures, for each stock *i* on each day *t*:

$$oibvol(i,t) = \frac{indbvol(i,t) - indsvol(i,t)}{indbvol(i,t) + indsvol(i,t)} \tag{1}$$

$$oibtrd(i,t) = \frac{indbtrd(i,t) - indstrd(i,t)}{indbtrd(i,t) + indstrd(i,t)} \tag{2}$$

$$oddoibvol(i,t) = \frac{oddindbvol(i,t) - oddindsvol(i,t)}{oddindbvol(i,t) + oddindsvol(i,t)} \tag{3}$$

$$oddoibtrd(i,t) = \frac{oddindbtrd(i,t) - oddindstrd(i,t)}{oddindbtrd(i,t) + oddindstrd(i,t)} \tag{4}$$

The first two measures are calculated using retail round lot executions before December 2013 and by aggregating round lot and odd lot executions thereafter, while the last two measures are calculated using only retail odd lots, and thus these measures begin in December 2013.

Summary statistics on the order imbalance measures are reported at the bottom of Table I. Across all stocks and all days, the mean order imbalance for share volume, *oibvol*, is -0.038, with a standard deviation of 0.464, and the mean order imbalance for trade, *oibtrd*, is -0.032, with a standard deviation of 0.437. The correlation between *oibtrd* and *oibvol* is around 85%, indicating great overlap in information covered in these two measures. Our later discussions mostly focus on *oibvol*, but the results using these two measures are quite similar given the high correlation between the two. Overall, the order imbalance measured in shares is close to zero on average, but with sells slightly more prevalent than buys, which is consistent with findings in Kaniel, Saar, and

Titman (2008). More importantly, the sizable standard deviation measures show that there is substantial cross-sectional variation in the activity levels and trading direction of retail investors. The odd lot order imbalance measures exhibit similar patterns.

In Figure 2, we plot the time-series of the cross sectional mean, median, the 25th percentile and the 75th percentile of the order imbalance measures over the six-year sample period. Across all four order imbalance measures, the means and medians are all close to zero, while the 25th percentiles are mostly around -0.3, and the 75th percentiles are mostly around 0.2. There are no obvious time trends or structural breaks in the time-series observations.

Because we work with TAQ data, we do not directly observe whether a trade is a buy or sell. We validate our identification of buy and sell orders through two channels: the proprietary data used in Kelley and Tetlock (2013) and a sample of proprietary NASDAQ data. The data in Kelley and Tetlock (2013) come from one large wholesaler, and they have buy/sell identifiers in the data. In a recent discussion of our paper, the correlation between our order imbalance measures and imbalances calculated from Kelley and Tetlock (2013)'s observed trade directions is in the range of 0.345 to 0.507, with an average of 0.45.[2] These correlations should be less than one, because their flow is from only one wholesaler. In contrast, our measure is from TRF, which covers nearly all retail order executions, but we will not be able to include unimproved or midpoint retail executions in our order imbalance measures. Interestingly and reassuringly, it turns out that in a regression horse race, our order imbalance measures and those of Kelley and Tetlock (2013) both significantly predict future stock returns positively.

[2] We are grateful to Eric Kelley for calculating and sharing these correlations with us.

The second channel for cross-validation is a dataset generously provided by NASDAQ which covers all intraday transactions on its TRF for 100 stocks during October 2010. The NASDAQ data separately reports trades by broker-dealers who are retail internalizers or wholesalers, and the dataset directly signs each buy and sell trade within this sample. The same dataset is used in Menkveld, Zhou, and Zhu (2017), who provide more detail about the data. For stocks with a share price below \$100, our subpenny approach matches the TRF's correct buy/sell sign 98.2% of the time, while the standard Lee and Ready (1991) trade-signing algorithm gets the trade sign right 96.7% of the time. When we calculate an order imbalance in shares using the NASDAQ data with direct trade-sign identifiers, the correlation with our order imbalance measure is 83.9%. There are two main reasons that this correlation is less than one. First, our order imbalance measure excludes trades at round and half pennies. Second, our order imbalance measure includes trades printed on the competing NYSE TRF, while the NASDAQ dataset does not. Nevertheless, these high correlations provide confidence that our order imbalance measures closely reflect the true buy and sell activities from retail investors.

Finally, to reduce the impact of microstructure noise on our results, we choose to focus on weekly horizons, even though we have daily data.. That is, our main variables of interest are average order imbalances over 5-day horizons, and 5-day returns. Blume and Stambaugh (1986) show that using the end-of-day closing price to compute daily returns can generate an upward bias, due to bid-ask bounce. Therefore, we compute two versions of weekly returns, one by compounding CRSP daily returns, which is based on daily closing prices, and one by compounding daily returns using the end-of-day bid-ask average price. We report results for both types of returns.

II. Empirical Results

A. What Explains Retail Investor Order Imbalances?

We start our empirical investigation by investigating what drives the trading of retail investors. More specifically, we examine whether retail investors' order flow is contrarian or momentum. To allow maximal time-series flexibility and focus on cross-sectional patterns, we adopt a Fama and MacBeth (1973) 2-stage estimation. At the first stage, for each day, we estimate the following predictive regression:

$$oib(i,w) = b0(w) + b1(w)'ret(i,w-1) + b2(w)'controls(i,w-1) + u(i,w) \qquad (5)$$

where we use various horizons of past returns, *ret*(i, w-1) and various control variables from the past to explain the order imbalance measures, *oib*(*i*,*w*) at week *w*. The first stage estimation generates daily time-series of coefficients, $\{b0(w), b1(w)', b2(w)'\}$. At the second stage, we conduct statistical inference using the time-series of the coefficients. Because we use overlapping daily frequency data for weekly order imbalance and return measures, the standard deviations of the time-series are calculated using Newey-West (1987) with 5 lags.

To explain the order imbalance over week *w*, from day 1 to day 5, we first include its own lag, the past week order imbalance measure from day -4 to day 0, or *oib*(*w*-1). We also include past returns over three different horizons: the previous week (*ret*(*w*-1)), the previous month (*ret*(*m*-1)), and the previous six months (*ret*(*m*-7,*m*-2)). For control variables, we use log market cap, log book-to-market ratio, turnover (share volume over shares outstanding), and daily return volatility, all observed from the previous month.

Results are in Table II, with regression I and II explaining the order imbalance measured in shares, and regression III and IV explaining order imbalance measured using the number of trades.

In the first regression, the order imbalance using share volume, *oibvol*, has a positive correlation with its own lag, with a highly significant coefficient of 0.22, indicating that directional trading activity is somewhat persistent over successive weeks, as suggested in Chordia and Subrahmanyam (2004). The coefficients for past week, past one month, and past six-month returns are -0.9481, -0.2778, and -0.0586, respectively. All three coefficients are highly significant, which shows that retail investors are contrarian, especially over short horizons. Control variables indicate that investors tend to buy more aggressively in larger firms, growth firms, and firms with higher turnover and higher volatility. All coefficients are highly significant.

We observe similar patterns for regressions II, III and IV, which use different return and order imbalance measures. At the weekly horizon, the results are similar across ways of computing returns and order imbalances. Henceforth, we focus our discussion on bid-ask midpoint returns, which do not have bid-ask bounce and thus exhibit a smaller degree of time-series predictability compared to returns based on transaction prices. We also include CRSP returns in the results for completeness and robustness.

In our sample, retail investors on average are contrarian over weekly horizons. Why are they contrarian? There are at least two possibilities: retail investors might be informed and know something about future stock returns, or retail investors are uninformed but tend to simply "buy the dips" by trading in the opposite direction of order imbalances and returns. In the latter case, retail investors could be worse than uninformed, making systematic mistakes by buying prior to further price declines or selling prior to further share price increases. We investigate these possibilities in the next sections by examining whether the retail order imbalance can predict future stock returns.

Our contrarian results match some of the findings in previous literature. For example, retail traders are found to be contrarian in Kaniel, Saar and Titman (2008) over monthly horizons, and by Barrot, Kaniel and Sraer (2016) over daily and weekly horizon. In contrast, Kelley and Tetlock (2013) paint a more complex picture. They find that at weekly horizons, retail order imbalance measures are contrarian and have negative coefficients on past returns. Over shorter (daily) horizons, however, they find that market order imbalances actually have a positive coefficient on the lagged 1-day return, which implies momentum rather than contrarian behavior. In unreported results, we find that for our retail order imbalance measures, the coefficient on the lagged 1-day return is either significantly positive or not significantly different from zero, which is broadly in line with the findings of Kelly and Tetlock (2013). Given that our main results are on weekly data, these daily results are available on request rather than reported in the paper.

