Retail Herding and Temporary Price Pressure: Effects of Retail Popularity on Stock Returns

Vincenzo Antonio Di Buono

Dissertation written under the supervision of Professor Tural Karimli

Dissertation submitted in partial fulfilment of requirements for the MSc in Economics, at Universidade Católica Portuguesa and for the MSc in Economics and Finance at Università degli Studi di Napoli Federico II, 10/09/2025

Retail Herding and Temporary Price Pressure: Effects of Retail Popularity on Stock Returns

Vincenzo Antonio Di Buono
Universidade Católica Portuguesa

Abstract

This dissertation studies whether investor attention translates into retail buying ("herding") that temporarily moves prices. I construct a holdings-based proxy for abnormal popularity (ABPOP) using Robintrack's data on the number of distinct retail accounts that hold each stock. ABPOP activates when the cumulative net inflows per share are high upon closure, with respect to the stock's rolling within-firm median for moving windows (one, three, or six months). Combining this information with CRSP, I examine the sequence attention → (flows) → herding → price pressure → attenuation. First, stock-day regressions suggest that short-term salience (the absolute abnormal one-week return) and trading intensity relative to each stock's own history (standardized abnormal turnover, SAT) predict increases in ABPOP. Second, sorting weekly portfolios on ABPOP (12 buckets; value-weighted), the impulse–fade pattern is present: the long–short (top 5% minus bottom 5%) delivers +1.20% in Week-1 and +0.92% in Week-2, and it dissipates in Weeks 3–4, with a lack of positive drift at quarterly frequencies. Third, the factor models show that a one-month momentum factor (SR_MOM) becomes the dominant source of performance, while the classic 12–2 momentum factor (UMD) is negligible, and the alphas are almost zero. Results are not sensitive to equal-weighting, changing the extremes to 0–10%/90–100%, an alternative market benchmark for abnormal returns, or longer ABPOP windows (which produce smoother, marginally smaller profiles). Taken together, ABPOP acts as a holdings-based measure of retail herding that generates temporary price pressure rather than a conditional risk premium.

Keywords: retail investors; attention; herding; price pressure; momentum; holdings data.

Retail Herding and Temporary Price Pressure: Effects of Retail Popularity on Stock Returns

Vincenzo Antonio Di Buono

Universidade Católica Portuguesa

Resumo

Esta dissertação investiga se a atenção dos investidores se traduz em compras de retalho ("efeito manada") que movem temporariamente os preços. Construo um proxy baseado em posições para a popularidade anormal (ABPOP) a partir dos dados da Robintrack sobre o número de contas de retalho por ação. O ABPOP ativa-se quando os influxos líquidos cumulativos por ação excedem a mediana móvel da firma em janelas de 1, 3 ou 6 meses. Combinando estes dados com a CRSP, examino a sequência atenção → (fluxos) → manada → pressão de preços → atenuação. Primeiro, regressões ação-dia mostram que a saliência de curto prazo (retorno semanal anormal absoluto) e a intensidade de negociação relativa ao histórico da ação (SAT) antecipam aumentos do ABPOP. Segundo, em carteiras semanais ordenadas por ABPOP (12 grupos; valor-ponderadas), observa-se padrão de impulso-desvanecimento: a long–short (5% superior − 5% inferior) rende ~+1,20% na Semana 1 e +0,92% na Semana 2; dissipa-se nas Semanas 3–4, sem deriva positiva trimestral. Terceiro, nos modelos fatoriais, o momentum de 1 mês (SR_MOM) explica o desempenho; o 12–2 (UMD) é residual e as alfas ficam próximas de zero. Os resultados são robustos a ponderação igual, extremos 0–10/90–100, outro índice de mercado e janelas de ABPOP mais longas (perfis mais suaves). No conjunto, o ABPOP comporta-se como medida baseada em posições do efeito manada que gera pressão temporária nos preços, não um prémio de risco condicional.

Palavras-chave: investidores de retalho; atenção; efeito manada; pressão de preços; momentum; dados de posições.

