Buying the Dip in Retail Trading

Xiao Yin and Dongchen Zou *

July 2026

Abstract

Retail investors have become a major force in U.S. equity markets. Their share of trading volume has tripled over the past decade, to about a third of the total. Using a comprehensive dataset that covers approximately 80 percent of actual U.S. retail trading, we document that retail investors buy the dip: retail flow rises significantly after price declines but is roughly flat after gains. The asymmetry holds across return horizons and is persistent across the sample period. A structural model with sparse attention to historical prices and a short-sale constraint reproduces the asymmetry. The model predicts that doubling the retail share increases return volatility from 19% to 26%. Despite its contrarian nature, buying the dip disproportionately amplifies volatility during market declines.

JEL codes: G11, G12, G14

Keywords: retail trading, buying the dip, retail order flow

*Yin: University College London, xiao.yin@ucl.ac.uk; Zou: Indiana University, dongzou@iu.edu. We thank Xavier Gabaix, Nick Roussanov, and seminar participants at the Chicago Booth RP Reunion conference, Deakin University, and Australian National University for helpful comments and suggestions. All errors are our own.

1 Introduction

Retail trading activity in U.S. equity markets has tripled over the past decade. Aggregate retail volume rose from less than 10 percent of total market volume in 2017 to over 30 percent in 2021 (Figure 1).[1] Since the meme-stock episode of early 2021, sustained public and academic attention has turned to retail trading. Yet the prior evidence is divided on a fundamental question: do retail investors lean against returns or chase them? How do retail investors trade in response to market movements, and what does that behavior imply for the rest of the market now that retail participation is substantial?

We answer these questions using a comprehensive dataset of daily retail order flows for individual U.S. equities. The series is constructed from actual retail transactions rather than inferred from microstructure signatures in public tape data. The dataset captures approximately 80 percent of total gross U.S. retail trading activity over the sample period from January 2017 through December 2024.

Using the data, we document a novel and robust stylized fact: retail investors systematically buy the dip. Retail order flow rises markedly following negative returns, but is roughly flat following positive returns. The asymmetric response is robust across return horizons. A piecewise specification that separates the slope of retail flow on positive and negative past returns yields loss-side coefficients roughly an order of magnitude larger in absolute value than their gain-side counterparts at every horizon from one month to twelve months. The pattern is also stable over time. In a tercile sort on the past twelve-month return, average retail flow into the bottom tercile exceeds that into the top tercile in 97 percent of the monthly sample, and the autocorrelation of bottom-tercile flow is 0.7. The behavior persists across the COVID-19 selloff and the meme-stock episode in 2021.

The dip-buying pattern is a specific feature of retail flow. We run the same exercise on four non-retail flow proxies: a large-trade institutional proxy, the change in 13F institutional ownership, the change in lendable supply, and total initiated order

[1] This trend is consistent with industry estimates. J.P. Morgan Asset Management (2025) reports a retail share of approximately 36 percent of U.S. equity market order flow by April 2025.

flow. On the loss side, all these measures slope upward: as past returns fall, their flow falls with them. In other words, non-retail investors are "selling the dip".

To examine the market implications of buying the dip, we develop a structural model. Two features of the dip-buying pattern discipline the model: the asymmetry across loss and gain states, and the response to signed past returns at multiple horizons. In our model, these features arise from two ingredients: sparse attention to a historical price "anchor" in the spirit of Gabaix (2014), and a binding short-sale constraint on retail trading. Retail investors evaluate the asset against the anchored prices and expect transitory components of prices to mean-revert. This generates contrarian demand at every horizon over which the displacement between the current and anchored prices remains. While after sustained gains, the short-sale constraint binds for a substantial fraction of investors. This flattens the retail's contrarian response to recent positive returns and generates the asymmetry. Neither ingredient alone delivers the asymmetry: sparse attention without the constraint generates a usual contrarian schedule, while the constraint without sparse attention generates a flat schedule over past returns. We calibrate the model by simulated method of moments and match the scale of empirically observed returns and flows. The model successfully reproduces the retail flow asymmetry.

With the model disciplined to the dip-buying fact, we use it to ask what the behavior implies for prices as the retail share grows. The natural conjecture is that buying the dip stabilizes prices on the loss side. Retail investors buy the falling price precisely when institutions and other investors sell, which implies they supply liquidity in down states and absorb the pressure that drives prices down. One would expect this liquidity provision to cushion declines and dampen volatility, echoing the "stabilizing speculation" argument of Friedman (1953).

The model implies the opposite: a higher retail share raises return volatility. As the retail share increases from 25% to 50%, the return standard deviation rises from 19% to 26%. The key mechanism is the delayed timing of retail investors' response. Contrarian retail demand that arrived on impact would indeed offset the shock concurrently, as the conjecture supposes. But sparse attention to the lagged

price anchor delays retail's response until after the price has begun to mean-revert. Therefore, their trades do not offset the original move but instead introduce a second, opposite force that overcorrects the price. Furthermore, because retail buys the dip without constraint but pulls back only modestly after gains, the overcorrection is larger in loss states. Volatility therefore rises most after price-depressing shocks, which the model identifies with supply pressure.

Related Literature. This paper relates to three strands of literature. The first concerns the measurement of retail trading. Early work infers retail activity from trade size, treating small trades as retail and large trades as institutional (Lee and Radhakrishna, 2000; Hvidkjaer, 2008). Subsequent evidence shows that fixed trade-size cutoffs are fragile, because institutions split large orders and informed trading also appears in small trades (Campbell et al., 2009; Cready et al., 2014). Boehmer et al. (2021) introduce the BJZZ algorithm, which identifies retail trades from sub-penny price improvements granted by wholesalers. Barber et al. (2024) document signing errors in BJZZ and propose a modified measure, BHJOS. Battalio et al. (2023) show that sub-penny-flagged trades have low correlation with known retail imbalances. We contribute a retail flow measure built from actual retail executions and verify it against publicly released retail trading statistics during the GameStop episode.

The second strand concerns retail investor behavior. Prior work documents that retail trading is attention-driven (Barber and Odean, 2008; Barber et al., 2022) and contrarian at short horizons (Kaniel et al., 2008; Kelley and Tetlock, 2013). A large literature documents stable cross-sectional tilts toward high idiosyncratic volatility and positively skewed payoffs (Mitton and Vorkink, 2007; Kumar, 2009; Han and Kumar, 2013; Bali et al., 2011) and toward financially distressed, lottery-like stocks (Campbell et al., 2008; Conrad et al., 2014; Barberis and Huang, 2008). Recent studies find that retail risk-taking is state-contingent, but on variables other than the stock's own past return: low rates push investors toward high-income assets (Daniel et al., 2021), and underperformance against a salient benchmark prompts asymmetric risk-

taking (Gelman et al., 2026). Our dip-buying pattern differs in three respects. First, retail flow responds asymmetrically to past returns, rising steeply after losses but staying flat after gains, whereas a linear specification averages the two sides and labels retail simply as contrarian. Second, the pattern survives controls for lottery demand and firm distress, pointing to a distinct mechanism. Third, it persists at horizons up to twelve months, beyond the short horizons at which contrarian retail trading has been documented.

The third strand concerns retail liquidity provision and market outcomes. Kaniel et al. (2008) and Barrot et al. (2016) study whether retail investors are compensated for providing liquidity to other market participants. A separate literature studies how the composition of demand shapes the informativeness of prices when informed capacity is constrained. Yuan (2005) and Glebkin et al. (2021) show that the equilibrium price function becomes more sensitive to noise in constrained down states. Our structural model connects to these studies and identifies retail demand as a source of down-state fragility. Although retail trades against past losses, its response is delayed by one period, and we show that this delay raises return volatility rather than calming it. This differs from the positive-feedback channel of De Long et al. (1990), where destabilization arises from noise traders chasing the trend. Here retail leans against the trend, so volatility rises from the timing of contrarian demand rather than its direction. Importantly, retail's sparse attention to lagged anchored prices shapes the market outcomes, amplifying volatility precisely in the down states where prices are already most sensitive to noise.

2 Data

2.1 Retail Flow

The primary data source is the S&P Retail Flow Database. This dataset provides daily retail buy and sell volume for individual U.S. equities from January 2017 through December 2024. Unlike the TAQ-based approaches described below, the S&P data are constructed from actual retail executions rather than inferred from patterns in market microstructure data. The dataset is accessible through the S&P Global Marketplace.

Coverage is comprehensive. According to S&P Global, the database captures approximately 80 percent of total gross U.S. retail trading activity over the sample period. Figure 1 documents aggregate retail buy and sell volumes over time (Panel A), the retail share of total market volume (Panel B), and aggregate net retail order flow (Panel C). The figure shows that retail trading activity has roughly tripled over the sample. Retail volumes rose from less than 10 percent of market volume in 2017 to over 35 percent during the COVID-19 pandemic and the meme-stock episode of early 2021, receded from the peak, and stabilized at levels above the pre-2020 period.

This featured dataset addresses a tradeoff in the retail-trading literature between coverage and measurement. One strand uses proprietary brokerage records. These data are valuable because they observe actual investor trades, but they often come from one broker, one platform, or a narrow investor group. As a result, the same empirical design can speak to different retail populations: Robinhood users, traditional discount-brokerage customers, or clients of a particular international broker may trade for different reasons and have different market effects (Barber and Odean, 2008; Kumar and Lee, 2006; Barber et al., 2022; Eaton et al., 2022). A second strand uses public TAQ data to infer retail trades. These measures are easier to replicate, but they depend on microstructure signatures that can fail when wholesaler pricing, routing, or market conditions change (Boehmer et al., 2021; Battalio et al., 2023; Barber et al., 2024). The S&P data occupy the middle ground. They are based on actual retail executions, have broad market coverage, and are available to researchers through a commercial data platform. This makes them useful for studying aggregate retail behavior while preserving a common, replicable measurement standard.[2]

We supplement the retail flow data with Trade and Quote (TAQ) data from WRDS to construct benchmark flow measures. We merge in CRSP and Compustat for stock and firm characteristics. We also use securities lending data from S&P Global and 13-F institutional holdings from LSEG Thomson Reuters.

[2]To our knowledge, this is the first paper that examines the S&P Retail Flow dataset.

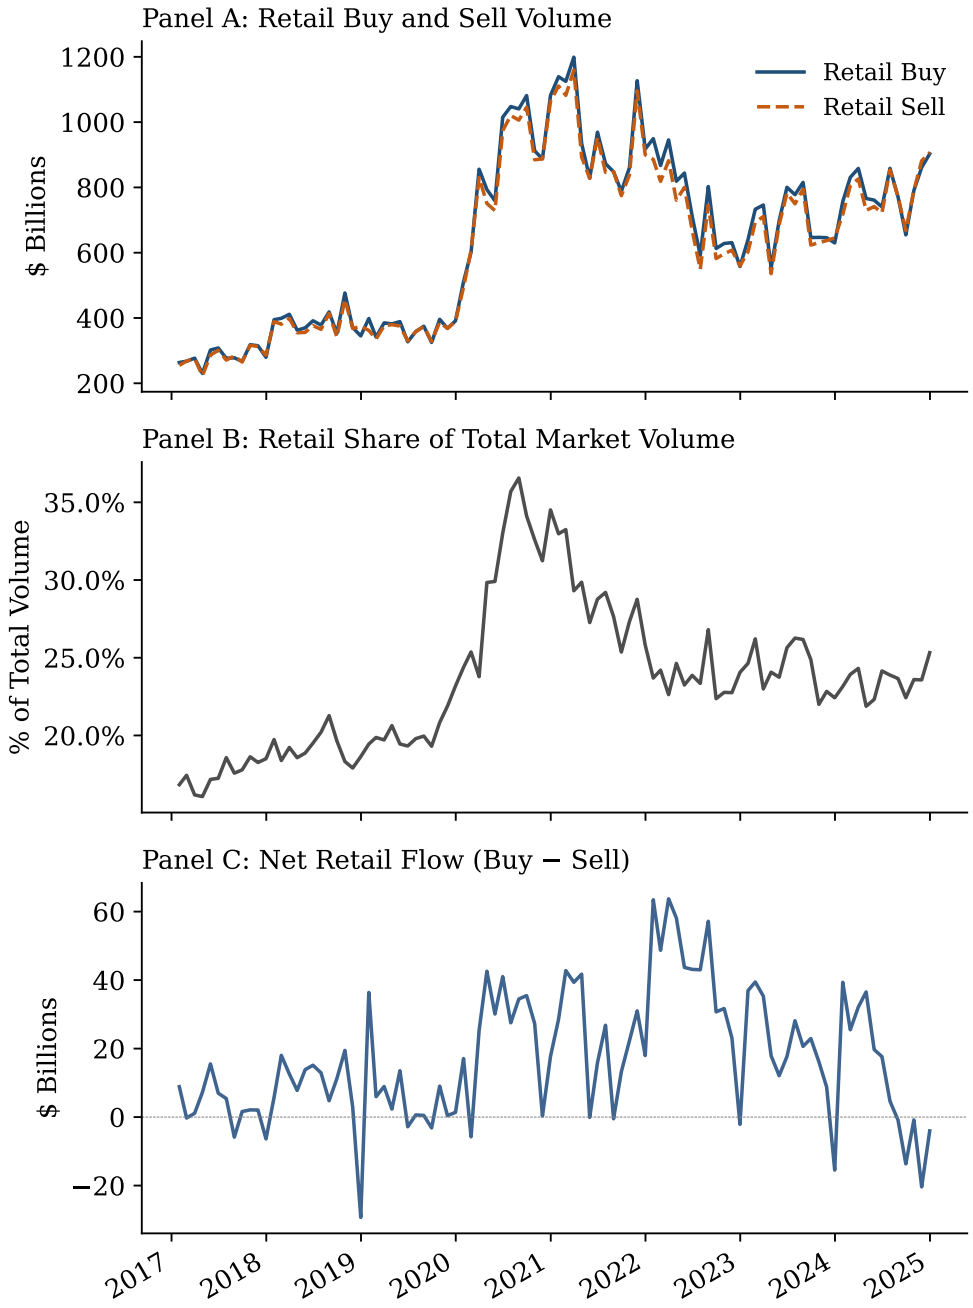


Figure 1. Retail Flows and Volumes. Panel A: aggregate monthly retail buy and sell volume ($ billions). Panel B: retail share of total market dollar volume (percent). Panel C: aggregate monthly net retail dollar order flow ($ billions), computed as buy volume minus sell volume. The sample covers January 2017 through December 2024. Retail flow is from the S&P Global Retail Flow Database.

2.2 Variable Definitions and Sample Construction

Our main variable, *Retail Flow*, is monthly retail buy volume minus retail sell volume from S&P Global. For aggregate time-series figures, we report either dollar values

in billions or percent of total market dollar volume. For stock-level analyses, we normalize the net flow by total shares outstanding.

We compare Retail Flow to five benchmark flow measures. Two are TAQ-based retail proxies. *BJZZ* is the sub-penny retail flow measure of Boehmer et al. (2021). A trade is classified as retail if it is an off-exchange or internalized execution and the execution price has a sub-penny component, $\text{cf}(p) = 100 \times \text{mod}(\text{p}, 0.01) \notin \{0, 1\}$, that is, the cent fraction of the price is neither zero nor exactly half a cent. Trades that pass this identification filter are signed by the sub-penny pattern of the price. Trades with $\text{cf}(p) \geq 0.6$ are classified as retail buys, trades with $\text{cf}(p) \leq 0.4$ are classified as retail sells, and trades with $\text{cf}(p) \in (0.4, 0.6)$ are excluded as ambiguous. Wholesalers internalize retail buys just below the round-cent ask (cent fraction near one) and retail sells just above the round-cent bid (cent fraction near zero).

Two limitations of the BJZZ measure have been documented in subsequent validation work. The first is identification: Battalio et al. (2023) compare BJZZ-flagged trades against a ground-truth sample of retail brokerage executions and report that BJZZ identifies fewer than one-third of actual retail trades and shows a low correlation with known retail imbalances. The second is signing. The cent-fraction rule aligns with the position of the trade between the bid and the ask only under the one-cent spread regime. When the spread is wider than one cent, the cent fraction of the price no longer reflects whether the trade sits in the buyer half or the seller half of the spread. Barber et al. (2024) highlight this spread-width sensitivity. As a stylized illustration, suppose the bid is \$1.98 and the ask is \$2.00, so the spread is two cents and the midpoint is \$1.99. An execution at \$1.997 has cent fraction 0.7 and sits 0.3 cents below the offer, in the upper half of the spread. An execution at \$1.987 also has cent fraction 0.7 but sits 0.7 cents above the bid, in the lower half of the spread. The BJZZ rule signs both trades identically as retail buys because they share the same cent fraction, even though the second trade is closer to the bid than to the ask and is more naturally classified as a retail sell. Barber et al. (2024) document the resulting signing-error rate in a controlled experiment with roughly 85,000 retail trades of known direction. The BJZZ cent-fraction rule produces signing errors on

approximately 28 percent of identified trades, while a position-relative-to-spread rule that uses the bid and the ask reduces signing errors to approximately 5 percent. Inferences about retail trading can be sensitive to these signing errors even when the identification step is held fixed.

BHJOS is the modified retail flow measure of Barber et al. (2024), which retains the BJZZ sub-penny identification filter but replaces the cent-fraction signing rule with a midpoint rule that explicitly uses the spread width. The trade-location statistic is $\text{loc} = (\text{price} - \text{nbb})/\text{spread}$, where nbb is the prevailing national best bid and spread is the prevailing NBBO spread. Trades with $\text{loc} > 0.6$ are classified as retail buys, trades with $\text{loc} < 0.4$ are classified as retail sells, and trades in $[0.4, 0.6]$ are excluded as ambiguous rather than forced into a buy or sell category. Returning to the stylized example above, the trade at \$1.997 has $\text{loc} = 0.85$ and is classified as a retail buy, while the trade at \$1.987 has $\text{loc} = 0.35$ and is reclassified from a BJZZ retail buy to a BHJOS retail sell. BHJOS also excludes trades with notional value above \$100,000 to remove probable institutional executions that occasionally satisfy the sub-penny identification filter.

The remaining measures capture non-retail or aggregate order flow. *Inst50k* is net flow in trades with notional value exceeding \$50,000. Trade-size rules were common in earlier work that inferred trader identity from transactions data (Lee and Radhakrishna, 2000; Hvidkjaer, 2008). Subsequent evidence shows that fixed trade-size cutoffs are fragile because institutions split orders and because small and large trades can both contain institutional information (Campbell et al., 2009; Cready et al., 2014). We therefore treat *Inst50k* as a benchmark, not as a definitive measure of institutional trading. *InstSec* is the monthly change in lendable supply divided by total shares outstanding, constructed from the securities-lending data. This measure serves as a proxy for institutional trades (Barardehi et al., 2024). ΔIO is the quarter-over-quarter change in the fraction of shares outstanding held by 13-F filers, carried forward within the quarter. *Total Flow* is total market net order flow.

We begin with all stocks in the Retail Flow Database that can be merged to CRSP using the PERMNO identifier. We aggregate daily flows to the stock-month level by

summing within each calendar month. We then apply three filters. First, we require the month-end stock price to be at least one dollar to exclude penny stocks. Second, we require the absolute value of normalized net retail flow (flow divided by shares outstanding) to be no greater than one. Third, we require positive market equity. This removes observations with invalid or nonpositive price-share information after the CRSP-Compustat merge. After filtering, the matched sample contains 417,519 stock-month observations covering 7,432 distinct stocks and 96 calendar months from January 2017 through December 2024. The average number of stocks per month is 4,349.

2.3 Summary Statistics

Table 1 reports summary statistics for the matched sample. Panel A describes the flow variables. All flow measures are normalized by shares outstanding and expressed in percent. The institutional ownership variable (ΔIO) is the quarter-over-quarter change in institutional ownership, also expressed as percent of shares outstanding.

Retail net flow has a positive mean of 0.23 percent of shares outstanding per month and a median near zero (0.01 percent). The distribution is right-skewed, with a standard deviation of 2.33 percent and a 90th percentile of 0.72 percent. The near-zero median reflects the fact that retail investors are roughly as likely to be net sellers as net buyers in a given stock-month. The positive mean is consistent with a slight secular tendency toward retail accumulation over the sample period.

The TAQ-based retail flow proxies (BJZZ and BHJOS) have smaller absolute means and medians close to zero, which is expected because they measure signed imbalances in a specific sub-penny subset of trades rather than total retail activity. The institutional large-trade proxy ($Inst50k$) has a negative mean of -0.13 percent of shares outstanding, consistent with institutional investors being net sellers on average over this period. The securities lending proxy ($InstSec$) and ΔIO are both near zero on average.