B. Predictive Regressions for Future Stock Returns

Can retail investors' activity provide useful information for future stock returns? In this section, we examine the predictive power of our order imbalance measures using Fama-MacBeth regressions as follows:

$$ret(i,w) = c0(w) + c1(w)oib(i,w-1) + c2(w)'controls(i,w-1) + e(i,w) \qquad (6)$$

where we use the order imbalance measure from week $t-1$, $oib(i, w-1)$ and various control variables to explain the next week's return, $ret(i,w)$. As in the previous section, because we use overlapping daily frequency data for weekly order imbalance and return measures, the standard deviations of the time-series are adjusted using Newey-West (1987) with five lags. If retail investors are informed about future stock returns, past order imbalance should be able to predict future returns in the right direction and we expect the coefficient c1 to be significantly positive. If retail investors are uninformed or wrongly-informed, the coefficient c1 might be close to zero or significantly negative.

We again include past returns as control variables, using three different horizons: the previous week, the previous month, and the previous six months (from month *m*-7 to month *m*-2). In addition, we include log market cap, log book-to-market ratio, turnover, and daily return volatility, all from the previous month. We report estimation results in Table III. In regression I, we use order imbalance based on share volume, *oibvol*, to predict next week's return based on bid-ask midpoints. The coefficient on *oibvol* is 0.0009, with a *t*-statistic of 15.60. The positive and significant coefficient shows that if retail investors buy more than they sell in a given week, the return on that stock in the next week is significantly higher. In terms of magnitude, we report at the bottom of the table that the inter-quartile range for the *oibvol* measure is 1.1888 per week. Multiplying the interquartile difference by the regression coefficient of 0.0009 generates a weekly return difference of 10.89 basis points (or 5.66% per year annualized) when moving from the 25^{th} to the 75^{th} percentile of the *oibvol* variable. When we use different order imbalance and return measures, the same pattern is present, and the weekly interquartile difference in the conditional mean return ranges from 9.31 basis points to 11.44 basis points (4.84% to 5.94% per year).

For the control variables, we observe negative coefficients on *ret*(*w*-1), which shows the presence of weekly return reversals, and positive coefficients on the other longer-horizon returns, which shows momentum. Size, book to market, turnover and volatility all carry their expected signs, while most of them are not statistically significant. This also confirms that the predictability we have found is not simply a manifestation of some other size, book-to-market, turnover, or volatility anomaly.

To summarize, order imbalance measures from retail investors strongly and positively predict one-week ahead stock returns. As a group, these retail investors seem to be informed about future stock return movements.

C. Subgroups

Each day, our sample includes on average more than 3,000 firms. Is the predictive power of retail investor order imbalances restricted to a particular type of firm, or do informed retail investors have preferences for particular types of firms? To investigate this, we analyze various firm subgroups in this section. We first sort all firms into three groups, based on a firm or stock characteristic observed at the end of previous month. Then we estimate equation (6) within each characteristic group. That is, we allow all coefficients in equation (6) to be different within each group, which allows substantial flexibility in the possible predictive relationship across these different groups.

The results are in Table IV. To save space, we only include results on weekly returns computed using end-of-day bid-ask average price. We first sort all stocks into three different size groups based on market capitalization: small, medium and big. The results are reported in Panel A of Table IV. In the left panel, we report coefficients on *oibvol*, the order imbalance computed from share volume. When we move from the smallest one-third of firms by market cap to the largest tercile, the coefficient on *obivol* decreases from 0.0013 to 0.0003, and the *t*-statistic decreases from 13.90 to 3.68. Clearly, the predictive power of retail order imbalance is much stronger for smaller firms than for larger cap firms, but the predictability remains reliably present in all three groups. Economically, the interquartile difference in weekly returns is 21.9 basis points for the smallest firms (11.39% per year), and 2.6 basis points for the largest firms (1.35% per year). The results in the right panel using order imbalance based on the number of trades (*oibtrd*) are quite similar.

In Panel B of Table IV, we sort all firms into three groups based on previous month-end share price. In the left panel, moving from the lowest share price firms to the highest share price firms, the coefficient on *oibvol* decreases from 0.0014 to 0.0002, and *t*-statistics go from 13.34 to 3.23. In terms of magnitude, the interquartile weekly return difference is 20.5 basis points (10.66% per year) for the lowest price firms and only 2.0 basis points for the firms with the highest share price (1.04% per year). For specifications using *oibtrd*, which are reported in the right panel, the results are similar, with slightly lower coefficients and *t*-statistics. The pattern is clear: retail investor order imbalances predict returns more for low price firms.

Next we sort all firms based on previous month turnover, which is a proxy for liquidity. In the left panel, moving from the tercile of low trading activity to the firms with more turnover, the coefficient on *oibvol* decreases from 0.0011 to 0.0007, and the *t*-statistic decreases from 15.60 to 4.98. In terms of magnitude, the interquartile weekly return difference is 20.5 basis points (10.66% per year) for the firms with the lowest turnover and 6.5 basis points for the firms with the highest turnover (3.38% per year). For specifications based on *oibtrd* in the right panel, the results are similar, with slightly lower coefficients and *t*-statistics. Overall, retail investor order imbalances predict returns better for firms with lower trading activity.

To summarize this section, we find that the predictive power of the retail investor order imbalance is significant and positive for all but one subgroup, which shows that the predictive power is not particularly driven by special subgroups. But there is a clear cross-sectional pattern for the predictive power. The predictive power of retail order imbalance is much stronger for small, low share price, and low liquidity firms.

D. Longer Horizons

The results in the previous section show that retail investor order imbalances can predict next week's returns positively and significantly. One natural question is whether the predictive power is transient or persistent. If the predictive power quickly vanishes or reverses, retail investors may be capturing price reversals, while if the predictive power lingers, retail investors are potentially informed about information related to firm fundamentals. To answer this question, we extend equation (6) to longer horizons as follows:

$$ret(i, w + k) = c0(w) + c1(w)oib(i, w) + c2(w)'controls(i, w) + e(i, w + k) \qquad (7)$$

That is, we use one week of order imbalance measures to predict k-week ahead returns, $ret(i,w+k)$, with k=1 to 12. To observe the decay of the predictive power of retail order imbalance, the return to be predicted is a weekly return over a one-week period, rather than a cumulative return over *n* weeks, which is an average over all weeks involved. If retail order imbalances have only short-lived predictive power for future returns, we might observe the coefficient c1 decrease to zero within a couple of weeks. Alternatively, if the retail order imbalance has longer predictive power, the coefficient c1 should remain statistically significant for a longer period. In our empirical estimation, we choose k ranging from two weeks to 12 weeks.

We report the results in Table V, with results based on bid-ask average returns in Panel A, and results based on closing transaction prices in Panel B. In Panel A, when we extend the window from two weeks to 12 weeks, the coefficient on *oibvol* monotonically decreases from 0.00055 to 0.00007, and the coefficient on *oibtrd* gradually decreases from 0.00048 to 0.00006. The coefficients are statistically significant up to six or eight weeks ahead. Results in Panel B are similar.

Figure 3 plots the coefficients over different horizons. With *oibvol* in Panel A and *oibtrd* in Panel B, the general pattern is similar: the predictive power of retail order imbalances gradually decreases to zero over six to eight weeks. Potentially the retail order imbalances capture information longer than a one-month quick reversal.

E. Long-Short Portfolios

One might wonder whether we can use retail order imbalances as a signal to form a profitable trading strategy. As discussed earlier, both *oibvol* and *oibtrd* are publicly available information. In this section, we form quintile portfolios based on the previous week's average order imbalance, then we hold the quintile portfolios for up to 12 weeks. If retail investors can select the right stocks to buy and sell, then firms with higher or positive retail order imbalance would outperform firms with lower or negative order imbalance.

Table VI reports long-short portfolio returns, where we buy the stocks in the highest quintile of scaled order imbalance, and short the stocks in the lowest order imbalance quintile. Portfolio returns are value-weighted using the previous month-end market cap. Because the holding period can be as long as 12 weeks, we report both the raw returns and risk-adjusted returns using the Fama-French three-factor model. Given the overlapping data, we adjust the standard deviations of the portfolio return time-series using Newey-West (1987) with the corresponding number of lags.