Acknowledgments

First, I would like to express my sincere gratitude to my supervisor and thesis advisor, Prof. Tural Karimli, for his guidance throughout all stages of this work. I am also deeply grateful to Prof. Saverio Simonelli, Coordinator of the Master's Degree Program in Economics and Finance at the University of Naples Federico II, and to the entire faculty of the program. Thanks to this program, I was able to participate in the double-degree program and spend a wonderful year at Católica Lisbon School of Business and Economics; my thanks also go to the excellent professors I met during my second academic year in Lisbon. A special thanks to my friends, fellow students, and housemates, who shared with me the challenges and the good and difficult moments of this year. Finally, my gratitude goes to my whole family, and in particular to my parents and my sister: without you, none of this would have been possible.

Contents

1. Introduction ... ***5***

2. Data ... ***8***

2.1. Variable Specifications ... *8*

2.2. Summary Statistics ... *10*

2.3. Correlations ... *13*

3. Methodology ... ***15***

3.1. Predictive specifications for ABPOP ... *16*

3.2. Portfolio construction on ABPOP and measurement of returns ... *17*

3.3. Factor exposures ... *18*

3.4. Robustness and diagnostics ... *18*

4. Empirical Results ... ***19***

4.1. Regressions ... *20*

4.2. Herding → Returns: Portfolios on ABPOP ... *20*

4.3. Factor exposures ... *22*

4.4. Robustness ... *22*

5. Conclusions ... ***26***

References ... ***31***

Appendix ... ***32***

AI Usage Disclaimer ... ***34***

1. Introduction

Over the past decade, the diffusion of zero-commission brokers and mobile trading platforms has made access to equity markets markedly more democratic for retail investors. Anyone with an internet connection can now invest with minimal frictions. This development has significantly increased the influence of retail investors and their presence in the markets. In this environment, some stocks periodically become popular among individual investors, exhibiting rapid increases in the number of accounts that hold them. When many small investors buy simultaneously over a short window, the resulting order imbalance can push prices away from fundamentals; as these flows abate, valuations tend to mean-revert. This study examines whether abnormal retail popularity, measured from holdings data as the cumulative change in the number of distinct retail owners, serves as a proxy for retail herding and whether such episodes have systematic short-horizon return implications. Two ideas motivate the analysis. First, given limited time and information, individual investors rely on salience to construct their purchase sets and tend to buy attention-grabbing names: stocks with unusual recent returns, elevated turnover, or heavy media coverage[1]. A large empirical literature shows that attention shocks attract net retail inflows, prices rise in the following days or weeks, and a partial reversal emerges at longer horizons[2]. Second, the herding literature documents that attention-induced trading by investor groups can move prices; in the case of institutions, such flows often appear information-driven and need not reverse, as they may accelerate the adjustment of prices toward fundamentals[3]. A now classic strand of research links attention spikes to net retail flows and short-run return patterns featuring an initial push followed by partial reversal. Barber and Odean (2008) show that individual investors, when buying, must choose from a vast universe of stocks; to reduce complexity, they purchase what "glitters", names with extreme price moves, unusual volume, or extensive media coverage, whereas selling occurs within the narrower set of stocks already held. This search asymmetry generates net inflows into high-attention stocks and temporary price pressure, consistent with positive short-run returns and subsequent normalization. Da, Engelberg, and Gao (2011) propose an innovative retail-attention indicator based on abnormal Google Search Volume for individual tickers: when search intensity deviates positively from its recent norm, the stock exhibits positive returns in the

[1] Barber, B. M., & Odean, T. (2008). *All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investor.*
[2] Da, Z., Engelberg, J., & Gao, P. (2011). *In Search of Attention.*
[3] Wermers, R. (2002). *Mutual Fund Herding and the Impact on Stock Prices.*