Panel B describes return variables. The average monthly stock return is 1.19 percent, with a standard deviation of 22.75 percent. Return volatility as a stock characteristic is measured using daily returns over a 21-trading-day window. Panel C

Table 1. Summary Statistics

This table reports summary statistics for the sample, which covers January 2017 through December 2024 (N = 417,519 stock-months, 7,432 stocks, 96 months). Panel A reports flow variables, each measured as monthly net order flow divided by shares outstanding (or by market equity for the Mkt Cap row) and multiplied by 100 (percent of denominator); ΔIO is the quarter-over-quarter change in institutional ownership from 13-F filings, expressed as percent of shares outstanding and carried forward within the quarter. Panel B reports return variables, where $\text{Ret}[t-k, t-1]$ is the cumulative return from month $t-k$ to $t-1$, and idiosyncratic volatility and idiosyncratic skewness are the standard deviation and skewness of the Fama-French three-factor residual returns over a 21-trading-day window. Panel C reports firm characteristics, with market equity in millions of dollars and firm age in months since first CRSP appearance.

Variable	N	Stocks	Mean	Std	p10	p25	p50	p75	p90
Panel A: Flow Variables (normalized × 100, percent of denominator)									
Retail Flow / Shares	417,519	7,432	0.232	2.325	−0.419	−0.126	0.010	0.207	0.715
Retail Flow / Mkt Cap	417,519	7,432	0.268	2.624	−0.413	−0.124	0.011	0.211	0.731
BHJOS Flow / Shares	417,519	7,432	0.045	1.379	−0.140	−0.051	−0.006	0.029	0.135
BJZZ Flow / Shares	417,519	7,432	0.045	1.601	−0.136	−0.043	−0.004	0.026	0.119
Inst50k Flow / Shares	417,519	7,432	−0.129	1.202	−0.561	−0.178	−0.001	0.074	0.296
InstSec Flow / Shares	417,519	7,432	0.149	1.337	−0.649	−0.151	0.023	0.390	1.028
Total Flow / Shares	417,519	7,432	−0.470	3.389	−1.543	−0.611	−0.140	0.161	0.650
ΔIO (qtrly, %)	388,635	7,236	0.004	36.045	−3.851	−1.293	0.102	1.802	5.024
Panel B: Return Variables									
$\text{Ret}[t]$	414,239	7,392	0.012	0.228	−0.161	−0.069	0.002	0.070	0.171
$\text{Ret}[t-2, t-1]$	406,946	7,342	0.013	0.233	−0.160	−0.068	0.002	0.071	0.172
$\text{Ret}[t-3, t-1]$	407,701	7,321	0.024	0.318	−0.225	−0.096	0.005	0.104	0.251
$\text{Ret}[t-6, t-1]$	398,213	7,229	0.054	0.507	−0.348	−0.153	0.012	0.176	0.416
$\text{Ret}[t-12, t-1]$	379,960	7,014	0.133	0.929	−0.488	−0.222	0.030	0.301	0.695
Idiosyncratic Volatility (21d)	413,237	7,402	0.027	0.033	0.008	0.012	0.020	0.033	0.051
Idiosyncratic Skewness (21d)	413,237	7,402	0.131	0.947	−0.888	−0.364	0.102	0.597	1.205
Panel C: Firm Characteristics									
Market Equity ($M)	417,519	7,432	9,374.9	59,941.8	50.7	175.3	743.1	3,596.2	15,069.6
Log Market Equity	417,519	7,432	6.703	2.189	3.925	5.167	6.611	8.188	9.620
Book-to-Market	371,684	6,498	0.728	1.313	0.103	0.231	0.492	0.875	1.399
Market Beta	320,045	5,687	1.197	0.714	0.417	0.752	1.120	1.548	2.074
Gross Profitability	384,238	6,550	2.088	554.985	0.000	0.043	0.194	0.380	0.604
ROA	391,408	6,677	−0.096	11.432	−0.432	−0.098	0.010	0.052	0.108
Debt-to-Asset	394,409	6,709	0.301	8.975	0.003	0.048	0.200	0.395	0.579
Firm Age (months)	417,519	7,432	280	237	41	85	238	396	617

describes firm characteristics. The average stock in the sample has market equity of $9.4 billion, and the median is $743 million. The average book-to-market ratio is 0.73.

2.4 Validation with Citadel Securities Statistics

We validate the S&P Retail Flow data against public statistics released by Citadel Securities in January 2021. Citadel Securities is one of the largest wholesale market makers, which intermediates a substantial share of U.S. retail order flow via payment for order flow (PFOF) arrangements with retail brokerages. In the aftermath of the GameStop (GME) short squeeze and trading restrictions, Citadel Securities faced heightened congressional and public scrutiny over its role as an intermediary for

Table 2. Citadel Securities: GME Retail Flows, January 25–28, 2021

This table reports daily GME retail order flow intermediated by Citadel Securities over January 25–28, 2021. Price is the daily closing price of GME from CRSP. Retail Buy and Retail Sell are the number of retail shares bought and sold through Citadel Securities on each day, and Net is their difference (Buy − Sell). Mkt. Volume is total GME market volume in shares. Mkt. Share is the retail volume intermediated by Citadel Securities divided by total GME market volume (percent). Net Flow is net retail shares divided by total market volume (percent). Source: Citadel Securities, published no later than January 29, 2021.

Date	Price ($)	Retail Buy	Retail Sell	Net (shares)	Mkt. Volume	Mkt. Share (%)	Net Flow (%)
2021-01-25	76.79	26,558,557	24,489,122	+2,069,435	177,874,000	28.7	+1.16
2021-01-26	147.98	24,888,375	26,794,942	−1,906,567	178,587,974	28.9	−1.07
2021-01-27	347.51	12,966,267	13,743,184	−776,917	93,396,666	28.6	−0.83
2021-01-28	193.60	9,972,227	10,078,110	−105,883	58,816,595	34.1	−0.18

retail order flow. To provide transparency about its intermediation activity during the episode, Citadel Securities published daily retail buy and sell volumes in shares for GME for January 25–28, 2021. These statistics were publicly available no later than January 29, 2021.

Table 2 reproduces the Citadel Securities statistics. Across the four trading days, retail investors routed 10 to 27 million shares per day through Citadel Securities, which intermediated 28.7 to 34.1 percent of total GME market volume. The episode is informative because the public narrative and the transaction data point in very different directions. Coordinated discussion on the Reddit forum and elsewhere urged retail investors to hold positions through the squeeze. Public commentary at the time treated retail demand as single-sided. The statistics show otherwise. Net retail flow through Citadel Securities was modestly positive on January 25 when buyers slightly outnumbered sellers, turned decidedly negative on January 26 and 27 as the price continued to climb, and returned to near balance on January 28 when the price retreated. On the squeeze days, retail behavior did not match the buy-and-hold rhetoric: a material share sold into the price increase.

Figure 2 compares our three flow measures—S&P Retail Flow, BJZZ, and BHJOS—against the rescaled Citadel Securities statistics over the broader event window (January 22 through February 5, 2021). Citadel Securities states that it intermediates approximately 39 percent of U.S. retail trading volume, while the S&P Retail Flow data captures approximately 80 percent of gross retail activity. We

therefore apply a rescaling factor of two ($80/39 \approx 2.05$) to the Citadel series to make it comparable in magnitude to the S&P series. We note that "Mkt. Share (%)" in Table 2 refers to Citadel Securities' intermediated retail volume relative to *total* GME market volume on each day (28.7–34.1 percent), which is a different denominator from the firm-wide retail-trading share used for the rescaling.

The S&P Retail Flow series matches the Citadel Securities statistics closely in both sign and magnitude on January 25–27. Both series show modest net retail buying on January 25, and a shift to net selling on January 26–27. The agreement in direction and scale across three independent trading days provides direct event-level validation that the S&P database captures retail order-flow imbalances in a salient high-volume episode.[3]

The TAQ-based measures tell a very different story. Both BJZZ and BHJOS are inconsistent in sign and magnitude with the reported statistics across the event window. The BJZZ series stays near zero throughout the short squeeze, and does not track either the net buying on January 25 or the net selling on January 26–27 that appear in both the S&P and the Citadel Securities data. The BHJOS series moves more than BJZZ, but its day-to-day signs also disagree with the statistics. After applying the same rescaling factor that aligns the S&P series, BHJOS sits at roughly one-half to one-third the size of the reported net flows. Both measures depend on a sub-penny identification step, which in turn depends on the wholesaler quoting and price-improvement patterns. Widening spreads and altered routing could render the sub-penny signature unreliable, especially during the meme-stock episode. The results are aligned with Battalio et al. (2023), who document that sub-penny-flagged trades show low correlation with known retail imbalances.

[3]The one notable discrepancy between the Net Retail Flow series and the Citadel Securities statistics occurs on January 28, 2021. On that day, Robinhood, Webull, and several other retail brokerages placed GME into position-closing-only (PCO) status, meaning retail customers could sell but not open new buy positions. Citadel Securities receives a large share of its retail order flow from Robinhood via PFOF. As a result, Citadel's data for January 28 is disproportionately short of retail buy orders relative to the broader retail universe. The S&P Retail Flow Database aggregates order flow across brokerages, including other brokerage companies that did not restrict buying on that day. Accordingly, the S&P series shows more retail buying on January 28 than the Citadel statistics imply. This discrepancy is a direct consequence of broker-specific restrictions rather than a data quality issue. Outside of this specific episode, the close alignment between the two series provides direct validation that the S&P database accurately measures retail order flow.

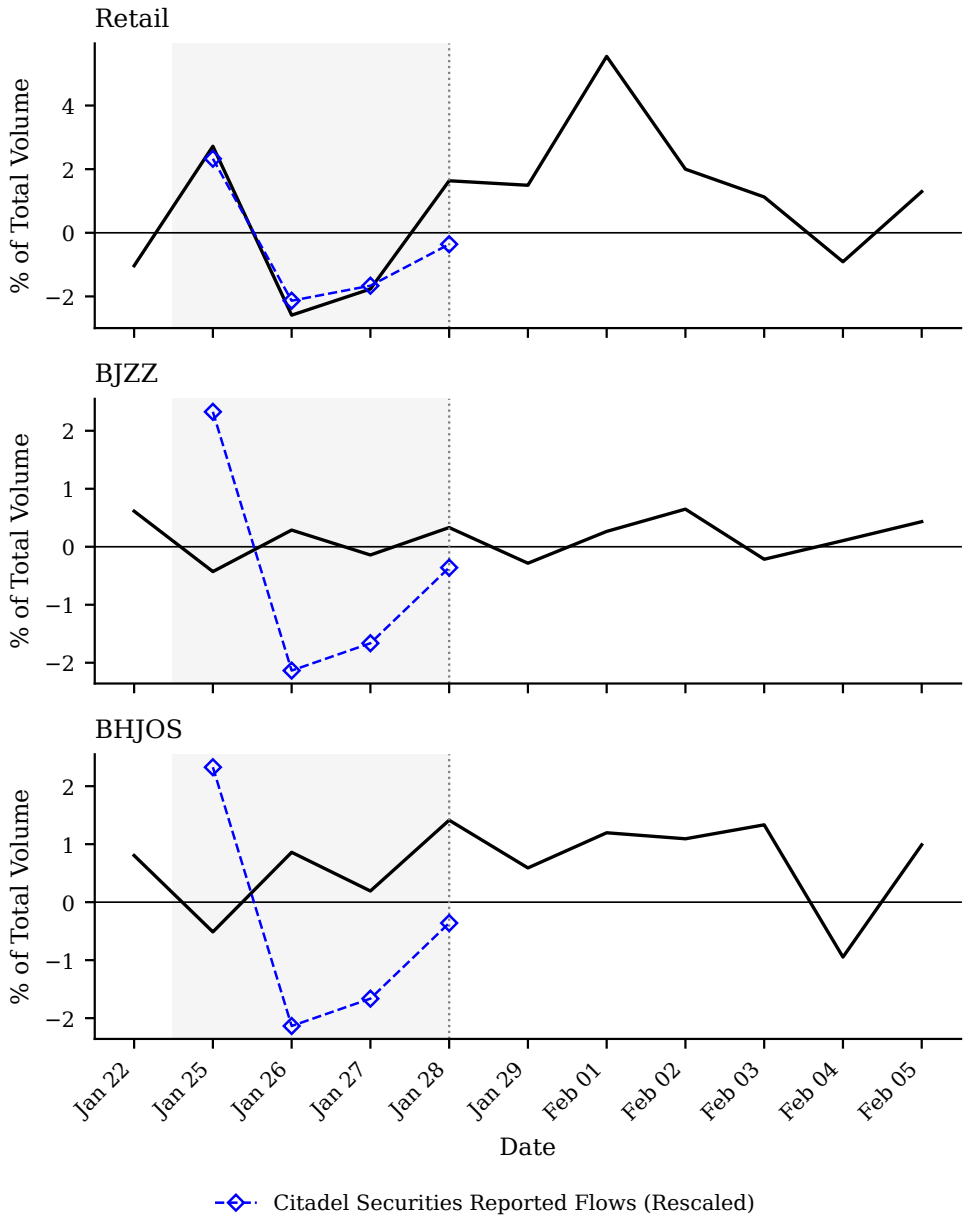


Figure 2. Net Retail Flows vs. Citadel Securities on GameStop. Each panel plots the daily net retail order flow as a percent of total GME market volume. *S&P Retail Flow*: actual retail flow from the S&P Global Retail Flow Database. *BJZZ*: sub-penny retail flow of Boehmer et al. (2021). *BHJOS*: modified retail flow of Barber et al. (2024). The dashed line with hollow markers shows the Citadel Securities statistics from Table 2, rescaled by a factor of 2 ($\approx$ 80%/39%). The shaded region spans January 25–28, 2021. The vertical dotted line marks January 28, 2021, when several retail brokerages imposed position-closing-only restrictions.

These failures occur during the episode that attracted the most public attention to retail trading, underscoring why inferred proxies can lead to materially misleading inferences about retail behavior in the cases researchers care most about.

Table 3. Correlations of Retail Flow with Other Flow Measures

This table reports time-series (TS) and average monthly cross-sectional (CS) correlations between S&P Retail net order flow and five alternative flow measures. *BJZZ*: sub-penny retail flow of Boehmer et al. (2021). *BHJOS*: modified retail flow of Barber et al. (2024). *Inst50k*: net flow in trades with notional value above $50,000, following the trade-size proxy tradition but subject to known classification concerns (Lee and Radhakrishna, 2000; Hvidkjaer, 2008; Campbell et al., 2009; Cready et al., 2014). *InstSec*: monthly change in lendable supply divided by shares outstanding (Barardehi et al., 2024). ΔIO: quarter-over-quarter change in 13-F institutional ownership fraction (carried forward within the quarter). *Total*: total market net order flow. TS correlations use monthly flows aggregated across stocks and scaled by total market dollar volume. CS correlations are averaged monthly Pearson correlations with flows normalized by shares outstanding. Panel A uses the full sample. Panels B and C sort on log market equity and prior 12-to-1 month return, respectively (Q1 = smallest/lowest). Panel D splits by sub-period: Pre-COVID (Jan 2017–Dec 2019), COVID/Meme (Jan 2020–Dec 2021), Post-meme (Jan 2022–Dec 2024).

	Time-Series Correlation						Cross-Sectional Correlation					
	BJZZ	BHJOS	Inst50k	InstSec	ΔIO	Total	BJZZ	BHJOS	Inst50k	InstSec	ΔIO	Total
Panel A: Full Sample												
Full sample	0.017	0.347	−0.092	−0.016	−0.274	−0.006	0.478	0.508	−0.009	0.053	−0.037	−0.092
Panel B: By Market Cap Quintile												
Q1 (Small)	0.272	0.306	−0.336	−0.033	−0.174	−0.164	0.493	0.511	−0.026	0.115	−0.064	−0.115
Q2	0.410	0.456	−0.399	0.056	0.221	−0.448	0.442	0.511	0.028	0.043	−0.058	0.065
Q3	0.216	0.537	−0.449	0.037	−0.036	−0.101	0.346	0.446	−0.015	0.081	−0.028	0.069
Q4	0.062	0.328	0.063	−0.051	−0.336	0.077	0.269	0.446	−0.011	0.048	−0.015	0.076
Q5 (Large)	−0.040	0.302	−0.028	−0.110	−0.110	0.028	0.156	0.455	−0.027	0.053	−0.005	0.053
Panel C: By Momentum Quintile (prior 12-to-1 month return)												
Q1 (Losers)	0.280	0.374	−0.487	0.039	−0.433	−0.335	0.509	0.537	−0.026	0.067	−0.079	−0.133
Q2	0.394	0.483	0.124	−0.249	−0.106	0.421	0.340	0.405	0.014	0.026	−0.049	−0.003
Q3	0.553	0.622	−0.116	−0.139	−0.193	−0.039	0.386	0.426	0.018	0.038	−0.029	0.059
Q4	0.554	0.516	−0.011	−0.144	−0.231	0.109	0.294	0.379	0.019	0.049	−0.031	0.007
Q5 (Winners)	0.102	0.446	−0.667	0.501	−0.237	−0.403	0.351	0.460	−0.016	0.066	−0.014	0.062
Panel D: By Sub-Period												
Pre-COVID (2017–19)	0.597	0.783	0.200	−0.104	−0.089	0.064	0.560	0.626	−0.006	0.026	−0.032	−0.050
COVID/Meme (2020–21)	0.497	0.449	0.375	0.244	−0.231	0.375	0.466	0.475	−0.020	0.070	−0.042	−0.030
Post-meme (2022–24)	−0.075	0.507	−0.803	−0.114	−0.277	−0.074	0.403	0.411	−0.004	0.069	−0.038	−0.176

2.5 Comparison with TAQ-Inferred Measures

We compare the S&P Retail Flow series to the five flow measures defined above. Table 3 reports time-series and cross-sectional correlations between S&P Retail Flow and each measure. Time-series correlations aggregate monthly flows across stocks, scale by total market dollar volume, and compute the Pearson correlation of the resulting monthly series. Cross-sectional correlations are the average monthly Pearson correlation across stocks, with all flows normalized by shares outstanding. Panels A through D of Table 3 report results for the full sample, by market capitalization quintile, by momentum quintile, and by sub-period, respectively.

In the full-sample time series (Panel A), retail flow is nearly uncorrelated with BJZZ ($\rho = 0.017$). The sub-penny identification algorithm fails to capture the aggregate timing of retail activity. The BHJOS modification improves this

substantially ($\rho = 0.347$), suggesting that the signing refinement recovers information about aggregate retail direction that BJZZ discards. Retail flow is essentially uncorrelated with total market flow in the time series ($\rho = -0.006$), confirming that retail and non-retail order flows are driven by distinct aggregate forces. The negative correlation with ΔIO ($\rho = -0.274$) is consistent with retail investors accumulating shares when institutional investors are reducing their positions at the aggregate level.

Cross-sectional correlations are significantly higher than time-series correlations for both TAQ proxies. In the full sample (Panel A), the average cross-sectional correlation between Retail Flow and BJZZ is 0.478; for BHJOS it is 0.508. These values confirm that the TAQ proxies capture meaningful stock-level variation in retail activity even when they miss aggregate timing. As a benchmark, Barber et al. (2024) report an R^2 of approximately 35 percent between BJZZ-identified flow and actual retail flow in their controlled experiment, implying a correlation of roughly 0.59. Our sample cross-sectional correlation of 0.478 falls somewhat below this level, in part because the S&P database draws from a broader population of retail trades and brokerages than the controlled sample in that study.

When do TAQ-based retail flow proxies work better? The conditional correlations in Panels B–D reveal systematic heterogeneity in proxy quality. Panel B shows that cross-sectional correlations are higher among small stocks ($\rho_{\text{BJZZ}} = 0.493$ for the smallest quintile) than among large stocks ($\rho_{\text{BJZZ}} = 0.156$ for the largest quintile). The BHJOS modification largely closes this size gap, with correlations ranging from 0.511 (Q1) to 0.455 (Q5). The size pattern reflects the fact that sub-penny price improvement is a less reliable marker of retail activity for large, liquid stocks where wholesaler quoting practices differ from those in smaller stocks.

Panel C shows that cross-sectional correlations are highest among past losers. For stocks in the bottom momentum quintile (Q1), the BJZZ correlation is 0.509, compared to 0.351 for past winners (Q5). This pattern has a direct implication for this paper's main analysis. The dip-buying behavior we document is concentrated precisely among past-loser stocks. The higher proxy reliability in that subsample means that evidence in the prior literature based on BJZZ is most credible when

studying retail contrarian activity, even if overall alignment has deteriorated in recent periods.

Panel D shows that time-series (cross-sectional) correlations between Retail Flow and BJZZ declined from 0.597 (0.560) in the pre-COVID period from 2017–2019, to 0.497 (0.466) during the COVID and meme-stock period from 2020–2021, and to -0.075 (0.403) in the post-meme period from 2022–2024. BHJOS shows a similar decline in the cross-section. This deterioration is consistent with structural changes in retail broker routing and wholesaler market share following the 2021 meme-stock episode. Researchers applying TAQ-based proxies to post-2021 data face substantially greater measurement noise than calibrations based on earlier samples would suggest.