In Panel A, the long-short strategy is based on the previous week's *oibvol*, and we report bid-ask average returns. Over a one-week horizon, the long-short portfolio return is 0.104%, or 5.41% per year annualized. The *t*-statistic is 3.28. Risk adjustment using the Fama-French three-factor model does not make much difference: the weekly Fama-French alpha for the long-short portfolio is 0.101%, with a *t*-statistic of 3.09. When we increase the holding horizon to 12 weeks, the mean return becomes 0.617%, with a *t*-statistic of 2.35. The general pattern is that holding-period returns (and alphas) continue to grow at a decreasing rate over time. We never observe evidence of a reversal in returns. In terms of statistical significance, the *t*-statistics are significant even at the longest 12-week horizon. The results are slightly weaker compared to the Fama-MacBeth regressions, and the main reason is that in this section we value-weight the portfolio returns across firms, while the Fama-MacBeth approach implicitly weights each stock equally.

When we restrict portfolio formation to one of the three different market cap groups, the one-week return is 0.472% (or 24.54% per year) with a *t*-statistic of 9.20 for the smallest size firms, while the one-week return is 0.071% (or 3.69% per year) with a *t*-statistic of 2.16 for the largest size firms. When the holding horizon becomes longer, the return on the long-short strategy is still significant and positive out to 12 weeks for the smallest third of firms, but significance is marginal for the largest tercile.

Results in Panel B, using *oibtrd*, are qualitatively similar, but with smaller magnitude and lower statistical significance. This is expected, because as discussed earlier, *oibvol* provides similar but finer information than *oibtrd*.

To make sure that the statistical significance in return differences is not driven by particular sample periods, we provide a time-series plot of the return differences between quintile 1 and 5 in Figure 4, where the portfolios are sorted on *oibvol* and the holding period is one week. Over our 6-year window, we observe both time-variation in the return differences and positive and negative spikes. Most of the data points are positive, however, and the positive returns are not driven by particular sample subperiods. Unreported plots of alphas show the same pattern.

III. Discussion

Retail order imbalances are informed about future stock returns. The predictive ability lasts up to 8 weeks and is stronger for smaller firms and lower priced firms. Overall, retail investors show some stock picking ability. In this section, we discuss several related issues to put the retail order imbalance's predictive power in perspective. In Section 4.1, we discuss whether retail investors can time the market, and we examine whether the predictive power is related to overall market conditions in Section 4.2. We investigate the predictive power of odd lot retail orders in section 4.3. Retail trades occur with different sizes, and we examine the informativeness of large vs. small trade sizes in Section 4.4. It is important to understand the role of wholesalers in this setup, and in Section 4.5 we look into the magnitude of price improvement and the profitability of interacting with retail order flow. To understand the nature of the information that retail order flows capture, we link retail order imbalances to earnings news in section 4.6. Finally, we examine whether the retail order imbalances can still predict future returns if we control for overall market order imbalances in section 4.7. To save space, all returns in this section are bid-ask returns.

A. Aggregate Retail Order Imbalance

If retail order imbalances can predict future stock returns in the cross section, retail investors may also be able to time aggregate market moves. To investigate this possibility, we aggregate retail order imbalances across all firms in order to predict aggregate stock market returns. We estimate the following equation,

$$mkt(w+1, w+k) = d0 + d1 * aggoib(w) + u(w+1, w+k) \tag{8}$$

where *mkt*(w+1,w+k) is the future k-week cumulative market return from week w+1 to week w+k, and *aggoib*(w) is the current aggregated retail order imbalance measure for week w. We compute *aggoib* using either value-weighted or equal-weighted *oibvol* or *oibtrd* measures.

Results are in Panel A of Table VII. Regardless of the weighting scheme and the order imbalance measure, the result is the same: there is no evidence that retail investors can reliably predict future market returns. Although retail investors display stock selection skills, they do not seem to be able to do market timing.

Our approach can also be applied to identify retail trading of exchanged-traded funds (ETFs). In Table VII Panel B, we examine retail order flow in a large cross-section of ETFs over the same time period. In cross-sectional predictive regressions of the form in equation (6), the coefficient is mostly around or below 1 basis point, which is much smaller than the comparable coefficients from Table III, and the *t*-statistics are mostly insignificant. This suggests that retail traders do not have much information about sector valuation or overall equity market values. To separate out sector-oriented information from broader market-wide information, we pick the six largest ETFs that focus on the overall U.S. equity market by tracking comprehensive U.S. equity indexes: SPY, IVV, VTI, VOO, IWM, and IWB. The results are reported in the last row of Panel B. Consistent with the Panel A market timing results, we find little evidence that retail order flow is informed about future returns on broad equity market ETFs.

B. Market Conditions

Barrot, Kaniel and Srear (2016) find that retail trades contain more information when markets are volatile, specifically when the VIX option-implied volatility index is high. Their sample is from 2002 to 2010, during which the VIX experiences dramatic changes. In contrast, our sample is from 2010 to 2015, and the VIX is far less volatile. Still, we separate our sample into two parts, when VIX is higher than the historical median of 18%, and when VIX is lower than the historical median.

We re-estimate equation (6) for the high VIX and low VIX subsample, and results are presented in Panel C of Table VII. Comparing the low and high VIX regimes, the coefficient on *oibvol* is quite similar, yet the *t*-statistic is higher when VIX is low rather than high. This might not be surprising, given that volatility of all variables increases when VIX is high. Overall, the predictive power in both high and low VIX regimes is positive and significant.

C. Odd Lots

In this section, we investigate the behavior of odd lot retail trades over the post-December 2013 period when odd lot transactions are reported to the consolidated tape. In untabulated results, we find that odd lot retail order imbalances are fairly similar to overall retail order imbalances in two dimensions. First, they are also contrarian with respect to the previous one-week return. Second, odd-lot order imbalances are contemporaneously correlated with market returns, but can't predict future market returns. Can odd lot retail order flow predict future firm level returns? We estimate regression (6) using odd lot retail order imbalances and present the results in Panel D of Table VII. Both coefficients are positive but not statistically significant. In unreported results, we find that daily odd lot order imbalance measures can significantly predict returns for the next trading day. We conclude that odd lot order imbalance predictive power is much weaker than retail order imbalance predictive power.

D. Retail Order Sizes

From Figure 1 Panel A, we know that a median market order submitted by a retail investor is around $7,000. The median retail trade is about 400 shares. Previous "stealth trading" literature argues that medium size orders are more likely to be informed, and large orders are usually broken into smaller size orders.

To see if there is differential information content by order size, we partition orders into large vs. small orders using 400 shares as the cutoff, and we estimate the predictive regression for each group separately. Results are reported in Panel E of Table VII. We find that both larger orders and small orders predict future stock returns, but the larger orders' predictive power is stronger. Our results suggest that more informed retail investors demand immediacy by using larger market orders, and stealth trading does not seem to characterize the trading of retail investors.

E. The Profitability of Marketable Retail Order Flow

If marketable retail order flow is sufficiently informed, trading with these orders would be unprofitable. This might make readers wonder whether our results are consistent with the apparently profitable business model of internalizers and wholesalers. Ultimately, as long as the information content of retail order flow is less than the bid-ask spread being charged, internalizers and wholesalers on average earn positive revenues by trading with these orders. For example, if a retail buy and sell order arrive at the same time, they offset each other, and a wholesaler earns the full bid-ask spread charged (the quoted spread less the price improvement given). Ultimately, internalizers and wholesalers are only exposed to adverse selection on retail order imbalances. The summary statistics in Table I show that there is a substantial amount of offsetting retail order flow. The interquartile range for the volume-based daily order imbalance measure goes from -0.301 to 0.217, indicating that even at the ends of these ranges, more than two-thirds of the retail order flow in such a stock on a given day is offsetting buys and sells.

To get a better sense of the profitability of interacting with retail order flow, we compute standard microstructure information-content measures for the retail trades in our sample. Specifically, we calculate proportional effective spreads, 1-minute price impacts, and 1-minute realized spreads for all retail buys and sells during 2015. Realized spreads are a standard proxy for trading revenue earned by a liquidity provider. We apply standard data filters, eliminating all trades where effective spreads exceed \$1, and we calculate dollar-volume weighted averages across all stocks. We find that the mean effective half-spread is 16 basis points. The 1-minute price impact is 4 basis points, leaving a realized half-spread of 12 basis points. Said another way, interacting with our identified retail order flow is profitable because the bid-ask spreads are sufficiently large. The liquidity provider (in this case, the wholesaler or internalizer) loses about 4 basis points (the price impact) of the bid-ask spread to short-term price moves, but this leaves about 12 basis points (the realized spread) of the bid-ask spread as average trading revenue to the liquidity provider. Note that the realized spread is a very crude measure of trading revenue. Furthermore, we cannot measure payments made by wholesalers to introducing brokers, nor can we measure the other costs associated with a wholesaleing or internalization operation. However, these realized spreads are considerably higher than the realized spreads associated with on-exchange transactions, so we feel comfortable in concluding that the price improvement business model is quite profitable for wholesalers and internalizers.