The correlations also reveal how the relation between retail and institutional flows varies with past returns. The correlation between retail flow and changes in institutional ownership (ΔIO) is negative in the aggregate time series and near zero across stocks. In Panel C, the correlation between Retail Flow and ΔIO is more negative for past losers and near zero for past winners. The asymmetry is suggestive: the negative co-movement between retail and institutional flows is concentrated among past losers, where institutional investors are most likely to demand for immediacy.

3 Buying the Dip

How do retail investors trade? One body of evidence points to return chasing (Barber et al., 2009a). Another body of evidence points that retail investors are contrarians at (Kaniel et al., 2008). The two readings have coexisted for years.

In this section, we revisit the question with a comprehensive panel of U.S. retail order flow. We find that retail investors are "buying the dip": retail demand is contrarian, but the response is concentrated almost entirely in down states. Retail order flow rises substantially after negative contemporaneous and past returns, whereas positive returns generate little systematic net selling. The asymmetry is persistent across the sample period and holds across various horizons, while

non-retail flow measures on average move in the opposite direction.[4]

3.1 Retail Response to Returns

Figure 3 plots bin-averaged retail net flow (as a percent of shares outstanding) against past returns at four horizons. Panel A uses the contemporaneous monthly return $\text{Ret}[t]$. Panels B, C, and D use lagged returns over windows ending at $t-1$: the past month $\text{Ret}[t-2, t-1]$, the past six months $\text{Ret}[t-6, t-1]$, and the past twelve months $\text{Ret}[t-12, t-1]$. The construction follows a standard binscatter approach. Stocks are partitioned into 50 equal-frequency bins by past return within each panel, and the figure plots the mean retail flow within each bin.

The figure delivers two visual findings. First, the relation is strongly negative in the left tail and the slope steepens as returns become more negative. Retail inflows accelerate disproportionately as losses deepen. Second, retail net flow remains close to zero or slightly positive across the upper half of the return distribution. Stocks with positive past returns generate little systematic retail selling. The contrarian response is therefore concentrated almost entirely in down states. Because this measure is net flow, the downward tilt could in principle reflect a netting artifact rather than directional demand. Section 3.4 decomposes retail buying and selling and shows that it does not.

The pattern holds at various return horizons. Panel A shows that retail demand responds to contemporary price declines. Panels B through D show that the response also tracks lagged returns measured over windows that exclude the current month. The lagged-return panels are important for interpretation. They show that the relation is not driven by co-movement between contemporaneous flows and returns within the same month, given that contemporary returns have low correlations with their past.

[4]The "buying the dip" phenomenon has recently been widely observed by market participants and financial media. See, for example, the eToro Retail Investor Beat survey (eToro, 2023), which noted that despite a prolonged bear market, "67% were either ambivalent or positive about 2022 and its impact on their investing mindset...16% saw 2022 as an opportunity to buy the dip, whilst 15% say the bear market actually increased their appetite for investing." J.P. Morgan Asset Management (2025) similarly reports that retail investors continued to purchase U.S. equities despite the market drawdown of early 2025. The financial press frequently echoes these observations; see, e.g., Bloomberg News (2025). Relatedly, Welch (2022) finds that Robinhood investors did not panic-sell during the March 2020 market crash, and were rewarded in the subsequent bull market.

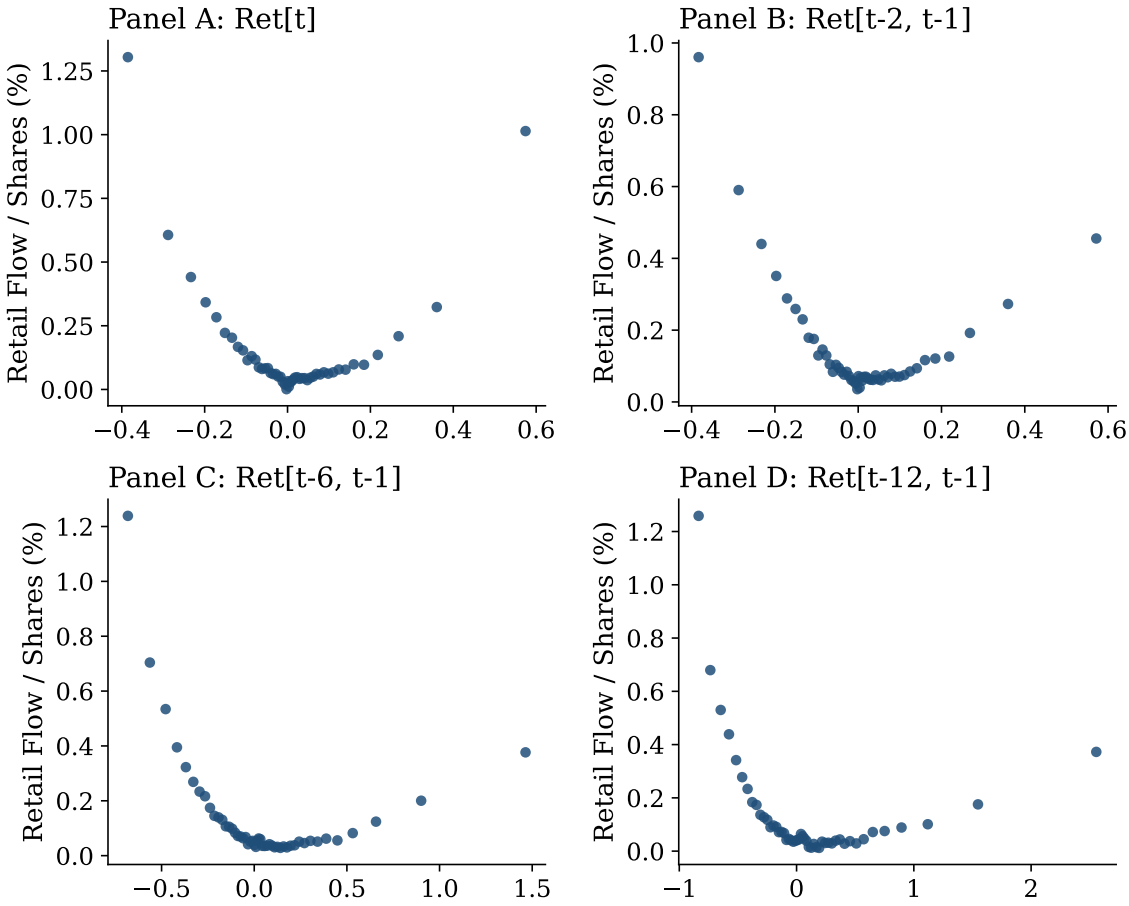


Figure 3. Buying the Dip: Retail Net Flow Versus Returns Across Horizons. Each panel is a binscatter with 50 equal-frequency bins. The vertical axis is monthly retail net order flow scaled by shares outstanding (percent). The horizontal axis is the past return measured over the indicated horizon. Panel A uses the contemporaneous return Ret$[t]$; Panels B, C, and D use lagged returns ending at $t-1$. The sample is January 2017 through December 2024. Both the past return and retail flow are winsorized at the 1% and 99% percentiles within each panel.

Retail investors continue to buy stocks that fell over the prior month, prior six months, or prior year, even after the most recent month's return is set aside.

To formalize the visual evidence and to assess robustness to firm characteristics and other flow measures, we estimate panel regressions of retail flow on a piecewise-linear function of past returns. For any return r, write $r^- \equiv \min(r, 0)$ and $r^+ \equiv$

$\max(r, 0)$ for its negative and positive parts. The specification is

$$\begin{aligned}\text{RetailFlow}_{i,t} = a_t + b_i + \beta_0^- \,\text{Ret}[t]^- + \beta_0^+ \,\text{Ret}[t]^+ \\ + \sum_{h\in\{2,3,6,12\}} \left[\beta_h^- \,\text{Ret}[t-h,\, t-1]^- + \beta_h^+ \,\text{Ret}[t-h,\, t-1]^+\right] \\ + \gamma' X_{i,t} + \varepsilon_{i,t},\end{aligned} \tag{1}$$

where $\text{Ret}[t]$ is stock i's contemporaneous monthly return at t, $\text{Ret}[t-h,\, t-1]$ is its cumulative return from month $t-h$ to month $t-1$, a_t and b_i are time and stock fixed effects, and $X_{i,t}$ is a vector of stock-month controls. Splitting each return into its negative and positive parts allows the slope of retail flow on returns to differ between losses and gains: a symmetric linear response would require $\beta_h^- = \beta_h^+$ for every h, while asymmetric dip-buying implies $\beta_h^- < 0$ and $\beta_h^+ \approx 0$.

The control vector $X_{i,t}$ is built up across the columns of Table A.1. Column (1) includes only the contemporaneous return components. Columns (2)–(4) progressively add three sets of controls: return-distribution moments (idiosyncratic volatility, idiosyncratic skewness, and market beta), firm characteristics (log market equity, book-to-market, debt-to-assets, return on assets, firm age, and gross profitability), and three contemporaneous institutional flow proxies discussed in Section 2.5 (the quarterly change in 13-F institutional ownership, the change in securities-lending supply, and the institutional large-trade proxy). Standard errors are two-way clustered by stock and by calendar month.

Figure 4 visualizes the regression coefficients across horizons.[5] Three features stand out. First, the loss-side coefficients β_h^- are negative and tightly estimated at every horizon. The single-horizon estimate at the contemporaneous return is approximately -1.3, indicating that a 1-percentage-point decline in $\text{Ret}[t]$ raises monthly retail flow by roughly 1.3 basis points of shares outstanding. Loss-side coefficients decline in absolute value as the horizon lengthens, with the twelve-month coefficient still significantly negative but roughly one-tenth the size of the contemporaneous one.

[5]The single-horizon estimates run one regression per horizon and include only that horizon's splits plus the composition controls in $X_{i,t}$. They isolate the loss-side and gain-side slopes at horizon h from the composition channel without conditioning on the other horizons' returns. The multivariate estimates correspond to column (4) of Table A.1 and include all five horizons jointly, so they report the marginal coefficient at horizon h holding all other horizons fixed.

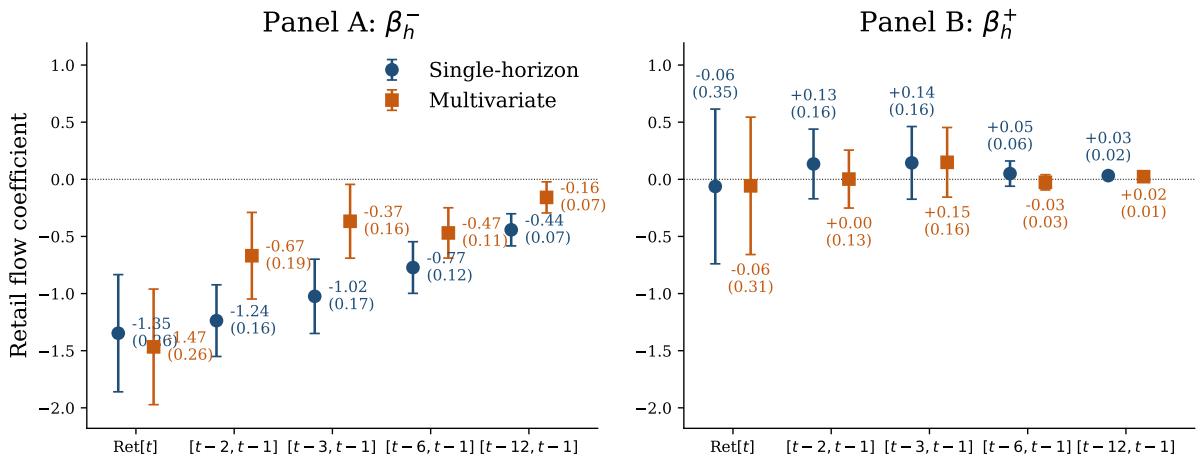


Figure 4. Regression Coefficients of Retail Flows on Returns Across Horizons. Panel A reports the loss-side coefficients β_h^- and Panel B the gain-side coefficients β_h^+ from estimates of equation (1), for horizons $h \in \{0, 1, 3, 6, 12\}$ months. Navy circles ("Single-horizon") come from a separate panel regression at each horizon that includes only the focal-horizon return splits, the composition controls in $X_{i,t}$, and stock and time fixed effects. Orange squares ("Multivariate") come from column (4) of Table A.1, which includes all five horizons jointly together with $X_{i,t}$ and the same fixed effects. Error bars are 95% confidence intervals based on standard errors clustered by stock and by calendar month. Each marker is annotated with the point estimate and standard error (in parentheses). The sample is January 2017 through December 2024.

Second, the gain-side coefficients β_h^+ are an order of magnitude smaller in absolute value, and their 95% confidence intervals span or hug zero at every horizon. The asymmetry between β_h^- and β_h^+ is therefore not a feature of the contemporaneous return alone; it holds at every past-return horizon up to twelve months. Third, the single-horizon and multivariate loss-side estimates track each other on the gain side but diverge on the loss side: the multivariate β_h^- is uniformly smaller in absolute value than its single-horizon counterpart. The two specifications agree that the gain-side response is essentially flat, and they together imply that the longer-horizon loss-side response operates partly through the shorter horizons in the multivariate fit.

In terms of economic magnitude, the cross-sectional standard deviation of retail net flow is 2.33 percentage points of shares outstanding per month. The standard deviation of Ret$[t]$ is 22.8 percentage points and that of Ret$[t-12, t-1]$ is 92.9 percentage points (Table 1). Using the column (4) estimates, a one-standard-deviation negative shock to Ret$[t]$ is associated with $1.47 \times 0.228 \approx 0.34$ percentage points of shares outstanding in additional retail buying, or roughly 14% of the cross-sectional retail-flow standard deviation. The gain-side coefficient at the contemporaneous

horizon is -0.06 with a standard error of 0.31, statistically indistinguishable from zero. At the past-twelve-month horizon, the marginal monthly effect is smaller, but it accumulates: the coefficient of -0.16 implies that a stock that fell 50 percent over the prior twelve months attracts roughly 0.08 percentage points of shares in additional retail buying every month in which that past-return state persists, or between three and four percent of the retail-flow standard deviation per month.

The positive-side upward tilt visible in Panel A of Figure 3 is therefore not evidence of retail return-chasing on the upside. The bivariate positive-return coefficient at the contemporaneous horizon (1.90 from Table A.1) falls to -0.06 once controls are added, consistent with the gain-side coefficient in Figure 4. The positive-return slope in the binscatter reflects the correlation between the contemporaneous return and other variables that drive retail demand, including idiosyncratic volatility, skewness, and market beta. Once these are controlled for, the positive-side slope disappears, turning marginally negative and statistically indistinguishable from zero. Section 3.2 documents a related feature in the time-series sort and discusses why a contemporaneous-return cut alone is misleading.

Appendix Table A.1 reports the full set of estimates including the control coefficients. The coefficient on idiosyncratic volatility is large and significantly positive across various specifications. Retail investors concentrate buying in stocks with greater idiosyncratic risk, consistent with prior evidence that retail demand is tilted toward speculative and high-volatility stocks (Kumar, 2009; Han and Kumar, 2013). The coefficient on idiosyncratic skewness is modestly *negative*, and the coefficient on market beta is also negative. Conditional on idiosyncratic volatility, retail flow is slightly higher for stocks with less positive return skewness and lower market beta. This nuances the lottery-stock framing in the literature: the volatility tilt is preserved, but the positive-skewness and high-beta tilt that would identify lottery-like demand is absent. Among the institutional flow proxies, the securities-lending flow enters positively, indicating that retail buying co-moves with lendable supply.

3.2 Buying the Dip over Time

The cross-sectional asymmetry documented in Section 3.1 could in principle be driven by a few episodes in which retail buying was concentrated. The COVID-19 lockdown period and the 2020–2021 meme-stock surge are two natural candidates. To assess whether buying the dip is structural or episodic, we examine the time series of average flows for stocks sorted into return terciles within each calendar month.

Figure 5 reports the results for two sorting variables. Panels A and B sort stocks into terciles by the contemporaneous monthly return $\text{Ret}[t]$ and plot the cross-sectional average of, respectively, retail flow and total initiated flow within each tercile. Panels C and D repeat the exercise sorting by the past twelve-month return $\text{Ret}[t-12, t-1]$. All flows are scaled by shares outstanding and expressed in percent. The shaded regions mark the COVID-19 selloff (February–May 2020) and the meme-stock window (January–March 2021).

Panel A reports the time-series of average retail flow by contemporaneous-return tercile. Both the bottom and the top tercile receive net retail buying. The middle tercile receives the least. Averaged across the 96 months in the sample, monthly retail flow is 0.37% of shares outstanding in the bottom $\text{Ret}[t]$ tercile, 0.05% in the middle tercile, and 0.26% in the top tercile. The bottom-minus-top gap is only 0.10% ($t = 4.09$), and the bottom tercile leads the top tercile in 74% of months. The pattern is best described as a U-shape with a slightly higher left tail.

A simple contrarian-toward-mean interpretation would predict retail flow to fall monotonically across contemporaneous return terciles. The pattern rejects that view. One plausible reason is that the contemporaneous monthly return is jointly determined with within-month order flow and is therefore correlated with several variables that drive retail demand independently of dip-buying. Stocks with highly positive returns in the same month also tend to have high recent volatility, recent news flow, and elevated retail attention. Some of those retail flows show up in the top tercile in Panel A, even though the underlying dip-buying behavior is fundamentally one-sided. The corresponding regression evidence is in Section 3.1: the bivariate positive-return coefficient falls substantially once additional controls are added. The

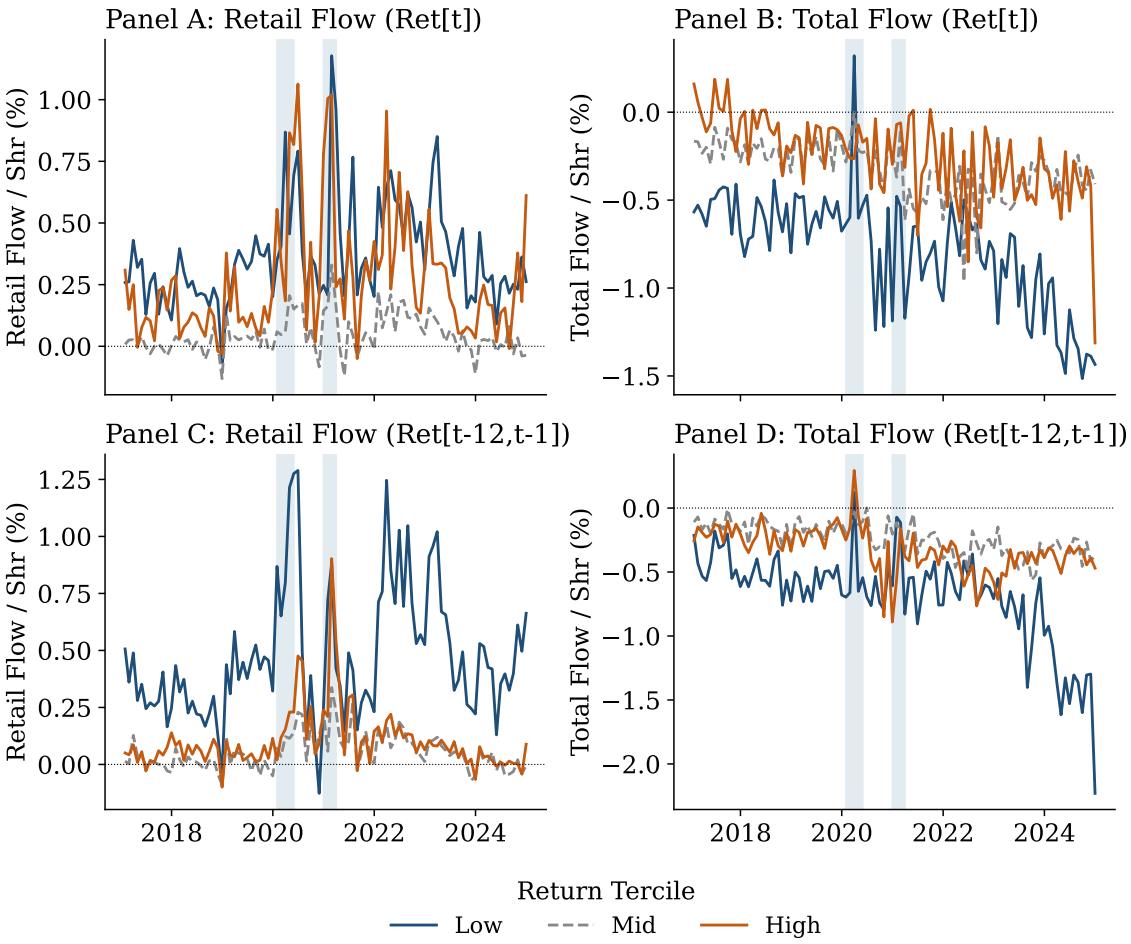


Figure 5. Retail and Total Flows by Return Tercile. Each panel plots the equal-weighted cross-sectional average of monthly net flow scaled by shares outstanding (percent) for stocks sorted into terciles within each calendar month. Panels A and B sort by contemporaneous return Ret$[t]$. Panels C and D sort by past twelve-month return Ret$[t-12, t-1]$. Panels A and C plot retail flow; Panels B and D plot total initiated flow. Shaded regions mark the COVID-19 selloff (February–May 2020) and the meme-stock window (January–March 2021). The sample is January 2017 through December 2024. Appendix Figure A.1 reports the market-cap weighted version.

contemporaneous-sort right tail of Panel A reflects this same composition issue. Therefore, it is necessary to also look at the past return sort.