We can also examine some of the segmentation that is performed by these liquidity providers. For instance, the magnitude of price improvement is chosen by the internalizers/wholesalers. They can rationally incorporate the potential information embedded in retail orders and only offer price improvement to the point that they themselves can still profit from the trade. That is, if they infer there might be relevant information in the retail order, they might offer less price improvement, and on the other side, if they conclude that the retail order is unlikely to contain relevant information, they might be willing to offer more price improvement. If this is true, the predictive power of retail order imbalances should be higher for retail trades with less price improvement.

In the main empirical analysis, we group all orders with subpenny prices between 0.6 and 1 to be retail-initiated buy orders, and between 0 and 0.4 to be retail-initiated sell orders. In this section, we further separate orders into "less price improvement" and "more price improvement." For transactions with less improvement, we define buyer-initiated trades as transactions with prices between 0.8 and 1, and seller-initiated trades as trades with transaction prices between 0 and 0.2. For the "more price improvement" category, we define buyer-initiated trades as trades with transaction prices between 0.6 and 0.8, and seller-initiated trades as trades with transaction price between 0.2 and 0.4. We compute retail order imbalances following equation (1) and (2). To compare predictive power of retail order imbalance for "more" or "less" price improvement, we estimate equation (4) on each order imbalance measure separately.

Recall that the distribution of subpenny price improvements is displayed in Figure 1 Panel B. Most transaction prices happen at a round penny or a half penny. Based on the other bins, each covering 10 bps, there is slightly more trading volume for the "less price improvement" category compared to the "more price improvement." Regression results for the cross-section of future returns are in Panel F of Table VII. For less price improvement, the coefficients range from 0.0004 to 0.0007, all with *t*-statistics above 5. For more price improvement, the coefficient ranges from 0.0001 to 0.0002, all with *t*-statistics below 4. Clearly, both sets of retail order imbalances have predictive power for future stock returns, but the retail trades with less price improvement have stronger predictive power, indicating that internalizers/wholesalers successfully price discriminate against retail orders with potentially more information content. Similar to the presence of large realized spreads, this observation also supports the viability of the internalization/payment for order flow business model.

F. The Information Content of Marketable Retail Order Flow

Our earlier results show that retail order flow predicts future stock returns in a positive way, indicating that retail investors are informed. The next natural question is: what information do retail investors have? Kelley and Tetlock (2013) use the Dow Jones news archive to identify whether retail investors are informed about cash flow news, and find that retail market orders can predict earnings surprises.

To explore this, we examine whether retail order flow becomes more predictive around earnings news. Specifically, we estimate a variant of equation (6) that allows the predictive relationship to differ based on the variable *eventday*, an indicator that takes a value of 1 if day t is an earnings announcement day and zero otherwise. The results are in Table VII Panel G. Our results show that the predictive power of retail order flow is larger on announcement days, but the difference is not statistically significant.

G. Controlling For Overall Order Imbalances

Previous studies, such as Chordia and Subrahmanyam (2004), find that overall order imbalances (calculated using all reported transactions, including individual and institutional) can predict future stock returns. In our data set, overall order imbalances and retail order imbalances are significantly correlated at around 30%. An interesting question to ask is whether overall and retail order imbalances are relatively orthogonal to each other. Specifically, if we control for the overall order imbalance, can the retail order imbalance still predict future stock returns?

We proceed in two steps to address this question and report the results in Panel H of Table VII. In the first step, we re-estimate equation (6), with the overall order imbalance from the previous week rather than retail order imbalance as a key predicting variable. Consistent with previous literature, we find that overall order imbalances significantly predict future stock returns, with a coefficient of 0.0004 and a significant *t*-statistic of 3.32.

In the second step, we estimate equation (6) using the retail order imbalance variables as key predicting variables, and include the overall order imbalance as a control. With both retail and market order imbalances in the model, retail imbalances are significantly positive, and they completely drive out the effect of market order imbalances. That is to say, the predictive power of retail order imbalance seems to be stronger than that of the overall order imbalance measure.

H. When Effective Spread is Less Than 1 Cent

Our identification for buy and sell orders relies on an implicit assumption that price improvements are always a small fraction (less than half) of a cent. If price improvements are larger, our method may not correctly sign trades. For example, if a stock has a bid price of $50.01 and an ask price of $50.04, and a retail market buy order arrives and is improved by 0.75 cents, the reported transaction price would be $50.0325, and our trade-signing approach would erroneously conclude that this is a sell order. To investigate whether our identification method is reliable, we take two approaches. First, we examine intraday quote data from TAQ. For all 2015 trades that we can sign using our approach, we compare our buy-sell assignment to the trade sign from the Lee and Ready algorithm, and we find that the trade signs match for 89.9% of the observations. In the second approach, we put in a strict filter that requires the average effective spread from the previous month to be narrower than one cent, and re-examine our results. For stocks with a one cent spread, our trade-sign approach for subpenny-priced trades should match the Lee-Ready algorithm exactly and should be virtually error-free overall. This strict filter gets rid of more than 80% of the data, and we only keep the most liquid stocks in the sample. The results are in Panel I of Table VII. We find the retail order imbalance still significantly predicts next week's stock returns, with a coefficient of 0.0008, and a significant *t*-statistic of 4.48, consistent with our findings in Table III.

IV. Conclusions

In this paper, we exploit the fact that most retail order flows in U.S. equity markets are internalized or sold to wholesalers. As a part of this routing process, retail orders are typically given a small fraction of a penny per share of price improvement relative to the national best bid or offer price, and this price improvement can be observed when the trade is reported to the consolidated tape. Institutional orders almost never receive this kind of price improvement, so it becomes possible to use subpenny trade prices to identify a broad swath of marketable retail order flow. It is also straightforward to identify whether the retail trader is buying or selling stock: transactions at prices that are just above a round penny are classified as retail sales, while transactions that are just below a round penny are retail purchases.

We use this methodology to characterize the trading behavior and the information content of retail orders. We find that retail investors are on average contrarian, buying stocks that have experienced recent price declines and selling stocks that have risen in the past week. More significantly, we find that these investors are quite well-informed as a group. Over the next week, stocks with more positive retail order imbalances outperform stocks with relatively negative retail order imbalances by about 10 basis points, which is on the order of 5% per year annualized. This predictability extends out to about 12 weeks before dying off.

An important advantage of our method is that it is based on widely available intraday transaction data: anyone with access to TAQ can easily identify retail buys and sells using our approach. We believe there are many possible research applications. For example, researchers can investigate behavioral biases to see whether individual traders as a group exhibit these biases. It should also be possible to identify the nature of the information possessed by these retail investors (for example, whether retail investors are informed about future earnings news), along the lines of Boehmer et al. (2016). Another possible direction is to study the seasonality and time-series variation of retail trading, including tax-related and calendar-driven trading as well as activity around corporate events such as dividends, stock splits, and equity issuance.

References

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

Barber, Brad M., Terrance Odean, and Ning Zhu, 2009, Do retail trades move markets? *Review of Financial Studies* 22, 151−186.

Barrot, Jean-Noel, Ron Kaniel and David Alexandre Sraer, 2016, Are Retail Traders Compensated for Providing Liquidity? *Journal of Financial Economics* 120, 146-168.

Blume, Marshall E. and Robert F. Stambaugh, 1983, Biases in computed returns: An application to the size effect, *Journal of Financial Economics,* 12, 387-404.

Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang, 2008. Which shorts are informed? Journal of Finance 63, 491-527.

Boehmer, Ekkehart, Charles M. Jones, Yuan Wu, and Xiaoyan Zhang, 2016. What Do Short Sellers Know? Working paper, Columbia.

Campbell, John Y., Tarun Ramadorai, and Allie Schwartz, 2009, Caught on tape: Institutional trading, stock returns, and earnings announcements, *Journal of Financial Economics*, 92(1), 66-91.

Chordia, Tarun, and Avanidhar Subrahmanyam, 2004, Order imbalance and stock returns: Theory and evidence, *Journal of Financial Economics* 72, 485−518.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607−636.

Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual investor sentiment and stock returns, *Journal of Finance* 63, 273−310.

Kelley, Eric K. and Paul C. Tetlock, 2013, How Wise Are Crowds? Insights from Retail Orders and Stock Returns, *Journal of Finance* 68, 1229-1265.

Lee, Charles M. and Balkrishna Radhakrishna, 2000, Inferring investor behavior: Evidence from TORQ data, *Journal of Financial Markets*, 3(2), 83-111.

Lee, Charles M.C. and Mark J. Ready, 1991, Inferring investor behavior from intraday data, *Journal of Finance*, 46(2), 733-746.

Menkveld, Albert J., Bart Z. Yueshen, and Haoxiang Zhu, 2017, Shades of darkness: A pecking order of trading venues, *Journal of Financial Economics*, 124, 503-534.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

O'Hara, Maureen, Chen Yao, and Mao Ye, 2014, What's not there: Odd lots and market data, *Journal of Finance*, 69(5), 2199-2236.

Table I. Summary Statistics

This table reports summary statistics of retail investor trading activity. Our sample period is from January 2010 to December 2015, and our sample firms are common stocks listed on all U.S. stock exchanges with a share price of at least $1. Across all stocks and all days, we report the pooled sample mean for number of shares traded (*vol*), retail buy volumes (*indbvol*), retail sell volumes (*indsvol*), number of trades (*trd*), retail buy trades (*indbtrd*), retail sell trades (*indstrd*), as well as their odd lot counterparts (prefix *odd*). Odd lot measures are available starting at the end of 2013. In this paper, we include odd lot-related data starting January 2014. We compute order imbalance measures (variables containing *oib*) as in equation (1) to (4).

	N	Mean	Std	Median	Q1	Q3
Round lots and odd lots						
vol	4628957	1,229,004	6,849,849	221,234	51,768	819,615
trd	4628957	5,917	13,909	1,505	312	5,502
indbvol	4628957	42,481	280,474	5,165	1,200	20,681
indsvol	4628957	42,430	264,704	5,635	1,369	21,828
indbtrd	4628957	110	410	22	5	79
indstrd	4628957	108	355	24	6	81
oibvol	4628957	-0.038	0.464	-0.027	-0.301	0.217
oibtrd	4628957	-0.032	0.437	-0.010	-0.276	0.205
Odd lots only						
oddvol	1446749	6,561	20,141	1,811	629	5,250
oddtrd	1446749	222	669	64	21	186
oddindbvol	1446749	1,108	5,054	211	58	690
oddindsvol	1446749	968	3,488	210	62	663
oddindbtrd	1446749	37	171	7	2	23
oddindstrd	1446749	33	114	7	2	23
oddoibvol	1446749	-0.004	0.559	0.014	-0.338	0.331
oddoibtrd	1446749	-0.017	0.506	0.000	-0.290	0.250

Table II. Determinants of Retail Investor Order Imbalances

This table reports determinants of retail investor trading activity. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least \$1. We estimate Fama-MacBeth regressions, specified in equation (5). The dependent variables are two order imbalance measures: *oibvol* (number of shares traded) and *oibtrd* (number of trades). As independent variables, we include previous week return, *ret*(*w*-1), previous month return, *ret*(*m*-1), and previous 6-month return, *ret*(*m*-7, *m*-2). For the weekly returns, we compute it in two ways, using end of day bid-ask average price or using CRSP closing price. The control variables are monthly turnover (*lmto*), monthly volatility of daily returns (*lvol*), log market cap (*size*) and log book to market ratio (*lbm*), all measured at the end of previous month. To account for serial correlation in the coefficients, the standard deviations of the time-series are adjusted using Newey-West (1987) with 5 lags.

reg	I		II		III		IV	
Dep.var	oibvol		oibvol		oibtrd		oibtrd	
return	Bid-ask return		CRSP return		Bid-ask return		CRSP return	
	Coef.	t-stat	Coef.	t-stat	Coef.	t-stat	Coef.	t-stat
Intercept	-0.4013	-21.19	-0.4065	-21.35	-0.4326	-23.19	-0.4357	-23.19
oib(w-1)	0.2200	99.34	0.2201	99.38	0.2865	158.97	0.2866	159.04
Ret (w-1)	-0.9481	-42.39	-0.9620	-43.24	-0.9003	-37.66	-0.9156	-38.50
Ret (m-1)	-0.2778	-20.39	-0.2784	-20.45	-0.2258	-15.75	-0.2262	-15.78
Ret (m-7, m-2)	-0.0586	-12.10	-0.0584	-12.07	-0.0380	-6.85	-0.0378	-6.83
lmto	0.0003	5.59	0.0003	5.46	0.0002	4.12	0.0002	4.02
lvol	0.8100	8.75	0.8478	9.20	0.4366	4.44	0.4633	4.72
size	0.0154	12.76	0.0157	13.03	0.0209	17.30	0.0211	17.41
lbm	-0.0275	-18.52	-0.0274	-18.46	-0.0274	-18.84	-0.0273	-18.80

Table III. Predicting Next Week Returns Using Retail Order Imbalances

This table reports estimation results on whether retail investor trading activity can predict one week ahead returns. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. We estimate Fama-MacBeth regressions, specified in equation (6). The dependent variable is weekly returns, computed in two ways, using end of day bid-ask average price or using CRSP closing price. The independent variables are two order imbalance measures: *oibvol* (number of shares traded) and *oibtrd* (number of trades). As independent variables, we include previous week return, *ret*(*w*-1), previous month return, *ret*(*m*-1), and previous 6-month return, *ret*(*m*-7, *m*-2). The control variables are log book to market ratio (*lbm*), log market cap (*size*), monthly turnover (*lmto*), and monthly volatility of daily returns (*lvol*), all measured at the end of previous month. To account for serial correlation in the coefficients, the standard deviations of the time-series are adjusted using Newey-West (1987) with 5 lags.

reg	I		II		III		IV	
Order imbalance	oibvol		oibvol		oibtrd		oibtrd	
Dep. var	Bidask return		CRSP return		Bidask return		CRSP return	
	Coef.	t-stat	Coef.	t-stat	Coef.	t-stat	Coef.	t-stat
Intercept	0.0050	2.58	0.0056	2.85	0.0050	2.58	0.0056	2.85
Oib(w-1)	0.0009	15.60	0.0010	16.29	0.0008	12.30	0.0008	13.20
Ret (w-1)	-0.0185	-5.83	-0.0220	-6.85	-0.0186	-5.88	-0.0222	-6.91
Ret (m-1)	0.0006	0.35	0.0006	0.34	0.0005	0.29	0.0005	0.29
Ret (m-7, m-2)	0.0008	1.16	0.0008	1.16	0.0008	1.12	0.0008	1.12
lmto	0.0000	-3.37	0.0000	-3.76	0.0000	-3.36	0.0000	-3.75
lvol	-0.0223	-1.41	-0.0205	-1.31	-0.0217	-1.37	-0.0198	-1.27
size	-0.0001	-0.86	-0.0001	-0.92	-0.0001	-0.90	-0.0001	-0.96
lbm	-0.0001	-0.39	0.0000	-0.07	-0.0001	-0.42	0.0000	-0.10
Interquartile	1.1888		1.1888		1.2292		1.2292	
Return diff	0.1089%		0.1144%		0.0931%		0.0997%	

Table IV. Retail Return Predictability within Subgroups

This table reports whether retail investor trading activity can predict the cross-section of returns for a subset of stocks. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. We first sort all firms into 3 groups based on previous month-end characteristics. Then we estimate Fama-MacBeth regressions, specified in equation (6), for each subgroup. The dependent variable is weekly returns, computed using end-of-day bid-ask average price. The independent variables are two order imbalance measures: *oibvol* (number of shares traded) and *oibtrd* (number of trades). To account for serial correlation in the coefficients, the standard deviations of the time-series are adjusted using Newey-West (1987) with 5 lags. For each regression, we also provide the interquartile range for the relevant explanatory order imbalance along with the difference in predicted week-ahead returns for observations at the two ends of the interquartile range. Control variables are the same as in Table 3; those coefficients are not reported.