Panel B shows that total initiated order flow is uniformly negative across all three Ret$[t]$ terciles, consistent with average net selling in the sample period. The most negative average flow is concentrated in the bottom tercile -0.78%. The bottom-minus-top gap in total flow is -0.53% ($t = -10.86$). Within a contemporaneous return sort, retail and total flows therefore have opposite signs in the bottom tercile. Retail investors buy net while non-retail investors sell net, even though both the bottom and

the top tercile of retail flow are positive on average.

Panels C and D sort by $\text{Ret}[t-12, t-1]$, which is predetermined relative to month t and is therefore not jointly determined with month-t flows. The contemporaneous-sort U-shape disappears; instead, we see a consistent L-shape pattern over the sample period. Average monthly retail flow is 0.47% of shares outstanding in the bottom past-return tercile, 0.05% in the middle tercile, and 0.10% in the top tercile. The bottom-minus-top gap is 0.37% ($t = 6.45$), nearly four times larger than under the contemporaneous-return sort. The bottom tercile leads the top tercile in 97% of the monthly sample, that is, in essentially every month. The few exception months cluster in the meme-stock window of early 2021, visible in Panel C of Figure 5. In those months a surge of return-chasing retail inflow into recent winners lifted top-tercile flow above bottom-tercile flow, briefly inverting the usual ordering before it resumed.

The first-order autocorrelation of the bottom-tercile retail flow series is 0.67. This high autocorrelation indicates that strong inflows into recent losers are not isolated to a few unusual months. Heavy retail buying in past-loser stocks persists across months and across distinct market regimes. Panel D shows the total flow is most negative in the bottom past-return tercile (-0.69% per month) and less negative in the top tercile (-0.33% per month). The bottom-minus-top gap is -0.36% ($t = -4.06$), and the bottom tercile lies below the top tercile in only 9.4% of months. The two panels together describe a stable feature of the panel: month after month, the same stocks that retail investors are buying most intensively are the stocks from which non-retail investors are withdrawing.

The shaded regions in Figure 5 mark two of the most disrupted episodes in the sample. The bottom-tercile retail flow series in Panel C lies above the middle and top terciles throughout the sample, including across the COVID-19 selloff and the meme-stock window. The bottom-tercile lead is visible in essentially every month and survives the COVID-19 selloff intact; the meme-stock window is the one brief interruption, after which the ordering resumes. The pattern is not driven by the meme-stock surge in retail participation, although that episode did raise the level of retail flow into all three terciles for several months. The opposing pattern in total flow

in Panel D is similarly persistent.

The market-cap weighted version of Figure 5, reported in Appendix Figure A.1, displays the same qualitative pattern. Buying the dip is therefore not concentrated in small-capitalization names; it is also visible when each tercile's average is weighted toward the largest stocks within the tercile. Together, Panels C and D and the persistence statistics establish that buying the dip is a structural feature of retail trading in this sample rather than a phenomenon tied to any single episode or cohort of new retail entrants.

3.3 Other Investor Flows

Sections 3.1 and 3.2 establish that retail flow rises after price declines and that the pattern persists through the full sample. A natural follow-up question is whether dip-buying is specific to retail investors or whether it reflects a broader response shared across initiated order flow. To address this, we apply the binscatter design of Figure 3 to non-retail flow measures using the lagged-return horizon. The lagged horizon mitigates the joint-determination concern raised in Section 3.2. We use the past twelve-month return $\text{Ret}[t-12, t-1]$ for consistency with the time-series sort in Section 3.2.

The four panels show that other investors react to returns differently: institutional proxies and total flow are momentum-following. Panel A plots the institutional large-trade proxy (Inst50k). The slope is positive across the return distribution: large trades are net buyers of past winners and net sellers of past losers. Panel B plots the securities lending supply (InstSec). The lendable supply rises in stocks with strong past returns and falls in past losers, indicating that institutional investors are chasing momentum. Panel C plots total initiated order flow, which exhibits a hump-shaped pattern. The pattern suggests that aggregate market flow tilts away from both past losers and past winners. Panel D plots the quarterly change in 13-F institutional ownership (ΔIO). The slope is positive, indicating that institutional investors sells losers and buy winners.

The retail flow slopes down on past returns; the four non-retail measures slope up. The opposing slopes imply that retail demand and the various institutional and aggregate measures of supply move in opposite directions across the past-return

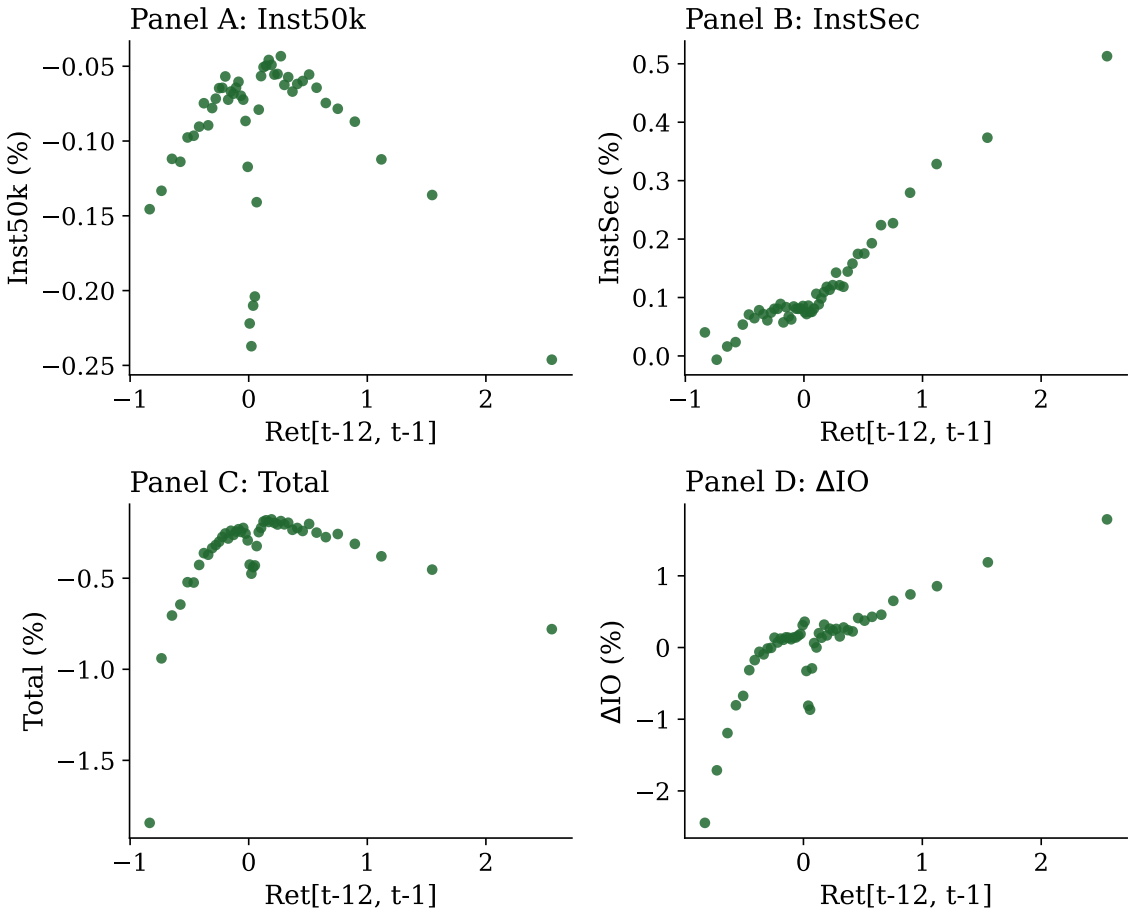


Figure 6. Other Net Flow Measures Versus Past Return. Each panel is a binscatter with 50 equal-frequency bins. The vertical axis is the indicated monthly net flow scaled by shares outstanding (percent). The horizontal axis is the past twelve-month return Ret$[t-12, t-1]$. Panel A: large-trade institutional proxy (Inst50k, trades above \$50,000 notional). Panel B: institutional securities-lending flow (InstSec). Panel C: total initiated order flow. Panel D: quarterly change in 13-F institutional ownership (ΔIO). Variable definitions follow Section 2.5. Both axes are winsorized at the 1% and 99% percentiles within each panel. The sample is January 2017 through December 2024.

distribution. Stocks that have fallen receive retail buying and institutional selling, while stocks that have risen receive institutional buying and only modest retail flow. The dip-buying response documented earlier in this section is therefore not a generic response of initiated order flow to past returns. It is a specific feature of retail flow.

3.4 Buy versus Sell

A long-standing concern with the retail trading pattern is that it may be a netting artifact (Barber and Odean, 2013). A downward-sloping relation between net flow and past returns need not mean that retail purchases respond contrarily. If gross purchases

and sales both move with past returns, net flow can slope down even though each leg moves up. We address this by decomposing retail net flow into its two gross legs. Our data identify retail buying and retail selling separately, and net flow is exactly their difference, RetailFlow $=$ RetailBuy $-$ RetailSell, so the decomposition is exact rather than estimated. Appendix Figure A.2 replicates the event-time exercise of Barber et al. (2009b, Figure 3). Weighting stock-months by retail buy versus sell intensity yields nearly the same pre-trade return path, so on average both trades follow stocks that experienced positive returns. The difference of the two legs tells a different story. The net flow path is hump-shaped: it peaks roughly a year before the trade and then declines steadily into the purchase month, ending below zero. Conditional on net buying, the stock has been losing over the year before retail investors step in. This is the event-time counterpart of the contrarian net behavior.

Separating the return into gain and loss sides, we find that retail purchases also reflect the buying-the-dip phenonmenon. Appendix Figure A.3 plots bin-averaged retail buying and selling against past returns at four horizons, with market turnover on a secondary axis. The gross response is not monotone: buying and selling are both elevated when the past return is large in absolute value, of either sign, and lowest near zero. Thus, a single average slope obscures the asymmetry. Moreover, the buy and sell legs are nearly coincident at every horizon, and market turnover traces the same shape, indicating a common turnover component in both legs. Appendix Figure A.4 estimates equation (1) separately for retail buying and selling, and market turnover, splitting each past return into its negative and positive parts. Gross retail purchases are contrarian to losses while chasing gains, and the buying, selling, and turnover coefficients are similar at each horizon. The common component is clearest in the gross-versus-net comparison: at the one-month lagged horizon the loss-side coefficient is -12.54 for buying and -11.30 for selling, with standard errors near 3.9, while netting the two legs leaves a loss-side coefficient of -1.24 with a standard error of 0.16. The roughly twenty-five-fold collapse in the standard error confirms that a large common component cancels in the net.

These results show that the dip-buying pattern among retail investors is not a

netting effect. Gross retail purchases are also contrarian particularly in the negative-return side. Furthermore, gross buying and selling activities are closely aligned, with a common movement related to market turnover. The contrarian tilt in net retail flow is the asymmetric component that remains once the common turnover variation is removed from the gross trades.

3.5 Robustness

The asymmetric dip-buying pattern is not an artifact of how retail flow is scaled or of the linear return measure. Appendix Figure A.5 reproduces the binscatter of Figure 3 using three alternative normalizations of net retail flow, at the contemporaneous and past-twelve-month horizons. The loss-side tilt is preserved under all three: dollar flow scaled by market capitalization traces the same steep left arm as the baseline measure, and the two signed order-imbalance ratios (net flow scaled by retail share volume, and dollar flow scaled by retail dollar volume) slope down monotonically in past returns, turning negative for past winners.

Appendix Figure A.6 re-estimates the coefficient plot of Figure 4 with log returns $\ln(1 + r_h)$ on the right-hand side; the loss/gain sign split is preserved because $\ln(1 + r_h) < 0$ if and only if $r_h < 0$. The asymmetry is unchanged: the loss-side coefficients β_h^- are negative and bounded away from zero at every horizon in both the single-horizon and multivariate specifications, while the gain-side coefficients β_h^+ are economically small with confidence intervals that span zero.

3.6 Relation to Prior Empirical Evidence

Contrarian retail trading. Existing work using TAQ-based proxies and proprietary brokerage data has documented that aggregated retail order flow leans against contemporaneous and short-horizon returns and predicts subsequent reversals (Kaniel et al., 2008; Kelley and Tetlock, 2013; Boehmer et al., 2021; Barrot et al., 2016; Luo et al., 2025). The loss-side response we document is consistent with this body of evidence. Two features push the pattern past what linear contrarian estimates can describe. The first is the asymmetric shape. The piecewise specification in equation (1) shows that the loss-side coefficient is roughly an order of magnitude

larger in absolute value than the gain-side coefficient at every horizon from one month to twelve months. The gain-side coefficient is statistically indistinguishable from zero at most horizons once standard controls are included. Linear contrarian specifications implicitly average across these two regimes and mask the kink. The second is the horizon profile. The loss-side coefficient remains negative and tightly estimated all the way out to twelve months, beyond the one-to-three-month windows that have served as the typical estimation horizon in prior work. The asymmetric, multi-horizon shape is the substantive addition that the retail flow data bring to the contrarian-retail literature.

Retail demand for distressed stocks. A separate asset-pricing literature documents that financially distressed stocks earn anomalously low returns and that retail investors are concentrated on the buy side. Distress in this work is a fundamental characteristic, measured from accounting and market data, not a return realization. Campbell et al. (2008) estimate firm-level failure probabilities and show that high-failure-probability stocks deliver low risk-adjusted returns inconsistent with rational compensation for distress risk. Conrad et al. (2014) attribute the anomaly to retail demand for jackpot-like upside in distressed firms, drawing on the stocks-as-lotteries framework of Barberis and Huang (2008). A closely related strand documents the same underlying preference for jackpot-like upside using individual-investor brokerage account data (Mitton and Vorkink, 2007; Kumar, 2009). Our stylized fact conditions retail demand on contemporaneous and past returns, which remain robust after controlling for fundamental distress or jackpot characteristics.

4 A Model of Buying the Dip

Section 3 documented two features of dip-buying that a model should reproduce. The first is the *asymmetry across loss and gain states*: retail flow rises markedly after past losses but is statistically indistinguishable from zero after past gains, so a mechanism that generates a symmetric response on either side of the return distribution cannot fit the data. The second is the *response to signed past returns at multiple horizons*: the

trigger is not the absolute return and it is not a stable stock characteristic; it is the signed return measured over windows from one month to twelve months.

This section develops a model that produces both features from two ingredients: sparse attention to a "default" historical price forecast in the spirit of Gabaix (2014), and a short-sale constraint on retail investors. We calibrate the model by simulated method of moments and use the calibrated benchmark for three exercises: the cross-section of retail flow against past returns, the dynamic response to supply shocks, and the dependence of return volatility on the retail share.

4.1 Environment

The economy has discrete time and an infinite horizon. A single risky asset trades each period. Its next-period cum-dividend payoff is $y_{t+1} = p_{t+1} + f_{t+1}$, where p_{t+1} is the ex-dividend price and f_{t+1} is the period-$t+1$ cash flow.

Cash flow follows

$$f_t = \bar{d} + \theta_t + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2),$$

where θ_t captures persistent variation in cash-flow fundamentals and ϵ_t is an idiosyncratic cash-flow innovation observed at date t. Net asset supply is

$$Q_t = \bar{Q} + \lambda_t,$$

where λ_t is a transitory supply-pressure component that shifts the quantity of the asset that investors must absorb. We collect the two latent components in the state vector $x_t = (\theta_t, \lambda_t)'$. The state evolves according to the first-order vector autoregression:

$$x_t = F x_{t-1} + u_t, \qquad F = \mathrm{diag}(\rho_\theta, \rho_\lambda), \qquad u_t \sim N(0, \Sigma_u),$$

with $\Sigma_u = \mathrm{diag}(\sigma_{u,\theta}^2, \sigma_{u,\lambda}^2)$. Equivalently,

$$\theta_t = \rho_\theta \theta_{t-1} + u_{\theta,t}, \qquad \lambda_t = \rho_\lambda \lambda_{t-1} + u_{\lambda,t}.$$

The asset is held by two classes of overlapping-generations investors. Each generation lives one period, is myopic, and has constant-absolute-risk-aversion (CARA) preferences over next-period wealth. *Institutional investors* observe the current state x_t and the public history. With CARA preferences and Gaussian payoffs their

aggregate demand is the standard mean-variance schedule,

$$Q_t^I = \tau^I(\mu_t^I - p_t), \qquad \mu_t^I = E_t^I[y_{t+1}],$$

where μ_t^I is the institutional forecast of the cum-dividend payoff and $\tau^I > 0$ is institutional risk-bearing capacity. *Retail investors* form a continuum. Retail investor i has idiosyncratic non-financial income whose covariance with the asset payoff generates a hedging need:

$$h_i = \bar{h} + \eta_i, \qquad \eta_i \sim N(0, \sigma_h^2).$$

A higher h_i means the investor's non-financial income covaries more positively with the payoff, so the desired stock position is smaller. Retail investors face a short-sale constraint, while instituional investors do not.

We index the retail share of the market by $\alpha \in (0,1)$. Writing γ_I and γ_R for the two CARA coefficients and σ_I^2 and σ_R^2 for the payoff variances each class prices risk against, the effective risk-bearing capacities are

$$\tau^I = \frac{1-\alpha}{\gamma_I \, \sigma_I^2}, \qquad \tau^R = \frac{\alpha}{\gamma_R \, \sigma_R^2}. \tag{2}$$

Raising α reallocates risk-bearing capacity from institutions to retail. This parameterization lets the volatility exercise hold the deep preference and payoff-risk parameters fixed while varying the composition of who bears risk.

4.2 Sparse attention and retail demand

Institutional investor observes x_t, while retail does not. Retail sees only the public history through the end of period $t-1$ and enters period t with the posterior mean $\widehat{x}_{t-1}$. Before trading, retail forms a forecast of the current price, $\widehat{p}_t = E_{t-1}^R[p_t]$, and processes the realized price only partially. The price retail acts on is a convex combination of the current price and the historical price forecast:

$$\widetilde{p}_t = m\, p_t + (1-m)\, \widehat{p}_t, \qquad m \in [0,1]. \tag{3}$$

The function form (3) adopts the sparsity-based inattention from Gabaix (2014). The parameter m is a primitive that measures attention to the contemporaneous price: at $m = 1$ retail fully processes the current price, and at $m = 0$ retail trades entirely on its lagged forecast.

Retail's unconstrained demand index is $\tau^R(\mu_t^R - \widetilde{p}_t) - h_i$, where $\mu_t^R = E_{t-1}^R[y_{t+1}]$ is retail's forecast of the cum-dividend payoff. Imposing the short-sale constraint, individual retail demand is

$$q_{i,t} = \max\{\tau^R(\mu_t^R - \widetilde{p}_t) - h_i,\ 0\}. \tag{4}$$

Let $w_t \equiv \tau^R(\mu_t^R - \widetilde{p}_t) - \bar{h}$ denote the hedging-adjusted excess-demand index, the common component of (4) net of average hedging demand. Integrating (4) over the Gaussian cross-section of η_i gives aggregate retail demand in closed form as a truncated-normal map,

$$Q_t^R = G(w_t) = w_t\,\Phi\left(\tfrac{w_t}{\sigma_h}\right) + \sigma_h\,\phi\left(\tfrac{w_t}{\sigma_h}\right),$$

where Φ and ϕ are the standard normal CDF and density. The map G is the source of the asymmetry. When w_t is sufficiently negative, almost the entire retail cross-section is at the constraint and aggregate demand is close to zero and nearly flat in w_t; when w_t is positive, a large mass is unconstrained and G is approximately linear with slope near one. Because each retail generation is newly born, its aggregate demand is also retail order flow, so Q_t^R is the model counterpart of the empirical flow measure.

The within-period sequence is as follows. The state x_t is realized; retail enters with $\widehat{x}_{t-1}$ and forms μ_t^R and $\widehat{p}_t$; institutions and retail submit demand schedules and the market clears at p_t; the cash flow f_t is realized; and retail updates its posterior from $\widehat{x}_{t-1}$ to $\widehat{x}_t$ using the period-t price and the cash flow.