Panel A. Market cap groups

Oib measure	oibvol				oibtrd			
Mkt cap	coef.	t-stat	interquartile	weekly return diff	coef.	t-stat	interquartile	weekly return diff
small	0.0013	13.90	1.662	0.219%	0.0012	11.58	1.736	0.207%
medium	0.0007	9.18	1.323	0.087%	0.0004	5.63	1.346	0.059%
big	0.0003	3.68	0.892	0.026%	0.0002	2.52	0.929	0.019%

Panel B. Share price groups

Oib measure	oibvol				oibtrd			
price groups	coef.	t-stat	interquartile	weekly return diff	coef.	t-stat	interquartile	weekly return diff
low	0.0014	13.34	1.432	0.205%	0.0012	10.34	1.586	0.185%
medium	0.0007	10.00	1.289	0.089%	0.0005	7.56	1.309	0.070%
high	0.0002	3.23	0.961	0.020%	0.0002	2.19	0.961	0.015%

Panel C. Turnover groups

Oib measure	oibvol				oibtrd			
turnover groups	coef.	t-stat	interquartile	weekly return diff	coef.	t-stat	interquartile	weekly return diff
low	0.0011	15.60	1.837	0.205%	0.0011	14.71	1.777	0.195%
medium	0.0008	10.21	1.219	0.094%	0.0006	7.05	1.228	0.071%
high	0.0007	4.98	0.910	0.065%	0.0004	2.55	1.005	0.037%

Table V. Predicting Returns k-weeks Ahead

This table reports estimation results on whether retail investor trading activity can predict individual stock returns in future weeks. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. We estimate Fama-MacBeth regressions, specified in equation (6). The dependent variable is n-week ahead weekly returns, computed in two ways, using end of day bid-ask average price (Panel A) or using CRSP closing price (Panel B). The independent variables are two retail order imbalance measures, *oibvol* (number of shares traded), and *oibtrd* (number of trades), respectively. To account for serial correlation in the coefficients, the standard deviations of the time-series are adjusted using Newey-West (1987) with 5 lags. Control variables are the same as in Table 3; those coefficients are not reported.

Panel A. predict bid-ask average return k weeks ahead

	oibvol		oibtrd	
# of week ahead	coef.	t-stat	coef.	t-stat
1 week	0.00092	15.60	0.00076	12.30
2 weeks	0.00055	9.35	0.00048	7.89
4 weeks	0.00031	5.56	0.00026	4.66
6 weeks	0.00022	3.90	0.00015	2.60
8 weeks	0.00021	3.47	0.00011	1.75
10 weeks	0.00010	1.82	0.00002	0.35
12 weeks	0.00007	1.29	0.00009	1.52

Panel B. predict CRSP return k weeks ahead

	oibvol		oibtrd	
# of week ahead	coef.	t-stat	coef.	t-stat
1 week	0.00096	16.29	0.00081	13.20
2 weeks	0.00058	9.99	0.00052	8.57
4 weeks	0.00032	5.92	0.00028	5.05
6 weeks	0.00024	4.18	0.00017	2.93
8 weeks	0.00021	3.50	0.00011	1.80
10 weeks	0.00011	2.04	0.00005	0.81
12 weeks	0.00008	1.39	0.00010	1.76

Table VI. Long-short strategy returns based on retail order imbalances

This table reports portfolio returns using a long-short strategy where we buy the stocks in the highest quintile of scaled order imbalance, and we short the stocks in the lowest order imbalance quintile. The order imbalance is computed from the previous week. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. Portfolio returns are value-weighted, and market cap terciles are based on the previous month-end market cap. Because the holding period can be as long as 12 weeks, we report both the raw returns and risk-adjusted returns using the Fama-French three-factor model. Given the data we use is overlapping, we adjust the standard deviations of the portfolio return time-series using Newey-West (1987) with corresponding lags.

Panel A. Form portfolios on previous week retail order imbalance based on number of shares traded

Holding	whole sample				small		medium		big	
period	mean	t-stat	alpha	t-stat	alpha	t-stat	alpha	t-stat	alpha	t-stat
1 week	0.104%	3.28	0.101%	3.09	0.472%	11.84	0.179%	7.73	0.071%	2.16
2 weeks	0.160%	2.93	0.151%	2.90	0.734%	10.70	0.304%	7.60	0.105%	2.00
4 weeks	0.262%	2.54	0.254%	2.80	1.163%	11.04	0.456%	6.65	0.174%	1.88
6 weeks	0.262%	2.35	0.232%	2.37	1.182%	11.59	0.462%	6.25	0.150%	1.53
8 weeks	0.508%	2.61	0.517%	2.97	1.779%	12.69	0.574%	4.63	0.393%	2.51
10 weeks	0.569%	2.45	0.520%	2.21	1.964%	10.77	0.545%	3.94	0.443%	2.14
12 weeks	0.617%	2.35	0.629%	2.14	2.266%	9.57	0.459%	2.53	0.515%	1.90

Panel B. Form portfolios on previous week retail order imbalance based on number of trades

Holding	whole sample				small		medium		big	
period	mean	t-stat	alpha	t-stat	alpha	t-stat	alpha	t-stat	alpha	t-stat
1 week	0.074%	2.14	0.081%	2.27	0.379%	8.73	0.103%	4.24	0.064%	1.90
2 weeks	0.145%	2.20	0.156%	2.48	0.602%	8.01	0.165%	4.11	0.133%	2.20
4 weeks	0.236%	1.95	0.256%	2.29	0.993%	8.60	0.264%	3.93	0.228%	1.96
6 weeks	0.236%	1.82	0.191%	1.48	1.015%	8.13	0.275%	4.12	0.153%	1.15
8 weeks	0.453%	2.05	0.543%	2.82	1.615%	6.62	0.308%	2.33	0.461%	2.57
10 weeks	0.456%	1.68	0.515%	1.93	1.733%	6.04	0.248%	1.59	0.472%	1.88
12 weeks	0.490%	1.56	0.635%	1.83	2.149%	6.27	0.182%	0.93	0.563%	1.70

Table VII. Predicting aggregate market returns and additional analysis

Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. Standard errors are calculated using Newey-West (1987). In Panel A, we estimate equation (8). The dependent variable is the n-week ahead weekly value-weighted market return. The independent variables are two retail order imbalance measures, *oibvol* (number of shares traded), and *oibtrd* (number of trade), respectively. For all other panels, the regression is specified in equation (6), and estimated using Fama-MacBeth regressions. In Panel B, the dependent variable is weekly returns on approximately 1000 ETFs. In Panel C, we estimate the coefficients for different VIX regimes. In Panel D, the independent variables are two odd lot retail order imbalance measures, *oddoibvol* (number of odd lot shares traded), and *oddoibtrd* (number of odd lot trades), respectively. In Panel E, we estimate the coefficients for different bins of price improvements. The dependent variable is weekly returns, computed in two ways, using end-of-day bid-ask average price or using CRSP closing price. The independent variables are two retail order imbalance measures, *oibvol* (number of shares traded), and *oibtrd* (number of trades), respectively. Control variables for the cross-sectional regressions are the same as in Table 3, except that we do not include a book-to-market variable in the ETF regression; those coefficients are not reported.

Panel A. predicting future n-week market return

Weights Horizon	oibvol value weight coef.	t-stat	oibvol equal weight coef.	t-stat	oibtrd value weight coef.	t-stat	oibtrd equal weight coef.	t-stat
1 week	0.0037	0.50	-0.0053	-0.57	0.0054	0.92	-0.0038	-0.46
2 weeks	0.0101	0.79	-0.0030	-0.20	0.0120	1.21	0.0007	0.06
4 weeks	0.0044	0.20	-0.0236	-1.04	0.0073	0.43	-0.0136	-0.63
6 weeks	-0.0061	-0.22	-0.0356	-1.25	0.0022	0.10	-0.0216	-0.80
8 weeks	0.0075	0.20	-0.0046	-0.10	0.0118	0.41	0.0044	0.11
10 weeks	0.0051	0.11	-0.0114	-0.23	0.0101	0.28	-0.0038	-0.08
12 weeks	-0.0059	-0.10	-0.0315	-0.58	0.0000	0.00	-0.0227	-0.46

Panel B. Using retail oib to predict ETF returns

Order imbalance	oibvol		oibtrd	
Dep. Var	Bidask return		Bidask return	
	Coef.	t-stat	Coef.	t-stat
All ETFs	0.0001	2.04	0.0001	1.68
Interquartile	1.4726		1.4737	
Return diff	0.0153%		0.0118%	
Broad market ETFs	-0.0004	-0.81	0.0005	1.52

Panel C. Different market conditions

	vix<=18%		vix>18%	
Dep. var	bidaskret		bidaskret	
Indep. var	coef.	t-stat	coef.	t-stat
oibvol	0.0009	13.49	0.0010	9.36
oibtrd	0.0007	10.32	0.0008	7.60

Panel D. Predicting stock returns using odd-lot order imbalances

Dep.var	Oibvol		oibtrd	
return	Bid-ask return		Bid-ask return	
	Coef.	t-stat	Coef.	t-stat
Odd lot	0.0001	1.41	0.0001	0.77
Interquartile	1.2734		1.1314	
Return diff	0.0154%		0.0086%	

Panel E. Different retail trade sizes

Dep.var	Oibvol		oibtrd	
return	Bid-ask return		Bid-ask return	
	Coef.	t-stat	Coef.	t-stat
Small trades (< 400 shares)	0.0004	5.77	0.0004	4.48
Large trades (≥ 400 shares)	0.0009	7.25	0.0008	5.85

Panel F. Different price improvement amounts

Order imbalance	oibvol		oibtrd	
Dep. var	Bidask return		Bidask return	
	Coef.	t-stat	Coef.	t-stat
less price improvement	0.00071	9.30	0.00042	5.57
more price improvement	0.00021	3.04	0.00018	2.43

Panel G. Earnings surprises

Order imbalance	oibvol		oibtrd	
Dep. Var	Bidask return		Bidask return	
	Coef.	t-stat	Coef.	t-stat
oib	0.0003	8.16	0.0004	11.98
oib* eventday	0.0003	1.47	0.0002	1.31

Panel H. Retail vs. overall order imbalance

Order imbalance	overall oib		oibvol		oibtrd	
Dep. var	Bidask return		Bidask return		Bidask return	
	Coef.	t-stat	Coef.	t-stat	Coef.	t-stat
Retail Oib			0.0011	6.14	0.0006	3.33
Overall oib	0.0004	3.32	0.0000	0.10	0.0001	0.51

Panel I. When effective spread is less than one cent

Order imbalance	oibvol		oibtrd	
Dep. var	Bidask return		Bidask return	
	Coef.	t-stat	Coef.	t-stat
Oib	0.0008	4.48	0.0004	2.45

Figure 1. Distribution of Trade Size and Subpenny Prices for Retail Orders

These figures report summary statistics for the retail investor trading we identify. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. In Panel A, we compute the trade size in dollars as the number of shares multiplied by transaction price, and for each year we report the cross-sectional median, q1 (25th percentile) and q3 (75th percentile). In Panel B, we separate trades into 12 groups based on subpenny increments: trades at the whole penny, at the half penny, and in buckets that are 0.1 cent wide. We report the cross-sectional median of the daily number of shares traded in each group.

Panel A. retail order trade size in dollars

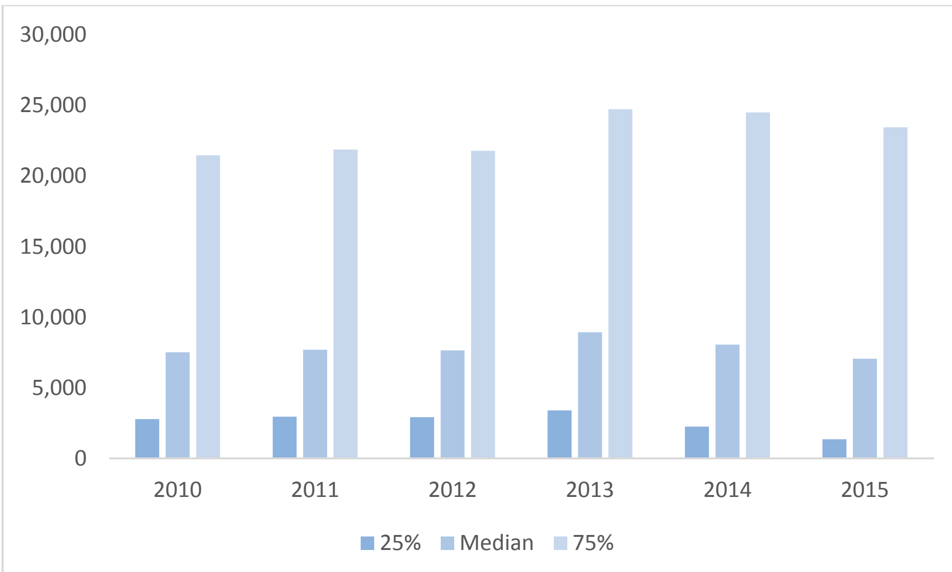

Panel B. median share volumes for different subpenny groups

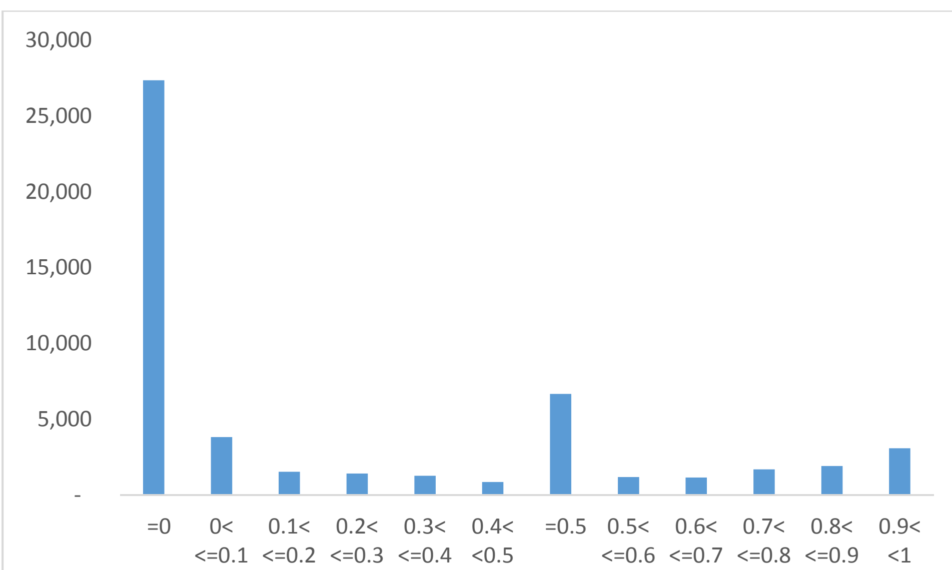

Figure 2. Time series of retail investor order imbalances

These figures report time series statistics of retail investor trading activities. Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least \$1. We present cross-sectional mean, median, q1 (25^{th} percentile) and q3 (75^{th} percentile) each day.