4.3 Equilibrium

A stationary equilibrium is a price rule mapping $(x_t, \widehat{x}_{t-1})$ to p_t and a retail belief-updating rule mapping $(\widehat{x}_{t-1}, p_t, f_t)$ to $\widehat{x}_t$ such that the market clears in every period,

$$\tau^I\left(\mu_t^I - p_t\right) + G(w_t) = \bar{Q} + \lambda_t, \tag{5}$$

the forecasts $\mu_t^I = E[y_{t+1} \mid x_t, \widehat{x}_{t-1}]$ and $\mu_t^R = E[y_{t+1} \mid \widehat{x}_{t-1}]$ are taken under the law of motion that the price and belief rules induce, and $\widehat{x}_t$ is the Bayesian filter implied by that same law of motion.

The equilibrium is nonlinear because the clearing price enters retail demand through the truncated-normal map G. We keep G exact and linearize the pricing

and filtering block around the stationary mean, at which the price is $\bar{p}$, retail's mean payoff forecast is $\bar{\mu}^R$, and the demand index takes the value $w_0 = \tau^R(\bar{\mu}^R - \bar{p}) - \bar{h}$. To first order the price is affine in the state and in retail beliefs,

$$p_t = \bar{p} + \Gamma x_t + \Pi \widehat{x}_{t-1} + R_t^p, \tag{6}$$

with R_t^p second order; all equalities below hold up to second-order remainder terms.

The coefficients (Γ, Π) are the joint fixed point of pricing and filtering. They fix the signal retail extracts from the price, and hence the Kalman gain that governs how retail beliefs evolve, and institutions price the asset knowing that gain. Appendix B derives retail's filtering problem, the institutional forecast, and the fixed-point system.

4.4 Results

We estimate the parameters in two steps.[6] The first step does not rely on solving the model. In this step, we pin down the persistence of the fundamental and of supply pressure, the standard deviations of returns and of retail order imbalance (which the measurement scales match by construction), the retail risk-aversion coefficient, taken from outside estimates (Beutel and Weber, 2026; Gorodnichenko and Yin, 2026), the signal noise σ_ϵ, which inverts the autocorrelation of the public cash-flow signal in closed form, and the average hedging need $\bar{h}$, which inverts the mean of retail flow relative to its dispersion, a scale-free moment that is strictly monotone in the location of the short-sale constraint. In the second step, we estimate the institutional trading sensitivity τ^I, the inattention parameter m, and the standard deviation of fundamental innovations $\sigma_{u,\theta}$ by simulated method of moments (SMM), targeting the loss- and gain-side slopes of regressing retail order imbalance on past return and the R^2 from regressing current return on current dividend growth. All analysis are based on a simulated path of 10,000 periods after a 2,000-period burn-in.

Table 4 reports the fit, and Figure 7 plots model retail order flow against the lagged return. The calibrated baseline ($\bar{h} = 8.11$, $m = 0.036$) reproduces the empirical dip-buying pattern of Section 3: loss- and gain-side slopes are -0.031 and 0.004 in the model against -0.031 and 0.008 in the data, and the return–dividend-growth R^2

[6] Appendix C reports the calibration and simulation details.

Table 4. Calibration Results

The table reports the calibration of the structural model. Panel A reports moments matched exactly: the scales $\bar{P}$ and a_Q set the simulated standard deviations of returns and retail flow to their empirical values; σ_ϵ is recovered in closed form from the cash-flow autocorrelation; and $\bar{h}$ inverts the flow-level ratio $E[Q_t^R]/sd(Q_t^R)$, targeted at the firm-level average of within-firm mean flow over its standard deviation. Panel B reports the three moments targeted by exactly identified SMM; columns (4)–(6) give the parameter most informative about each moment, its estimate, and its standard error. Fixed parameters include $\rho_\theta = 0.96$, $\rho_\lambda = 0.5$, $\sigma_{u,\lambda} = 1$, $\sigma_h = 10^{-6}$, $\alpha = 0.25$, and $\tau^R = 1.221$.

	(1)	(2)	(3)	(4)	(5)
	Moments		Parameters		
	Data	Model		Estimates	SE
Panel A. Moments matched by construction					
$sd(r_t^\ell)$	0.183	0.183			
$sd(Q_t^R)$	0.014	0.014	τ^R	1.221	
$\text{corr}(d_t, d_{t-1})$	0.820	0.820	σ_ϵ	0.278	
$E[Q_t^R]/sd(Q_t^R)$	0.150	0.150	$\bar{h}$	8.111	
Panel B. SMM moments					
β^- (loss-side slope)	-0.031	-0.031	τ^I	0.116	(0.007)
β^+ (gain-side slope)	0.008	0.004	m	0.036	(0.016)
$R^2(r_t^\ell \sim d_t)$	0.005	0.005	$\sigma_{u,\theta}$	0.188	(0.018)

is 0.005 for both empirical and model-simulated data.

The estimated attention parameter, $\widehat{m} \approx 0.04$, indicates that retail investors in period t trade mostly on their forecast of the current price formed from last period's information. This is a direct consequence of matching retail's strong contrarian response to past prices: when m is close to zero, the perceived cost of the asset is assessed almost entirely from lagged information and therefore loads strongly on past prices (Figure A.7). Bayesian learning in turn implies that the perceived price has more inertia than the traded price: in Figure A.8, $\tilde{p}_t$ tracks p_t on average (Panel A) but varies considerably less (Panels B).

A second observation is that the estimated institutional trading sensitivity, $\widehat{\tau}^I = 0.116$, is an order of magnitude below the calibrated retail sensitivity $\tau^R = 1.221$. The gap reflects the two information sets. Institutional investors observe the fundamental and supply realizations as well as the traded price, so their return expectations are far more volatile than retail's expectations, which are smoothed by Bayesian learning from lagged signals (Figure A.8). Matching the empirical standard deviations of returns and of retail order imbalance, together with the flow–return slopes, then

requires a small trading sensitivity for the investors with volatile expectations, and is consistent with a large sensitivity for the investors with smooth perceived returns.

Panel B of Figure 7 reports comparative statics. The hedging need $\bar{h}$ shifts the schedule in parallel, since it enters mean demand additively. Larger h shifts retail demand down. Since retail flow is mostly binding on the gain region, fewer retail demand in general only weakens the loss-side relationship between retail flow and past return.

The attention parameter m reshapes the curvature of the dip-buying pattern. Raising m away from zero introduces two forces. First, the retail demand index now loads more on the contemporaneous state, which makes it more volatile conditional on the past return. Because realized flow is truncated at zero, and the mean of a truncated variable rises with its dispersion, the more volatile demand raises average flow throughout the schedule. Second, when $m > 0$, retail also trades against current price. Stronger dip-buying after a past loss pushes the current price up, and retail observes a fraction m of that increase and perceives a smaller expected return. Thus, marginal demand is partially offset by an expected increase in current price, and the schedule flattens at deeper losses.

4.4.1 Price dynamics

Figure 8 reports the impulse responses to one-time supply-pressure shocks, where σ_λ is the stationary standard deviation of the supply state. Panels A and B report the price and retail-flow responses to a $\pm 2\sigma_\lambda$ shock, and Panels C and D report the price response to $\pm 1\sigma_\lambda$ and $\pm 3\sigma_\lambda$ shocks. For each shock we simulate three paths that share the same realizations of cash-flow shocks, supply shocks, and signal noise: a no-shock benchmark, a path with a positive innovation to the supply state at the event date, and a path with a negative innovation. The response at each horizon is the difference between a shocked path and the benchmark, averaged over 500 paths.

Retail flow is zero on impact and peaks one period after the shock (Panel B). In the baseline model, retail trades primarily on lagged beliefs, so it does not act on the contemporaneous price and moves only the following period, after it updates its forecast from the realized price and cash flow. At the event date, therefore, only

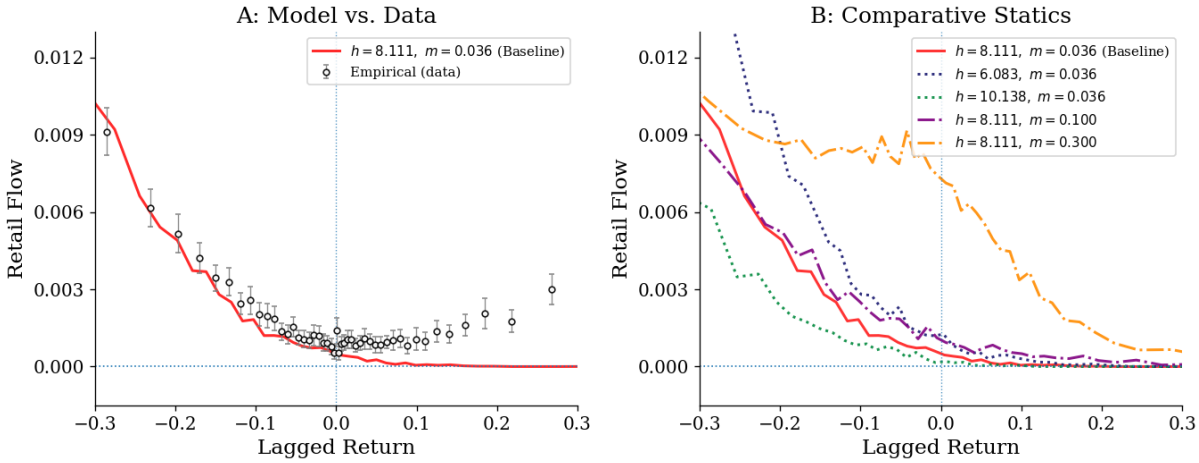


Figure 7. Retail Order Flow and the Lagged Return. Both panels plot model retail order flow against the lagged cum-dividend return, using 40 binned averages from a simulation of 80,000 periods after a 3,000-period burn-in, with retail flow scaled to its empirical standard deviation. Panel A overlays the calibrated benchmark on the empirical buying-the-dip schedule (open circles, with 95% confidence intervals); the empirical points are binned averages of retail net flow scaled by shares outstanding against the lagged return. Panel B reports comparative statics: holding all other parameters at their calibrated values, it varies the average hedging need $\bar{h}$ ($\pm 25\%$ around the benchmark) and the attention parameter m ($m = 0.10$ and $m = 0.30$).

institutions respond to the innovation, and they price the supply shock linearly. The impact is correspondingly symmetric: the price moves by about 30% in either direction (Panel A). The asymmetry appears only in the recovery. After the positive shock retail flow rises by about 4.9%, while after the negative shock it falls by only about 0.15%. Retail buys the dip strongly after a price decline but barely adjusts after an equal-sized increase, because the short-sale constraint binds on the sell side, consistent with the loss-gain asymmetry in Section 3.1.

The delayed and asymmetric retail response gives the two shocks different recovery shapes (Panel A). After the positive supply shock, the price recovers most of the way within one period. The retail dip-buying at period 1, together with the mean-reversion of the supply state, returns the price close to its pre-shock level. However, because the retail buying is concentrated in that single period, the price dips again at period 2 once the buying normalizes, and then converges through a damped oscillation. After the negative supply shock, retail is constrained from selling and adds almost no offsetting flow, so the price decays monotonically at roughly the supply state's own rate of mean-reversion ($\rho_\lambda = 0.5$). The same

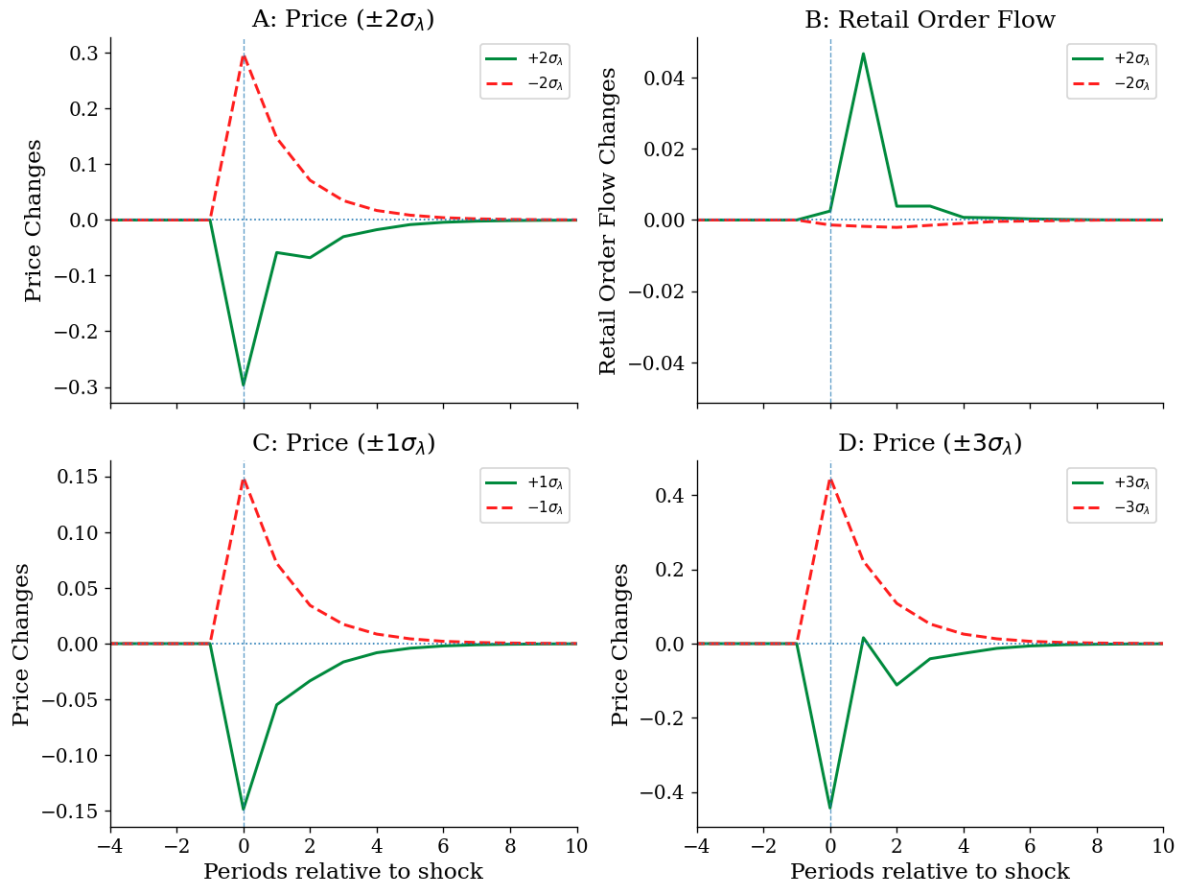


Figure 8. Price and Retail-Flow Responses to Supply Shocks. Generalized impulse responses to one-time supply-pressure shocks of size $\pm k\sigma_\lambda$, where σ_λ is the stationary standard deviation of the supply state. For each shock we simulate three paths that share the same realizations of cash-flow, supply, and signal shocks: a no-shock benchmark, a path with a $+k\sigma_\lambda$ innovation at the event date, and a path with a $-k\sigma_\lambda$ innovation. The response at each horizon is the difference between a shocked path and the benchmark, averaged over 500 paths. Panels A and B plot the price and retail-flow responses to the $\pm 2\sigma_\lambda$ shock; Panels C and D plot the price response to the $\pm 1\sigma_\lambda$ and $\pm 3\sigma_\lambda$ shocks. Solid (green) lines are positive shocks and dashed (red) lines are negative shocks; each panel is scaled to its own response magnitude.

asymmetric, one-period-delayed retail demand therefore maps a symmetric impact into an asymmetric recovery, fast but oscillatory after a price decline and slow but monotonic after a price increase.

Panels C and D reveal that the price response to supply shocks is nonlinear in the shock size. At a one-standard-deviation shock the positive and negative price paths are close to symmetric. This symmetry breaks as the shock grows. While on impact the price response stays roughly linear in the shock, the recovery after a positive supply shock becomes increasingly oscillatory and eventually overshoots its

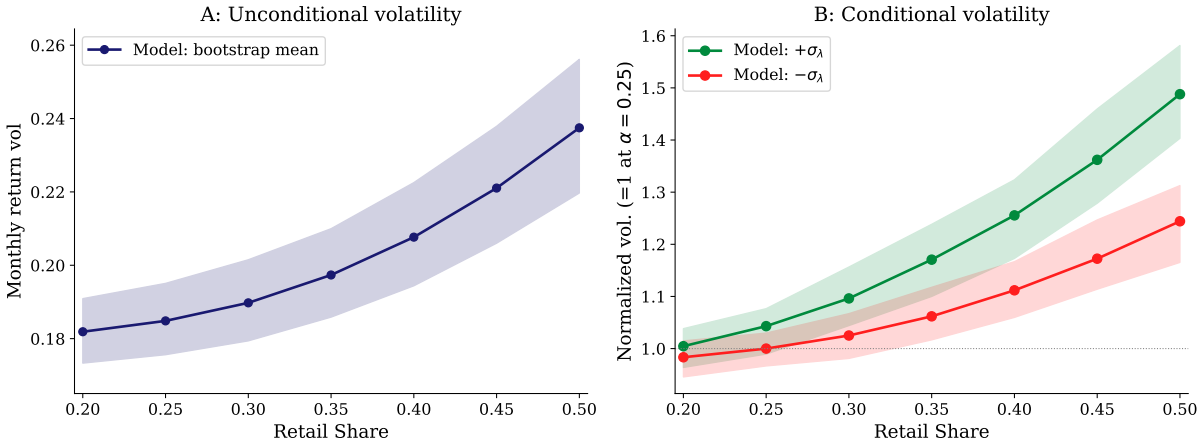


Figure 9. Retail Share and Return Volatility. Each panel plots return volatility against the retail share. In Panel A the model line is the unconditional return standard deviation, shown as the bootstrap mean with a 95% confidence band (1,500 simulated periods after a 500-period burn-in). In Panel B the model lines are the return volatility following a $+\sigma_\lambda$ (high) or $-\sigma_\lambda$ (low) supply-pressure shock. Both series in Panel B are normalized to one at the low-shock value at $\alpha = 0.25$.

pre-shock level when the shock is sufficiently large. The recovery after a negative shock remains a smooth monotonic decay, so the asymmetry between the two paths widens with the shock size. The nonlinearity is driven by the delayed and constrained dip-buying from the retail investor.

4.4.2 Return volatility

How does increased retail trading activity affect return volatility, and does the asymmetry of retail flow contribute to that effect? The following exercise examines unconditional volatility and conditional volatility under supply shocks by varying the retail share α while holding preferences and payoff-risk parameters fixed.

Panel A of Figure 9 plots unconditional return volatility against the retail share. The model predicts that the monthly return standard deviation rises from about 19% at $\alpha = 0.25$ to about 26% at $\alpha = 0.50$, and the relation is convex. Two forces raise volatility with α. First, because retail investors trade on lagged price information, institutions must absorb the contemporaneous shock; as α increases, their effective risk-bearing capacity $\tau^I = (1-\alpha)/(\gamma_I \sigma_I^2)$ shrinks, so the same shock moves the price more on impact. Second, and less intuitively, retail's contrarian demand destabilizes prices rather than stabilizing them. A contrarian trading contemporaneously would lean against the shock, and dampen the price movement on impact. But retail

investors are contrarian but *delayed*, taking action after the shock has begun to mean-revert. Their contrarian trading does not offset the original move and instead adds a second, opposite one (Section 4.4.1), overcorrecting the price and raising the variance. A larger retail share makes that correction larger. The delayed response thus reverses the usual sign: the contrarian demand that would have calmed the market on impact instead amplifies volatility over time.

Panel B conditions on the direction of the supply shock, normalizing each series to one at $\alpha = 0.25$. The model lines are the return volatility following a $+\sigma_\lambda$ (high) or $-\sigma_\lambda$ (low) shock. Both rise with the retail share, but the high-shock series lies above the low-shock series at every α, and the gap widens from around 13% at $\alpha = 0.25$ to 30% at $\alpha = 0.50$. The asymmetry in volatility comes from the asymmetric response of positive versus negative shocks. The delayed contrarian response that amplifies volatility is larger after a price decline, where retail buys the dip without constraint, than after a price increase, where retail can only pull back modestly. A positive supply shock therefore raises volatility more than a negative shock of the same size, and a larger retail share widens the gap.

4.5 Alternative explanations

The model delivers the two features of dip-buying, the asymmetry across loss and gain states and the response to signed past returns at multiple horizons, from a single set of primitives. We now consider whether alternative behavioral mechanisms in the prior literature reproduce both. Some of the mechanisms fail because they generate symmetric response in the sign of the return. The attention or salience framework of Barber and Odean (2008), Barber et al. (2022), van der Beck et al. (2026), and Graves (2024) predicts retail buying of attention-grabbing stocks identified by extreme returns, news, or volume, so it generates a response to $|r_{i,t}|$ rather than to $r_{i,t}$, and the attention measures decay too quickly to match the longer horizon. Mean-reverting beliefs, in which retail expects prices to revert toward fundamentals (Cohen et al., 2002; Bastianello and Fontanier, 2024, 2025), matches retail response at the loss side, but predicts retail selling after gains that is absent in the data. Extrapolative beliefs (Barberis et al., 2015) can generate retail responses to past returns but fail to create

contrarian trading. Our model delivers the asymmetry from a complementary route, a short-sale constraint layered on a mean-reversion belief that is itself symmetric in the supply-pressure state.