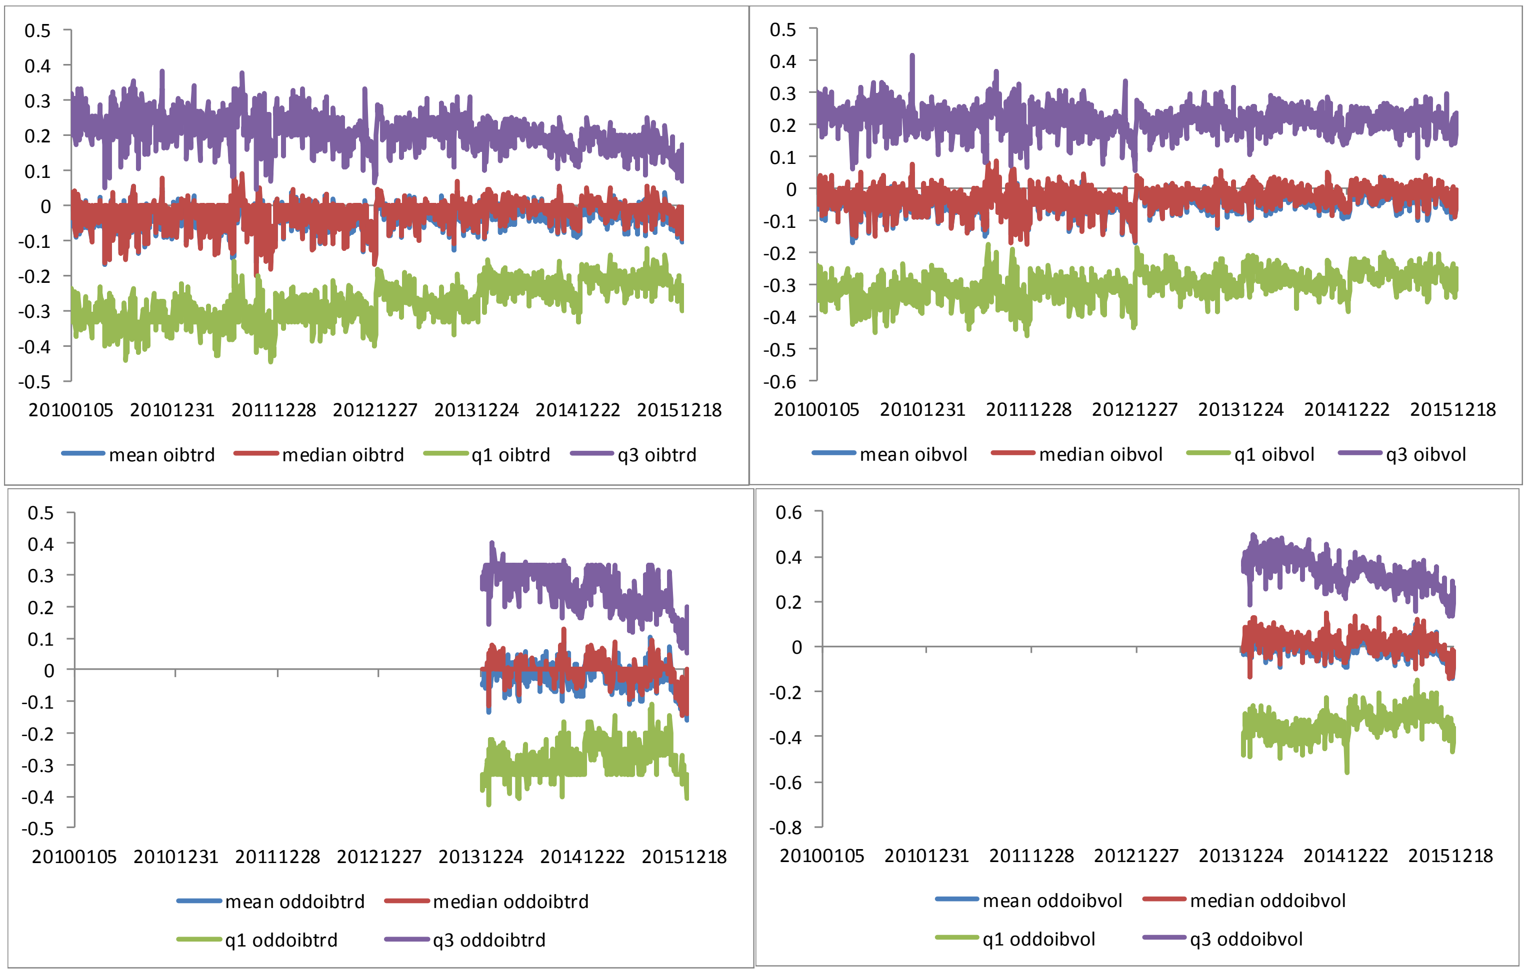

Figure 3. Predicting weekly returns n-weeks ahead, Fama-MacBeth regression coefficients

These figures plot the Fama-MacBeth coefficients on retail order imbalance measures in regression (6). Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least \$1. The dependent variable is weekly returns *n* weeks ahead, computed in two ways, using end of day bid-ask average price (*bidaskret*) or using CRSP closing price (*crspret*). The main independent variables are two retail order imbalance measures: *oibvol* (number of shares traded) and *oibtrd* (number of trades), respectively.

Panel A. Using retail order imbalance from number of shares traded (*oibvol*)

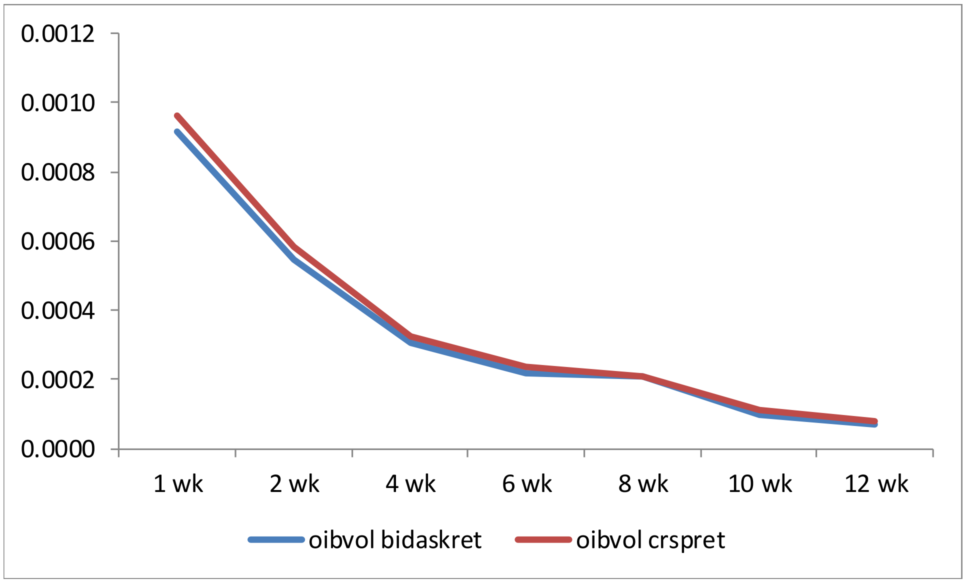

Panel B. Using retail order imbalance from number of trades (*oibtrd*)

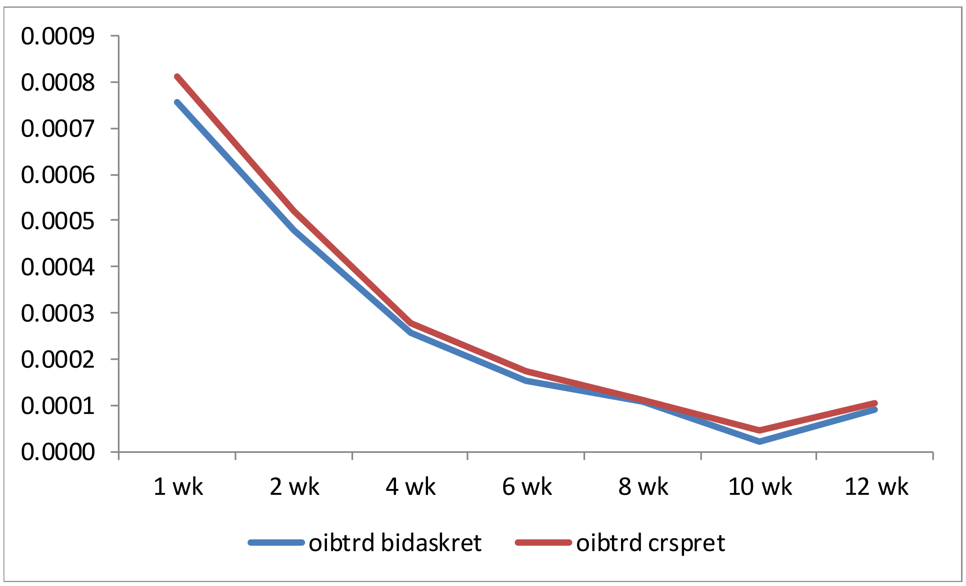

Figure 4. Portfolio return difference using previous week retail order imbalance

These figures plot weekly value-weighted portfolio return differences between quintile 5 and quintile 1, where stocks are sorted on the previous week retail order imbalance calculated using the number of shares traded (*oibvol*). Our sample period is from January 2010 to December 2015, and our sample firms are all common stocks listed on U.S. stock exchanges with a share price of at least $1. The portfolio returns are computed using the end of day bid-ask average price (*bidaskret*) in Panel A and the CRSP closing price (*crspret*) in Panel B.

Panel A. Weekly portfolio return difference using end-of-day bid-ask average prices

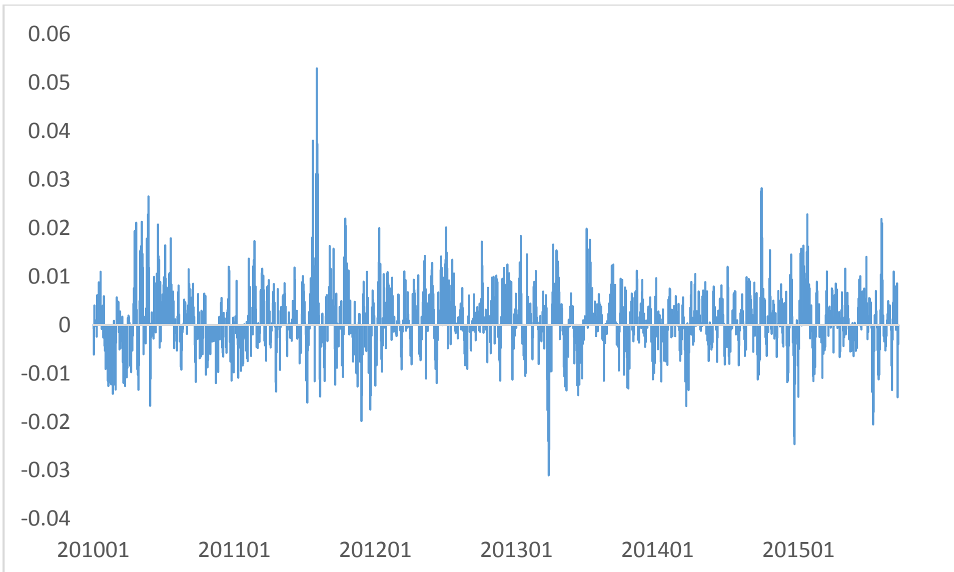

Panel B. Weekly portfolio return difference using CRSP closing prices