The other candidates fail because they are not return-conditional or predict the wrong sign. Habitat and lottery preferences describe stable cross-sectional tilts toward hard-to-value, high-volatility, or high-skewness stocks (Laarits and Sammon, 2025; Balasubramaniam et al., 2023; Dorn and Huberman, 2010; Kumar, 2009; Han and Kumar, 2013; Bali et al., 2011), not flow conditioned on the signed past return. Appendix Table A.1 further shows retail flow loading negatively on idiosyncratic skewness, the opposite of the lottery prediction. The disposition effect (Shefrin and Statman, 1985; Odean, 1998) predicts that retail selling slackens after losses, yet the buy-sell decomposition in Section 3.4 shows the loss-side gross sell response rising rather than falling, close in magnitude to the loss-side buy response (Appendix Figures A.3 and A.4). Each mechanism may contribute to the level of dip-buying in particular states or stocks, but none reproduces the asymmetric, multi-horizon shape that the single-mechanism model delivers.

5 Conclusion

This paper documents a stylized fact about retail trading: retail investors buy the dip. Retail net order flow rises markedly after negative contemporaneous and past returns and is roughly flat after positive returns. The sign-contingent asymmetry holds across horizons and is stable over time. We rationalize the asymmetry with a structural model of sparse attention and a binding short-sale constraint, calibrated by simulated method of moments. The model reproduces the loss-side slope and predicts that return volatility rises with the retail share, which we confirm in the cross-section of stocks.

Several questions remain for future research. First, identifying the supply-pressure channel requires pinning down which institutional friction generates the exogenous pressure that the model takes as a primitive. Second, our data measure aggregate retail order flow rather than individual accounts. They speak to retail trading behavior in aggregate but cannot trace holding periods or realized returns.

Whether retail investors also exit in time to capture the reversal premium therefore requires more granular data that we leave to future work. Third, a fuller welfare characterization will need to combine flow data with investor-level market timing. We view the asymmetric, state-contingent shape of retail demand as the central object that such work must confront.

References

Balasubramaniam, Vimal, John Y. Campbell, Tarun Ramadorai, and Benjamin Ranish. Who owns what? a factor model for direct stock holding. *The Journal of Finance*, 78 (3):1545–1591, 2023. doi: 10.1111/jofi.13220.

Bali, Turan G., Nusret Cakici, and Robert F. Whitelaw. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2):427–446, 2011. doi: 10.1016/j.jfineco.2010.08.014.

Barardehi, Yashar H., Zhi Da, Peter Dixon, and Junbo L. Wang. You can only lend what you own: Inferring daily institutional trading from lendable equity inventory. Working paper, SSRN, November 2024. Date written: November 15, 2024; posted November 25, 2024; last revised October 30, 2025.

Barber, Brad M. and Terrance Odean. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2):785–818, 2008. doi: 10.1093/rfs/hhm079.

Barber, Brad M. and Terrance Odean. The behavior of individual investors. In George M. Constantinides, Milton Harris, and René M. Stulz, editors, *Handbook of the Economics of Finance*, volume 2, chapter 22, pages 1533–1570. Elsevier, 2013. doi: 10.1016/B978-0-44-459406-8.00022-6.

Barber, Brad M., Terrance Odean, and Ning Zhu. Do retail trades move markets? *The Review of Financial Studies*, 22(1):151–186, 2009a. doi: 10.1093/rfs/hhn035.

Barber, Brad M., Terrance Odean, and Ning Zhu. Systematic noise. *Journal of Financial Markets*, 12(4):547–569, 2009b. doi: 10.1016/j.finmar.2009.03.003.

Barber, Brad M., Xing Huang, Terrance Odean, and Christopher Schwarz. Attention-induced trading and returns: Evidence from robinhood users. *The Journal of Finance*, 77(6):3141–3190, 2022. doi: 10.1111/jofi.13183.

Barber, Brad M., Xing Huang, Philippe Jorion, Terrance Odean, and Christopher Schwarz. A (sub)penny for your thoughts: Can analyzing order submissions improve our understanding of retail investor behavior? *The Journal of Finance*, 79(4): 2311–2356, 2024. doi: 10.1111/jofi.13334.

Barberis, Nicholas and Ming Huang. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5): 2066–2100, 2008. doi: 10.1257/aer.98.5.2066.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. X-capm:

An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1): 1–24, 2015. doi: 10.1016/j.jfineco.2014.08.007.

Barrot, Jean-Noel, Ron Kaniel, and David Sraer. Are retail traders compensated for providing liquidity? *Journal of Financial Economics*, 120(1):146–168, 2016. doi: 10.1016/j.jfineco.2016.01.005.

Bastianello, Francesca and Paul Fontanier. Partial equilibrium thinking, extrapolation, and bubbles. Working paper, SSRN, 2024.

Bastianello, Francesca and Paul Fontanier. Expectations and learning from prices. *The Review of Economic Studies*, 92:1341–1374, 2025. doi: 10.1093/restud/rdae059.

Battalio, Robert H, Robert H Jennings, Mehmet Saglam, and Jun Wu. Difficulties in obtaining a representative sample of retail trades from public data sources. *Available at SSRN 4579159*, 2023.

Beutel, Johannes and Michael Weber. Beliefs and portfolios: Causal evidence. *The Review of Financial Studies*, page hhag011, 2026.

Bloomberg News. YOLO crowd's record dip-buying binge calms a jumpy stock market. Bloomberg News, May 2025. URL https://www.bloomberg.com/news/articles/2025-05-19/yolo-crowd-s-record-dip-buying-binge-calms-a-jumpy-stock-market.

Boehmer, Ekkehart, Charles M. Jones, Xiaoyan Zhang, and Xinran Zhang. Tracking retail investor activity. *The Journal of Finance*, 76(5):2249–2300, 2021. doi: 10.1111/jofi.13051.

Campbell, John Y., Jens Hilscher, and Jan Szilagyi. In search of distress risk. *The Journal of Finance*, 63(6):2899–2939, 2008. doi: 10.1111/j.1540-6261.2008.01416.x.

Campbell, John Y., Tarun Ramadorai, and Allie Schwartz. Caught on tape: Institutional trading, stock returns, and earnings announcements. *Journal of Financial Economics*, 92(1):66–91, 2009. doi: 10.1016/j.jfineco.2008.03.006.

Cohen, Randolph B., Paul A. Gompers, and Tuomo Vuolteenaho. Who underreacts to cash-flow news? evidence from trading between individuals and institutions. *Journal of Financial Economics*, 66(2–3):409–462, 2002. doi: 10.1016/S0304-405X(02)00229-5.

Conrad, Jennifer, Nishad Kapadia, and Yuhang Xing. Death and jackpot: Why do individual investors hold overpriced stocks? *Journal of Financial Economics*, 113(3): 455–475, 2014. doi: 10.1016/j.jfineco.2014.04.001.

Coval, Joshua and Erik Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512, 2007. doi: 10.1016/j.jfineco.2006.09.007.

Cready, William, Abdullah Kumas, and Musa Subasi. Are trade size-based inferences about traders reliable? evidence from institutional earnings-related trading. *Journal of Accounting Research*, 52(4):877–909, 2014. doi: 10.1111/1475-679X.12056.

Daniel, Kent, Lorenzo Garlappi, and Kairong Xiao. Monetary policy and reaching for income. *The Journal of Finance*, 76(3):1145–1193, 2021. doi: 10.1111/jofi.13004.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Positive feedback investment strategies and destabilizing rational speculation. *The Journal of Finance*, 45(2):379–395, 1990. doi: 10.1111/j.1540-6261.

1990.tb03695.x.

Dorn, Daniel and Gur Huberman. Preferred risk habitat of individual investors. *Journal of Financial Economics*, 97(1):155–173, 2010. doi: 10.1016/j.jfineco.2010.03.013.

Eaton, Gregory W., T. Clifton Green, Brian S. Roseman, and Yanbin Wu. Retail trader sophistication and stock market quality: Evidence from brokerage outages. *Journal of Financial Economics*, 146(2):502–528, 2022. doi: 10.1016/j.jfineco.2022.08.002.

eToro. Retail investor beat. eToro Retail Investor Beat Survey, January 2023. Quarterly retail investor sentiment survey.

Friedman, Milton. The case for flexible exchange rates. In *Essays in Positive Economics*, pages 157–203. University of Chicago Press, Chicago, 1953.

Gabaix, Xavier. A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710, 2014. doi: 10.1093/qje/qju024.

Gelman, Michael, Liron Reiter-Gavish, and Nikolai Roussanov. FOMO economics: External reference-dependence and risk-taking in household portfolios. Working paper, SSRN, 2026.

Glebkin, Sergei, Naveen Gondhi, and John C. Kuong. Funding constraints and informational efficiency. *The Review of Financial Studies*, 34(9):4269–4322, 2021. doi: 10.1093/rfs/hhab001.

Gorodnichenko, Yuriy and Xiao Yin. Higher-order beliefs and risky asset holdings. *Available at SSRN: https://ssrn.com/abstract=4889206*, 2026.

Graves, Daniel. Attention and the retail alignment puzzle. Working paper, Yale University, 2024.

Han, Bing and Alok Kumar. Speculative retail trading and asset prices. *Journal of Financial and Quantitative Analysis*, 48(2):377–404, 2013. doi: 10.1017/S0022109013000100.

Hvidkjaer, Soeren. Small trades and the cross-section of stock returns. *Review of Financial Studies*, 21(3):1123–1151, 2008. doi: 10.1093/rfs/hhn049.

J.P. Morgan Asset Management. Who is buying U.S. equities? J.P. Morgan Asset Management Market Insights, 2025. URL https://am.jpmorgan.com/us/en/asset-management/liq/insights/market-insights/market-updates/on-the-minds-of-investors/who-is-buying-us-equities/. Accessed March 2026.

Kaniel, Ron, Gideon Saar, and Sheridan Titman. Individual investor trading and stock returns. *Journal of Finance*, 63(1):273–310, 2008. doi: 10.1111/j.1540-6261.2008.01316.

Kelley, Eric K. and Paul C. Tetlock. How wise are crowds? insights from retail orders and stock returns. *The Journal of Finance*, 68(3):1229–1265, 2013. doi: 10.1111/jofi.12028.

Kumar, Alok. Who gambles in the stock market? *The Journal of Finance*, 64(4):1889–1933, 2009. doi: 10.1111/j.1540-6261.2009.01483.x.

Kumar, Alok and Charles M. C. Lee. Retail investor sentiment and return comovements. *The Journal of Finance*, 61(5):2451–2486, 2006. doi: 10.1111/j.1540-6261.2006.01063.x.

Laarits, Toomas and Marco Sammon. The retail habitat. *Journal of Financial Economics*,

172:104144, 2025. doi: 10.1016/j.jfineco.2025.104144.

Lee, Charles M. C. and Balkrishna Radhakrishna. Inferring investor behavior: Evidence from TORQ data. *Journal of Financial Markets*, 3(2):83–111, 2000. doi: 10.1016/S1386-4181(00)00002-1.

Lou, Dong. A flow-based explanation for return predictability. *Review of Financial Studies*, 25(12):3457–3489, 2012. doi: 10.1093/rfs/hhs103.

Luo, Bryan, Enrichetta Ravina, Marco Sammon, and Luis M. Viceira. Retail investors' contrarian behavior around news, attention, and the momentum effect. Working Paper 34086, National Bureau of Economic Research, 2025.

Mitton, Todd and Keith Vorkink. Equilibrium underdiversification and the preference for skewness. *The Review of Financial Studies*, 20(4):1255–1288, 2007. doi: 10.1093/rfs/hhm011.

Odean, Terrance. Are investors reluctant to realize their losses? *The Journal of Finance*, 53(5):1775–1798, 1998. doi: 10.1111/0022-1082.00072.

Shefrin, Hersh and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, 40(3):777–790, 1985. doi: 10.1111/j.1540-6261.1985.tb05002.x.

van der Beck, Philippe, Cameron Cohen, and Coralie Jaunin. The equity market implications of the retail investment boom. Working Paper 26-054, Harvard Business School, 2026.

Vayanos, Dimitri and Paul Woolley. An institutional theory of momentum and reversal. *The Review of Financial Studies*, 26(5):1087–1145, 2013. doi: 10.1093/rfs/hht014.

Welch, Ivo. The wisdom of the robinhood crowd. *The Journal of Finance*, 77(3):1489–1527, 2022. doi: 10.1111/jofi.13128.

Yuan, Kathy. Asymmetric price movements and borrowing constraints: A rational expectations equilibrium model of crises, contagion, and confusion. *The Journal of Finance*, 60(1):379–411, 2005. doi: 10.1111/j.1540-6261.2005.00733.x.

A Additional Figures and Tables

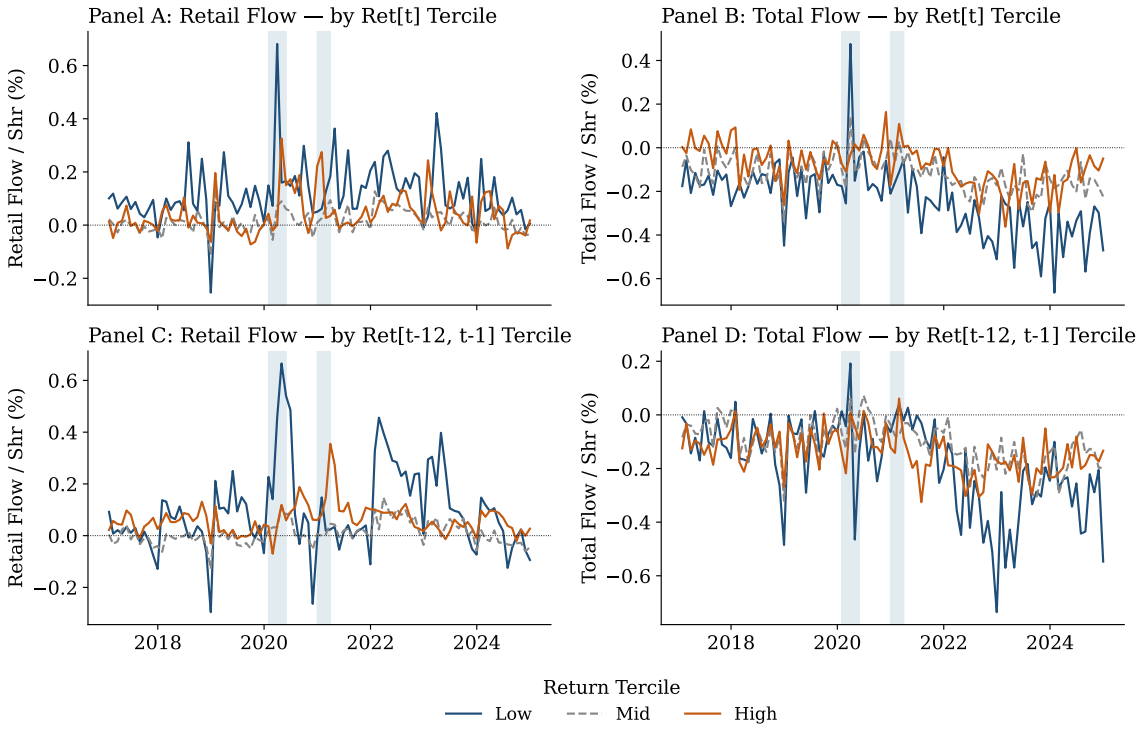


Figure A.1. Retail and Total Flows by Return Tercile, Market-Cap Weighted. Each panel plots the market-cap-weighted cross-sectional average of monthly net flow scaled by shares outstanding (percent) for stocks sorted into terciles within each calendar month. Weights are stock-level market capitalization at month t. Panels A and B sort by contemporaneous return $\text{Ret}[t]$. Panels C and D sort by past twelve-month return $\text{Ret}[t-12, t-1]$. Panels A and C plot retail flow; Panels B and D plot total initiated flow. Shaded regions mark the COVID-19 selloff (February–May 2020) and the meme-stock window (January–March 2021). The sample is January 2017 through December 2024. Figure 5 in the main text reports the equal-weighted version.

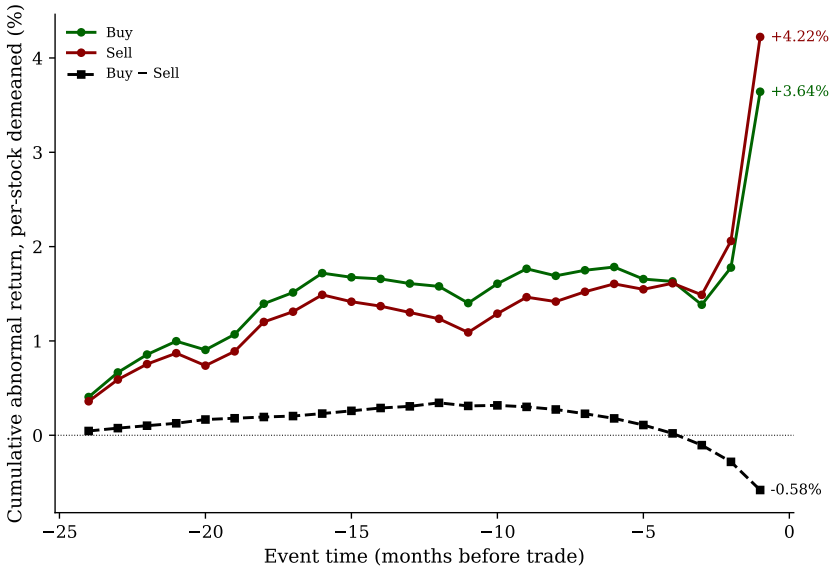


Figure A.2. Pre-Trade Cumulative Market-Adjusted Return: Retail Buys versus Sells. This figure is the panel analog of Barber et al. (2009b, Figure 3). Each stock-month is an event. The buy path weights each event by retail shares bought scaled by shares outstanding; the sell path weights by retail shares sold scaled by shares outstanding. The plotted quantity is the cumulative market-adjusted return over event months -24 through -1, where the market-adjusted monthly return is the stock return minus the value-weighted panel return that month, winsorized at the 1% and 99% percentiles and then demeaned by stock over the full sample. "Buy − Sell" is the difference of the buy and sell cumulative paths. The sample is January 2017 through December 2024.

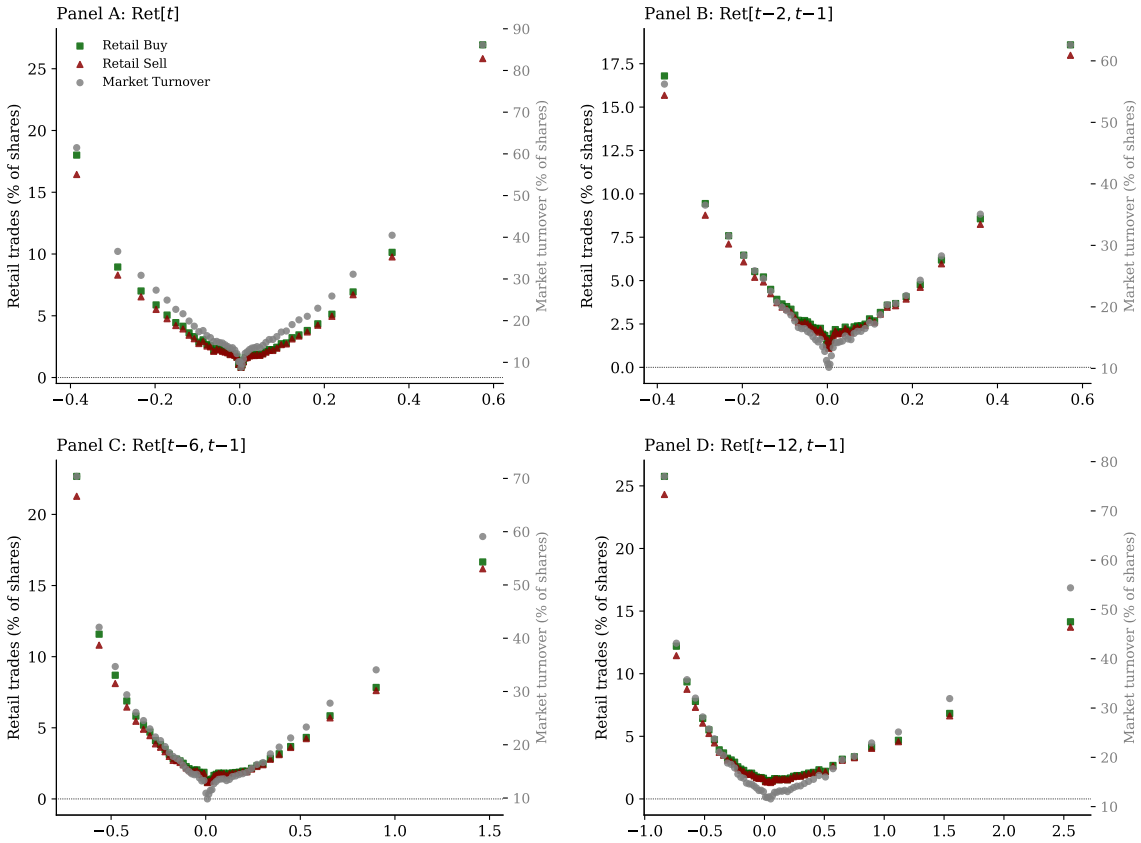


Figure A.3. Retail Buying and Selling Versus Past Returns, with Market Turnover. Each panel is a binscatter with 50 equal-frequency bins. "Retail Buy" is monthly retail shares bought and "Retail Sell" is monthly retail shares sold, each scaled by shares outstanding (percent) and read on the left axis. "Market Turnover" is total monthly share volume scaled by shares outstanding (percent), read on the right (secondary) axis. Series are pooled across stock-months. The horizontal axis is the past return over the indicated horizon. The past return and each plotted variable are winsorized at the 1% and 99% percentiles within each panel. The sample is January 2017 through December 2024.

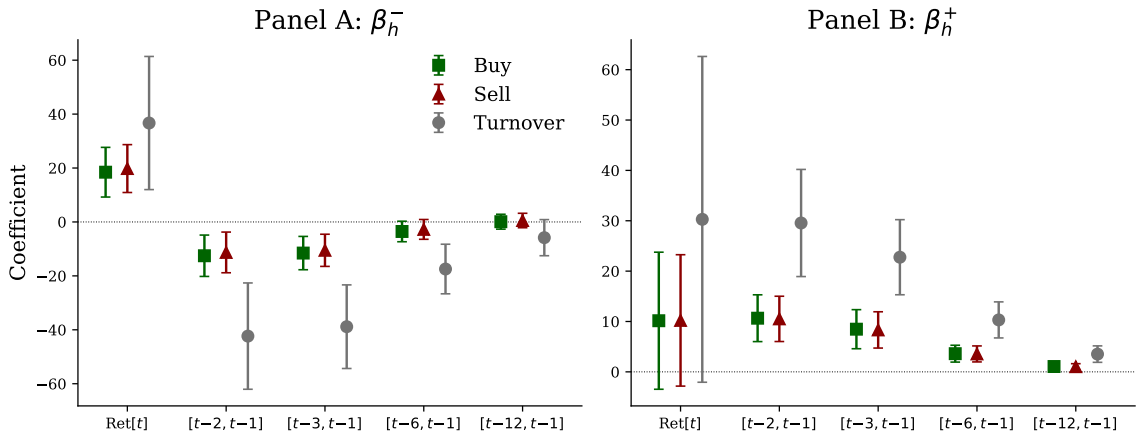


Figure A.4. Loss/Gain-Split Coefficients for Retail Buying, Selling, and Turnover. Markers report the loss-side coefficients β_h^- (Panel A) and gain-side coefficients β_h^+ (Panel B) from single-horizon panel regressions, splitting each past return into its negative and positive parts. Plotted series are retail buying, retail selling, and market turnover, each scaled by shares outstanding (percent). The net (buying minus selling) coefficients come from the same regressions; they are an order of magnitude smaller and are shown on their own scale in Figure 4. Each regression includes one horizon's negative- and positive-return parts, the composition cofntrols $X_{i,t}$ of equation (1), and stock and time fixed effects. Horizons are $h \in \{0, 1, 3, 6, 12\}$ months. Because net flow equals retail buying minus retail selling and the regressor matrix and sample are identical across the dependent variables at each horizon, the net coefficient equals the buy coefficient minus the sell coefficient by construction. Error bars are 95% confidence intervals from standard errors clustered by stock and by calendar month. Returns are not winsorized. The sample is January 2017 through December 2024.

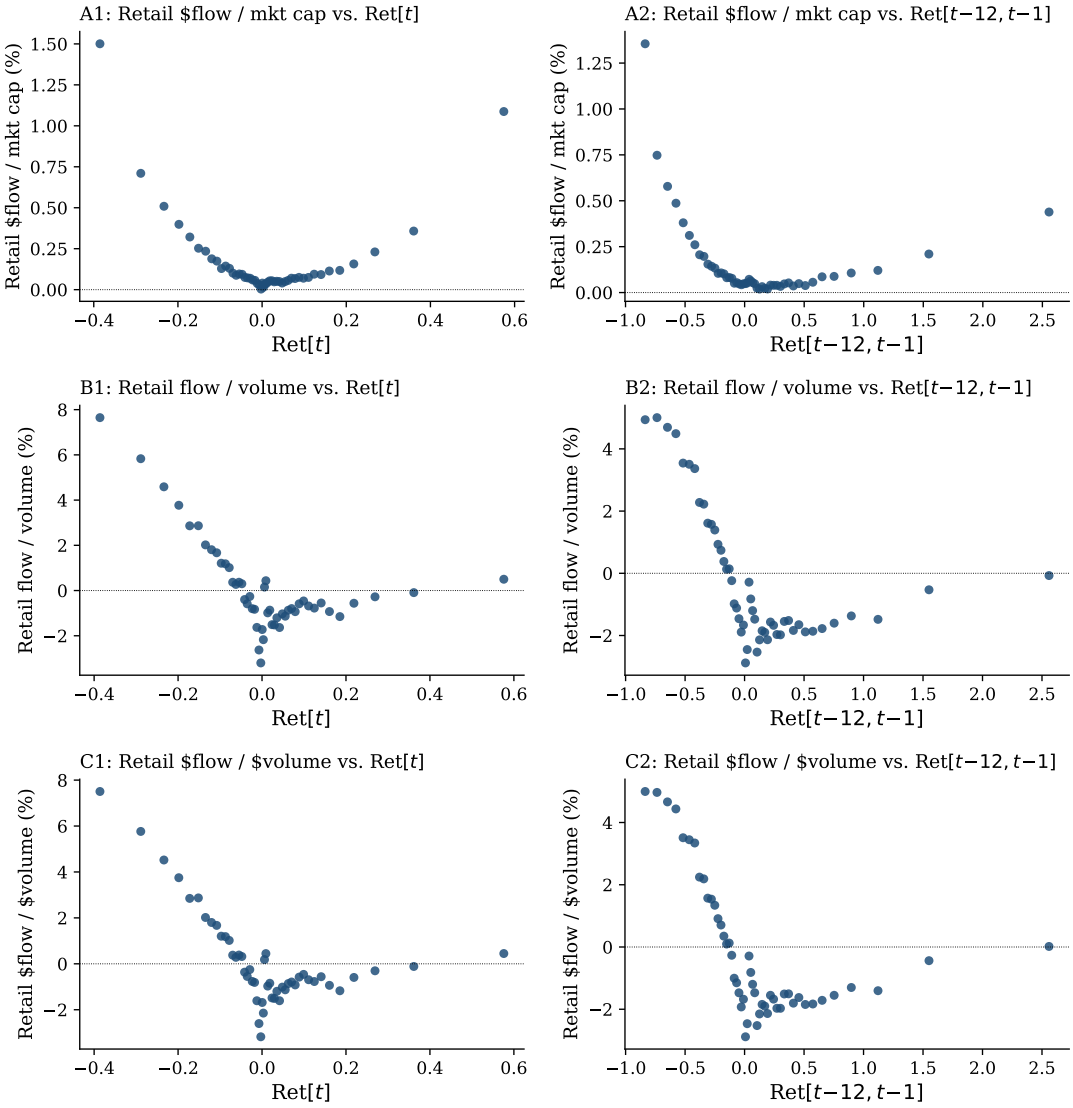


Figure A.5. Buying the Dip Under Alternative Flow Normalizations. Each panel is a binscatter with 50 equal-frequency bins. The rows use three alternative normalizations of retail net flow: retail dollar flow divided by market capitalization (Row A), retail net share flow divided by retail share volume (Row B), and retail dollar flow divided by retail dollar volume (Row C). The columns use the contemporaneous return $\text{Ret}[t]$ and the past twelve-month return $\text{Ret}[t-12, t-1]$. Both axes are winsorized at the 1% and 99% percentiles within each panel. The sample is January 2017 through December 2024.

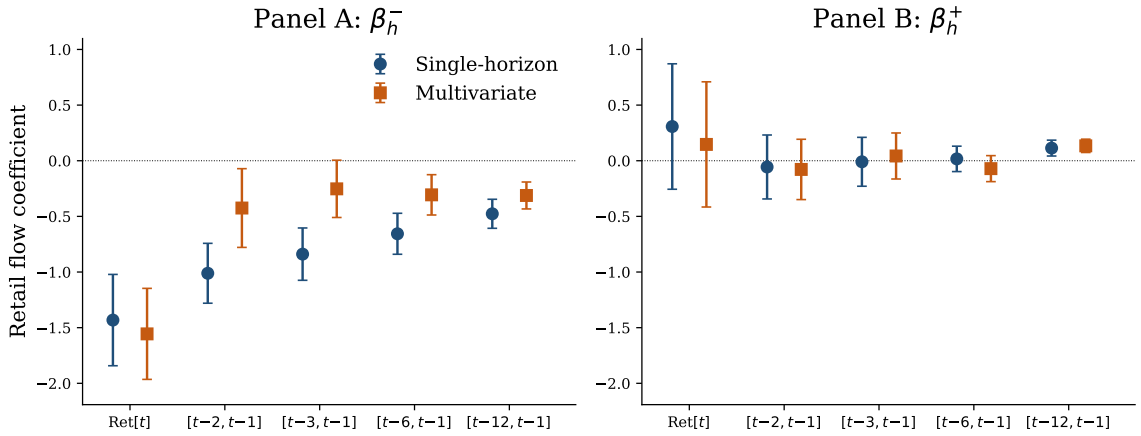


Figure A.6. Buying the Dip With Log Returns. The figure repeats the loss/gain-split coefficient plot of Figure 4 with each past return r_h replaced by its log transform $\ln(1 + r_h)$, defined where $1 + r_h > 0$. The sign split is preserved because $\ln(1 + r_h) < 0$ if and only if $r_h < 0$. Navy circles are single-horizon estimates and orange squares are the multivariate (all-horizons) estimate, each with the composition controls $X_{i,t}$ of equation (1) and stock and time fixed effects. Error bars are 95% confidence intervals from standard errors clustered by stock and by calendar month. The dependent variable is retail net flow scaled by shares outstanding (percent). The sample is January 2017 through December 2024.

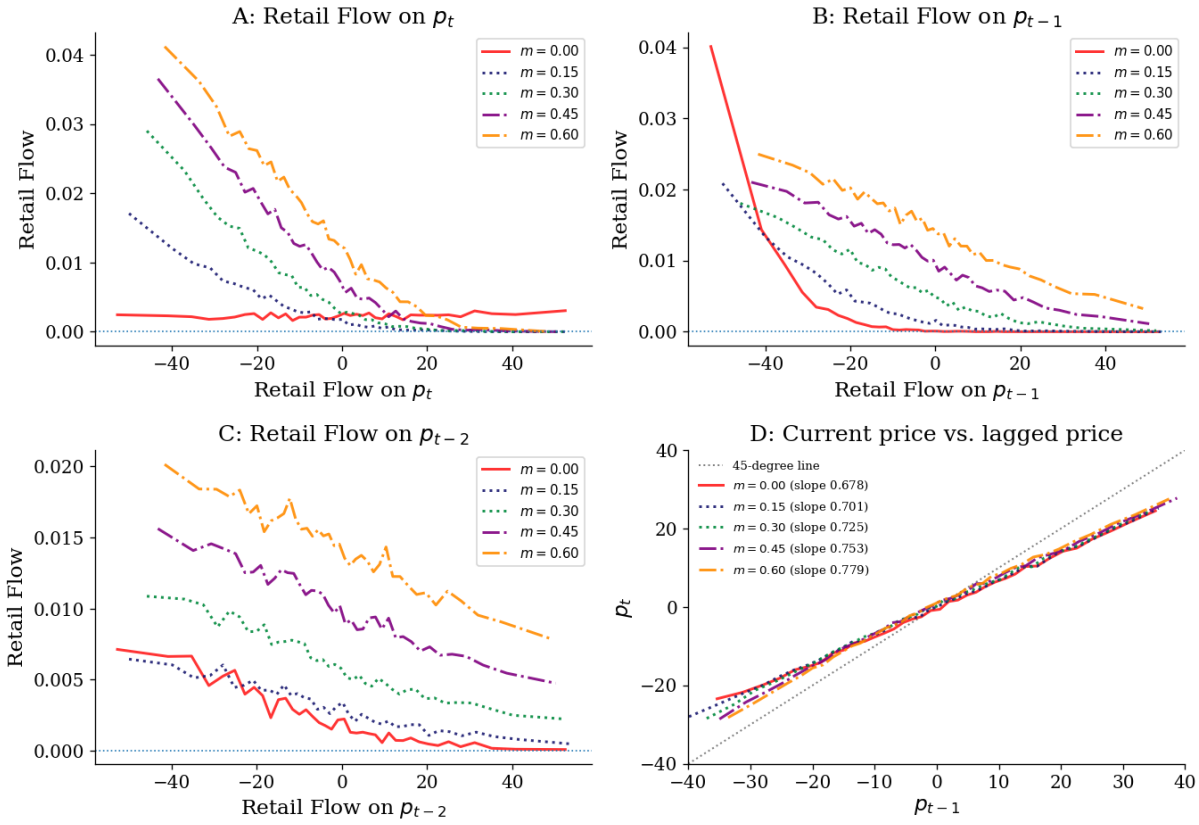


Figure A.7. Retail Order Flow and Prices. Panels A, B, and C respectively plot retail order imbalance in period t on price in period t, $t-1$, and $t-2$

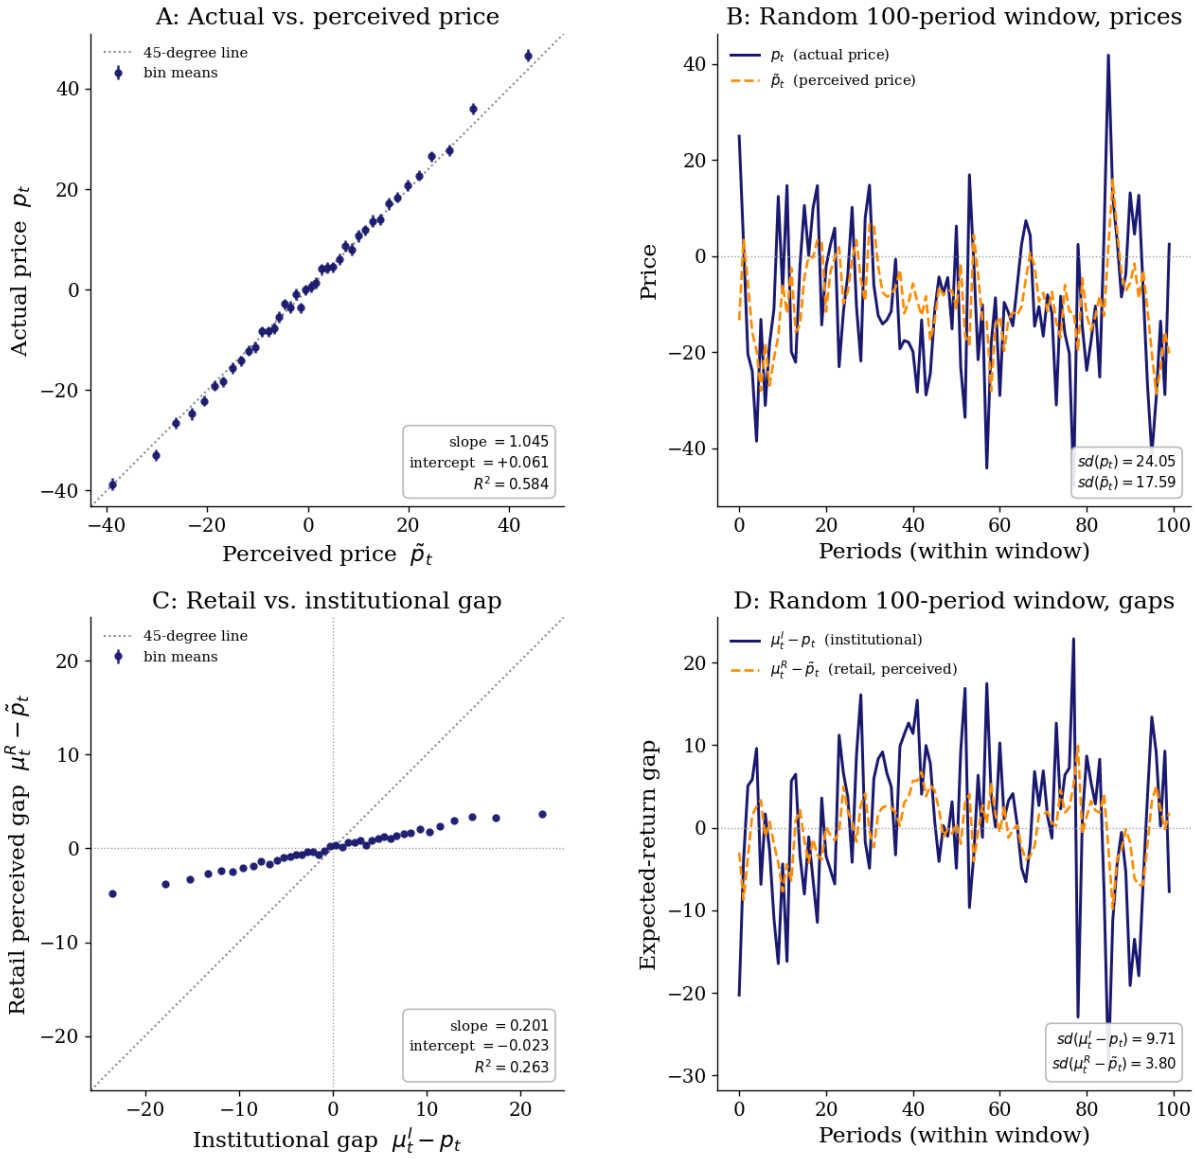


Figure A.8. Retail Beliefs and the Actual Price. The figure compares the realized price P_t to retail-belief, $E^R_{t-1}[P_t]$, from a single long simulation of the calibrated model. Both panels use a path of length $T = 10{,}000$ with the first 2,000 periods discarded as burn-in. Panel A is a binned scatter plot of P_t against retail's lagged-information forecast of the same period's price, $E^R_{t-1}[P_t]$. Panel B plots a randomly drawn 100-period window from the same simulation, showing the realized price p_t (solid) and retail's two-period-ahead lagged forecast $E^R_{t-1}[P_{t+1}]$ (dashed).

Table A.1. Buying the Dip: Panel Regressions of Retail Flow on Past Returns

Sample is January 2017 through December 2024 ($N \approx 414{,}000$ stock-months, 7,392 stocks, 96 months). The dependent variable is monthly retail net order flow as a percentage of shares outstanding (in pp). Past returns are in fractions, so a coefficient of β on $\text{Ret}[\cdot]^-$ means a 100pp move in the negative-return state shifts retail flow by β pp of shares outstanding. For any return r, $r^- \equiv \min(r, 0)$ and $r^+ \equiv \max(r, 0)$ denote its negative and positive parts; $\text{Ret}[\cdot]^-$ and $\text{Ret}[\cdot]^+$ apply this split to the indicated return. All specifications include stock and time fixed effects. Standard errors are clustered by stock and by time and reported in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	(1)	(2)	(3)	(4)
$\text{Ret}[t]^-$	-3.304***	-1.717***	-1.645***	-1.467***
	(0.258)	(0.268)	(0.272)	(0.258)
$\text{Ret}[t]^+$	1.904***	-0.203	-0.115	-0.057
	(0.298)	(0.273)	(0.299)	(0.307)
$\text{Ret}[t-2, t-1]^-$		-0.716***	-0.733***	-0.669***
		(0.191)	(0.192)	(0.193)
$\text{Ret}[t-2, t-1]^+$		0.015	0.077	0.002
		(0.125)	(0.141)	(0.129)
$\text{Ret}[t-3, t-1]^-$		-0.522***	-0.409**	-0.368**
		(0.167)	(0.172)	(0.165)
$\text{Ret}[t-3, t-1]^+$		0.071	0.088	0.149
		(0.131)	(0.155)	(0.155)
$\text{Ret}[t-6, t-1]^-$		-0.500***	-0.545***	-0.469***
		(0.117)	(0.121)	(0.112)
$\text{Ret}[t-6, t-1]^+$		-0.001	-0.013	-0.026
		(0.029)	(0.033)	(0.034)
$\text{Ret}[t-12, t-1]^-$		-0.175**	-0.171**	-0.158**
		(0.076)	(0.075)	(0.069)
$\text{Ret}[t-12, t-1]^+$		0.023*	0.026*	0.023*
		(0.012)	(0.013)	(0.013)
Idio. Volatility		23.814***	23.985***	22.663***
		(2.211)	(2.184)	(2.335)
Idio. Skewness		-0.035***	-0.041***	-0.039***
		(0.010)	(0.010)	(0.009)
Beta		-0.081***	-0.085***	-0.085***
		(0.022)	(0.022)	(0.021)
Log(ME)			0.025	0.028
			(0.032)	(0.029)
B/M			0.002	0.022
			(0.025)	(0.015)
Debt/Assets			0.170	0.145
			(0.113)	(0.098)
NI/Assets			-0.036	-0.024
			(0.034)	(0.024)
Age			-0.002	-0.004
			(0.002)	(0.004)
Gross Profitability			-0.040	-0.079*
			(0.032)	(0.040)
Δ Inst. Ownership				-0.031
				(0.027)
InstSec Flow / Shr				7.529***
				(1.613)
Inst50k Flow / Shr				0.650
				(5.071)
N	414,239	316,428	298,709	285,742
R^2 (within)	0.03	0.09	0.09	0.09
Stock FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

B Solving the Equilibrium and the Calibration Moments

This appendix derives the equilibrium solution and the model moments that the simulated method of moments in Appendix C targets. The price is the affine rule (6), and all equalities hold to first order around the stationary mean.

B.1 Retail forecasts and the demand wedge

Retail enters period t with the posterior mean $\widehat{x}_{t-1}$. Under the price rule (6), the state law of motion, and $y_{t+1} = p_{t+1} + f_{t+1}$, the tower property $E^R_{t-1}[\widehat{x}_t] = E^R_{t-1}[x_t] = F\widehat{x}_{t-1}$ gives retail's forecasts of the current price and the cum-dividend payoff,

$$\widehat{p}_t = \bar{p} + (\Gamma F + \Pi)\,\widehat{x}_{t-1}, \tag{7}$$

$$\mu^R_t = \bar{\mu}^R + \left(\Gamma F^2 + \Pi F + e_\theta' F^2\right)\widehat{x}_{t-1}. \tag{8}$$

Substituting these and the perceived price $\widetilde{p}_t = mp_t + (1-m)\widehat{p}_t$ into $w_t = \tau^R(\mu^R_t - \widetilde{p}_t) - \bar{h}$ and collecting the terms in x_t and in $\widehat{x}_{t-1}$ yields the demand wedge,

$$w_t - w_0 = -m\tau^R\Gamma x_t + \tau^R A\,\widehat{x}_{t-1}, \qquad A = \Gamma F^2 + \Pi(F - I) + e_\theta' F^2 - (1-m)\,\Gamma F.$$

The contemporaneous loading $-m\tau^R\Gamma$ collects the mp_t term in $\widetilde{p}_t$; the lagged loading A collects retail's forecasts and the $(1-m)\widehat{p}_t$ term.

B.2 Signal extraction and the Kalman filter

Define the cleaned price innovation $z_t = p_t - \bar{p} - \Pi\,\widehat{x}_{t-1}$, the part of the price not predictable from retail's priors; equation (6) gives $z_t = \Gamma x_t$. Retail observes (z_t, f_t), a linear-Gaussian system

$$z_t = \Gamma x_t, \qquad f_t = \bar{d} + e_\theta' x_t + \epsilon_t, \qquad x_t = Fx_{t-1} + u_t, \tag{9}$$

with observation matrix $H = [\,\Gamma\,;\, e_\theta'\,]$ and observation noise $\nu_t = (0, \epsilon_t)'$ of covariance $R = \operatorname{diag}(0, \sigma^2_\epsilon)$ — the price signal is noiseless, the cash flow carries ϵ_t. Retail beliefs follow the steady-state Kalman filter for (9): the predicted error covariance $P = \operatorname{Cov}(x_t - F\widehat{x}_{t-1})$ solves the discrete algebraic Riccati equation

$$P = F\big[\,P - PH'(HPH' + R)^{-1}HP\,\big]F' + \Sigma_u, \tag{10}$$

the steady-state gain is $K = PH'(HPH' + R)^{-1}$, and with $J = KH$ and $L = (I - KH)F$ the posterior update is

$$\widehat{x}_t = Jx_t + L\,\widehat{x}_{t-1} + K\nu_t, \qquad E[\widehat{x}_t \mid x_t, \widehat{x}_{t-1}] = Jx_t + L\,\widehat{x}_{t-1}. \tag{11}$$

B.3 The institutional forecast

Institutions observe x_t and the public history but trade before f_t is realized, so they condition on $(x_t, \widehat{x}_{t-1})$ and not on ϵ_t. Using (11) with $E[\nu_t] = 0$ and $E[x_{t+1} \mid x_t] = Fx_t$,

$$\mu^I_t = \bar{\mu}^I + (\Gamma F + \Pi J + e_\theta' F)\,x_t + \Pi L\,\widehat{x}_{t-1}. \tag{12}$$

B.4 The price fixed point

Linearizing retail demand around w_0 gives $Q^R_t = \bar{Q}^R + \kappa_0(w_t - w_0)$ with $\kappa_0 = G'(w_0) = \Phi(w_0/\sigma_h)$. Substituting the price rule, the institutional forecast, and the wedge into market clearing (5) and matching the coefficients on x_t and on $\widehat{x}_{t-1}$ gives

$$\left(\tau^I + \kappa_0\tau^R m\right)\Gamma = \tau^I\left(\Gamma F + \Pi J + e_\theta' F\right) - e_\lambda', \tag{13}$$

$$\tau^I\,\Pi(L - I) + \kappa_0\tau^R A = 0, \tag{14}$$

together with $J = KH$, $L = (I - KH)F$, and $H = [\,\Gamma\,;\, e_\theta'\,]$. Taking unconditional means in (5), and using $\bar{\mu}^I = \bar{\mu}^R = \bar{p} + \bar{d}$ (the stationary mean cum-dividend payoff), the stationary price satisfies $\bar{Q} = \tau^I \bar{d} + \bar{Q}^R$, and the stationary demand index reduces to $w_0 = \tau^R \bar{d} - \bar{h}$.

B.5 Stationary moments for the calibration

Collect the state and retail beliefs in the augmented vector $g_t = (x_t', \widehat{x}_t')'$. Combining $x_{t+1} = Fx_t + u_{t+1}$ with the posterior recursion (11),

$$g_{t+1} = \Phi_g\, g_t + \xi_{t+1}, \qquad \Phi_g = \begin{pmatrix} F & 0 \\ JF & L \end{pmatrix}, \qquad \xi_{t+1} = \begin{pmatrix} u_{t+1} \\ Ju_{t+1} + K\nu_{t+1} \end{pmatrix}, \tag{15}$$

with innovation covariance

$$\Omega_g = \begin{pmatrix} \Sigma_u & \Sigma_u J' \\ J\Sigma_u & J\Sigma_u J' + KRK' \end{pmatrix}.$$

The stationary covariance Σ_g solves the discrete Lyapunov equation

$$\Sigma_g = \Phi_g\, \Sigma_g\, \Phi_g' + \Omega_g, \tag{16}$$

and the autocovariances are $\mathrm{Cov}(g_t, g_{t-k}) = \Phi_g^k \Sigma_g$.

Every object the calibration uses is an affine functional of a finite window of g's and the i.i.d. shocks: the price $p_t = \bar{p} + \Gamma x_t + \Pi \widehat{x}_{t-1}$, the cum-dividend return $r^c_{t-1} = p_{t-1} + f_{t-1} - p_{t-2}$, and the demand index $w_t = w_0 - m\tau^R \Gamma x_t + \tau^R A\, \widehat{x}_{t-1}$. Their variances and covariances are therefore quadratic forms in Σ_g, the autocovariances $\{\Phi_g^k \Sigma_g\}$, Σ_u, and σ_ϵ^2. In particular,

$$\beta_w = \frac{\mathrm{Cov}(w_t, r^c_{t-1})}{\mathrm{Var}(r^c_{t-1})}, \qquad \widetilde{\sigma}^2 = \mathrm{Var}(w_t) - \beta_w^2\, \mathrm{Var}(r^c_{t-1}) + \sigma_h^2, \tag{17}$$

the second equality using the Gaussian conditional variance $\mathrm{Var}(w_t \mid r^c_{t-1}) = \mathrm{Var}(w_t) - \beta_w^2 \mathrm{Var}(r^c_{t-1})$. With w_0, β_w, and $\widetilde{\sigma}$ in hand, conditional expected retail flow is the truncated-normal mean

$$E[Q_t^R \mid r^c_{t-1} = x] = \bar{m}(x)\, \Phi\!\left(\tfrac{\bar{m}(x)}{\widetilde{\sigma}}\right) + \widetilde{\sigma}\, \phi\!\left(\tfrac{\bar{m}(x)}{\widetilde{\sigma}}\right), \qquad \bar{m}(x) = w_0 + \beta_w\big(x - E[r^c_{t-1}]\big), \tag{18}$$

which, together with the unconditional return variance read from Σ_g, supplies the model moments matched to their empirical counterparts in the SMM objective. The supply-shock impulse responses and the retail-share volatility statistics of Section 4.4 are computed by simulation from the same solved system, retaining the exact map G.

B.6 Numerical solution

The equilibrium is a joint fixed point: a guess for (Γ, Π) fixes H in (9) and, through the Riccati equation (10), the gain K; K enters (13)–(14), which return updated (Γ, Π). We iterate to convergence, then solve the Lyapunov equation (16) for Σ_g and form the moments above. The SMM loop wraps this solution in an outer search over the structural parameters.

C Calibration

This section reports the calibration of the structural parameters that determine them. The model solution algorithm and the per-exercise simulation design are described in the remainder of the online appendix.

We proceed in three steps. The first step fixes parameters that the retail-flow moments are not well suited to identify. The second step recovers two parameters, the signal noise σ_ϵ and the average hedging need $\bar{h}$, from two moment conditions that each pin down a single parameter exactly. The third step estimates the remaining three parameters, $\psi = (\tau^I, \sigma_{u,\theta}, m)$, by simulated method of moments on an exactly identified system, leaving the price-informativeness measure as an untargeted validation moment.

Fixed parameters

The two state persistences are set so that the cash-flow component is far more persistent than the supply-pressure component, the ordering on which the mechanism of Section 5.5 rests. The cash-flow persistence is fixed at $\rho_\theta = 0.96$, close to a random walk, so that cash-flow innovations are near-permanent. The supply-pressure persistence is fixed at $\rho_\lambda = 0.5$, which makes the supply-pressure state transitory, consistent with evidence that institutional flow-driven trades generate temporary price pressure that subsequently reverses (Coval and Stafford, 2007; Lou, 2012; Vayanos and Woolley, 2013). The supply-pressure innovation is normalized to $\sigma_{u,\lambda} = 1$, which fixes the unit of the supply-pressure state. The cross-sectional dispersion of hedging needs is set to $\sigma_h = 10^{-6}$, so the truncated-normal map G is numerically the hard kink $\max\{w_t, 0\}$. The retail share is $\alpha = 0.25$ and $\sigma_I = \sigma_R = 1$.

The retail risk-bearing capacity τ^R is calibrated rather than estimated. Because reported retail flow is rescaled so that its standard deviation matches the data (equation (C) below), every moment constructed from scaled flow is invariant to proportional rescalings of the retail demand index, and the flow moments are therefore locally uninformative about the retail demand scale.[7] We therefore anchor the retail risk-aversion coefficient to outside evidence. Following Beutel and Weber (2026) and Gorodnichenko and Yin (2026), we set $\gamma_R = 8$ and $\sigma_R^2 = 0.16^2$ to match the montly standard deviation of the US general stock market. This implies $\tau^R = \alpha/(\gamma_R \sigma_R^2) = 1.091$.

Measurement of returns and retail flow

The model is solved and simulated in local payoff units. A simulated path delivers the local price object p_t and the cash-flow object d_t at each date. These local objects enter the equilibrium, demand, and filtering equations. They should not be interpreted as globally positive dollar prices. In particular, because the first-order price rule is affine in Gaussian states, the local price object can take negative values in tail realizations.

To map the local solution into empirical return units, we introduce a positive

[7] Formally, the columns of the moment Jacobian corresponding to τ^R are an order of magnitude smaller than those of the estimated parameters, and the smallest singular value of the Jacobian loads almost entirely on τ^R; treating it as a free parameter produces standard errors that reflect this near-flatness rather than sampling information.

measurement layer. Let $\bar{p}$ and $\bar{d}$ denote the stationary means in the local model, and let $\bar{P} > 0$ and $\bar{D} > 0$ denote measurement-level price and cash-flow scales. Define

$$\mathcal{P}_t = \bar{P} \exp\left(\frac{p_t - \bar{p}}{\bar{P}}\right), \qquad \mathcal{D}_t = \bar{D} \exp\left(\frac{d_t - \bar{d}}{\bar{D}}\right).$$

The affine local price p_t can therefore remain Gaussian, while the measured traded price $\mathcal{P}_t$ is always positive. The exact measured cum-dividend log return is

$$r_t^\ell = \log \frac{\mathcal{P}_t + \mathcal{D}_t}{\mathcal{P}_{t-1}}.$$

For calibration and simulation we use the first-order log-linear approximation around $(\bar{P}, \bar{D})$. Let $\rho = \frac{\bar{P}}{\bar{P}+\bar{D}}$.

Then

$$r_t^\ell - \bar{r}^\ell \approx \rho \frac{p_t - \bar{p}}{\bar{P}} + (1-\rho)\frac{d_t - \bar{d}}{\bar{D}} - \frac{p_{t-1} - \bar{p}}{\bar{P}}, \tag{19}$$

where $\bar{r}^\ell = \log\left(1 + \frac{\bar{D}}{\bar{P}}\right)$.

All regression and SMM moments use the demeaned log-linear return in (19). The constant $\bar{r}^\ell$ is omitted because the calibration targets slopes, standard deviations, correlations, and R^2s.

We fix the steady dividend yield $y_D = \frac{\bar{D}}{\bar{P}}$ at the model frequency and concentrate out the price scale $\bar{P}$ in each SMM evaluation. Equivalently, for a fixed y_D, equation (19) can be written as

$$r_t^\ell - \bar{r}^\ell = \frac{1}{\bar{P}}\left[\rho(p_t - \bar{p}) + \frac{1-\rho}{y_D}(d_t - \bar{d}) - (p_{t-1} - \bar{p})\right], \qquad \rho = \frac{1}{1+y_D}.$$

Thus $\bar{P}$ is a measurement denominator. For every candidate structural parameter vector, we choose $\bar{P}$ so that the simulated standard deviation of r_t^ℓ equals the empirical target,

$$sd(r_t^\ell) = 0.183. \tag{20}$$

The corresponding $\bar{D}$ is then given by $\bar{D} = y_D \bar{P}$. In the numerical implementation, values such as $\bar{P} = 100$ and $\bar{D} = 2$ are only initial measurement scales; the calibration uses the volatility target in (20) to put returns in empirical units.

Retail flow is scaled separately. Let $\widetilde{Q}_t^R$ denote the simulated retail order imbalance generated by the nonlinear market-clearing condition. We define the reported retail flow as

$$Q_t^R = a_Q \widetilde{Q}_t^R, \qquad a_Q = \frac{0.014}{sd(\widetilde{Q}_t^R)}.$$

Hence the scaled retail-flow series satisfies

$$sd(Q_t^R) = 0.014.$$

The return scale $\bar{P}$ and the flow scale a_Q are measurement normalizations, not structural parameters. They convert local payoff-unit quantities into empirical return and normalized-flow units.

Recovering σ_ϵ and $\bar{h}$

Two parameters are recovered from two moment conditions that each hold exactly, in the same spirit as the scale normalizations above.

First, because the cash-flow process is exogenous, the autocorrelation of the public cash-flow signal identifies the signal noise in closed form. With $f_t = \bar{d} + \theta_t + \epsilon_t$ and θ_t an AR(1),

$$\operatorname{corr}(f_t, f_{t-1}) = \frac{\rho_\theta \, \sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2}, \qquad \sigma_\theta^2 = \frac{\sigma_{u,\theta}^2}{1 - \rho_\theta^2}, \tag{21}$$

which inverts to $\sigma_\epsilon^2 = \sigma_\theta^2 \, (\rho_\theta - \operatorname{corr}) / \operatorname{corr}$. We impose (21) at the empirical value $\operatorname{corr}(d_t, d_{t-1}) = 0.820$ inside every SMM evaluation, so the autocorrelation is matched exactly by construction and σ_ϵ moves with $\sigma_{u,\theta}$ during the estimation. At the point estimate $\widehat{\sigma}_{u,\theta}$ below, the implied value is $\widehat{\sigma}_\epsilon = 0.326$.

Second, the average hedging need $\bar{h}$ is recovered from the level of retail flow relative to its dispersion. The ratio $E[Q_t^R]/sd(Q_t^R)$ is invariant to the flow scale a_Q, since the scale multiplies numerator and denominator identically, and it is the natural moment for an additive shifter of the demand index: a higher $\bar{h}$ pushes more of the retail sector against the short-sale constraint, lowering mean flow relative to its dispersion. Under the stationary Gaussian distribution of the demand index, the ratio is a strictly increasing function of the standardized constraint location $z = w_0/\widetilde{\sigma}_w$ and maps $\mathbb{R}$ onto $(0, \infty)$,[8] so the moment condition

$$\frac{E[Q_t^R]}{sd(Q_t^R)} = 0.15 \tag{22}$$

identifies $\bar{h}$ uniquely given the remaining parameters. The empirical target of 0.12 is the firm-level average of this ratio: for each firm we compute the within-firm mean and standard deviation of retail flow and take their ratio, and we average the ratio across firms. Holding the estimated parameters below fixed, we solve (22) for $\bar{h}$ by root-finding on the simulated model, which yields $\widehat{h} = 8.11$. Consistent with the sizable truncation implied by the flow-level target, the short-sale constraint binds for the aggregate retail sector in the large majority of periods, the same force that generates the flat gain side of the conditional-flow schedule.

Simulated Method of Moments

The third step estimates the three remaining structural parameters $\psi = (\tau^I, \sigma_{u,\theta}, m)$ by simulated method of moments. We estimate the institutional risk-bearing capacity

[8] With $Q = \max\{w, 0\}$ and $w \sim N(w_0, \widetilde{\sigma}_w^2)$, the ratio equals $\psi(z)/\sqrt{v(z)}$ with $\psi(z) = z\Phi(z) + \varphi(z)$ and $v(z) = (z^2+1)\Phi(z) + z\varphi(z) - \psi(z)^2$. Its logarithmic derivative is proportional to $\Phi(z)^2 - \varphi(z)\psi(z) = \Phi(z)^2 \operatorname{Var}(X \mid X < z) > 0$ for a standard normal X, so the ratio is strictly increasing in z and the inversion admits a unique solution.

τ^I directly and recover the underlying risk-aversion coefficient from $\tau^I = \frac{1-\alpha}{\gamma_I \sigma_I^2}$.

The return and flow scales, the signal noise, and the average hedging need are concentrated out within every SMM evaluation as described above, so they add no dimensions to the search.

The estimator targets three moments. The first and second are the loss-side and gain-side slopes, β^- and β^+, from the piecewise regression of retail flow on the lagged log-linear return,

$$Q_t^R = a + \beta^- \min(r_{t-1}^\ell, 0) + \beta^+ \max(r_{t-1}^\ell, 0) + u_t.$$

This is the model analogue of the empirical piecewise specification. The third moment is the explanatory power of the public cash-flow signal for returns, $R^2(r_t^\ell \sim d_t)$.

With three moments and three parameters the system is exactly identified, and the estimator drives the moment distance to zero: each targeted moment is matched exactly at the estimate. Each moment is most informative about one part of the model: the loss-side slope about the institutional risk-bearing capacity τ^I, the return–cash-flow R^2 about the cash-flow innovation scale $\sigma_{u,\theta}$, and the shape of the flow response across the two return regions about the attention parameter m.

The estimator minimizes the diagonally weighted moment distance

$$\widehat{\psi} = \arg\min_{\psi} \frac{1}{3} \sum_{j=1}^{3} \left(\frac{m_j(\psi) - m_j^{\text{data}}}{\omega_j} \right)^2,$$

with weights $\omega = (0.0125, 0.0125, 0.01)$. The model is solved by an iterative simulate-and-regress fixed point. The solution algorithm, the weighting, and the common-random-numbers scheme are described in the Online Appendix.

The estimated attention parameter is $\widehat{m} = 0.036$, close to the full-inattention benchmark: retail trades almost entirely on lagged information, now as an estimate rather than an assumption.