following weeks and a correction over subsequent months, consistent with non-informational buying pressure. Gervais, Kaniel, and Mingelgrin (2001) document that shocks to trading volume (daily or weekly) precede positive short-run returns, suggesting that abnormal volume acts as a beacon of attention that attracts additional buyers. These studies converge on the mechanism attention → retail inflows → price. The proxy developed here sits closer to the portfolio decision than indirect attention measures. Unlike trading volume, media coverage, or search indices, which record traces of attention (what is discussed or clicked) but not necessarily its translation into portfolio positions, my measure is directly holdings-based and therefore one step further down the chain attention → retail inflows → price. It does not merely mean a stock is newsworthy; it means that the stock has made its way into retail portfolios. Methodologically, I define abnormal popularity as a windowed measure of holdings, the cumulative net inflow of retail accounts across a fixed horizon (say, one month, one quarter, two quarters), scaled by shares outstanding, and relative to the stock's rolling median of the same cumulative inflow. Normalization removes the heterogeneity arising from the size of the issue, allowing a comparison of small-stocks and big-stocks. Overall, the measure is focused more closely towards attention-induced retail buying (herding) rather than its informational or visibility determinants, and its windowed construction is designed to capture horizons over which retail flows are plausibly able to move prices. This approach is related (but orthogonal) to recent event-style evidence based on Robintrack holdings. Barber et al. (2022) document attention-driven trading and sharp event-day price reactions when holder counts surge, exploiting daily shocks to attention or trading intensity; the effects are immediate and fade as attention wanes. In contrast, the present measure does not observe a one-day flare. It turns on only when purchases are repeated and new positions are held in portfolios for part of the period, documenting retail accounts have built and held positions, vis à vis the stock's own recent norm. From a microstructure point of view, serially correlated order flow (i.e., attention-induced buying clustering across multiple days) leads to slow price pressure distributed over weekly horizons rather than at a point in time. Window aggregation mitigates measurement noise and day-to-day timing noise (for example, reporting lags, day-of-week effects), focusing on episodes when retail interest is systematic and persistent as opposed to episodic. The institutional-herding literature provides some correction here: when large investors move together, price effects can be more enduring, as they are frequently information-based. Nofsinger and Sias (2002) associate changes in institutional ownership with bigger price moves compared to retail or individual counterparties and argue that institutions engage in positive-feedback trading with stronger effects. For mutual funds, Wermers (2002) defines herding as the quarterly divergence between buying and selling funds and demonstrates that while not pervasive, herding is significant in e.g., small–cap and growth; when funds buy together, abnormal performance is positive over the short to medium term (and the opposite for coordinated sales). It is important to note the difference for the interpretations of the results here: if abnormal popularity mainly reflects retail herding, which is often a result of attention, salience, and imitation, we would only predict rapid and reversible effects (price pressure) instead of long-term, information-based patterns observed on institutional herding. Together, those ingredients situate the paper at the nexus of two literatures, retail attention and price pressure (Barber & Odean, 2008; Da, Engelberg & Gao, 2011; Gervais, Kaniel & Mingelgrin, 2001) and institutional herding and price dynamics (Nofsinger & Sias, 2002; Wermers, 2002), thus offering a view that complements that of platform-level attention shocks (see e.g., Barber et al., 2022). More narrowly focused on the question of cumulative holdings inflows per share, and using the stock itself as a benchmark, the study tries to find evidence for or against the hypothesis that if retail participation (not just one shiny day of it) has short-run systematic price effects, how quickly or not-so-quickly they wear away. The contribution is threefold. First, I introduce and operationalize a holdings-based, windowed measure of abnormal popularity: The cumulative change over some fixed window (one, three, and six months) in the number of retail holders per share, normalized by the stock's within-firm rolling median of cumulative changes. This design is different from both indirect attention proxies (volume, news, searches) and event-day attention shocks (e.g., Barber et al., 2022). Second, I position the proxy within the existing toolkit, showing near-orthogonality to market factors and stronger co-movement with short-horizon attention/trading variables (turnover, SAT[4], absolute abnormal returns), thereby establishing incremental behavioral content beyond volume/news/search and daily spikes. Third, I study return implications via portfolio sorts and multifactor regressions (FF3, FF5, classic and s*h*ort-term momentum), documenting positive short-run excess returns followed by attenuation, consistent with temporary price pressure from serially correlated retail order flow; results are robust to weighting schemes, alternative windows, and outlier treatments. The remainder of the thesis is organized as follows. Section 2 describes data and measures, the stock universe and sample period, and the construction of abnormal popularity (shares-outstanding normalization and median

[4] For a more comprehensive discussion of the variable, see Chordia, Huh, and Subrahmanyam (2007).

benchmarks over 1m/3m/6m), alongside descriptive statistics. Section 3 outlines the methodology: predictive regressions for popularity (recent returns, absolute abnormal returns, turnover, SAT), portfolio formation via a long–short strategy on abnormal popularity, and multifactor regressions. Section 4 presents the empirical results. Section 5 concludes with implications for retail behavior and asset pricing, key limitations, and avenues for future research.

2. Data

Robintrack records, at an hourly cadence, how many Robinhood users hold each stock listed on the platform[5]. I aggregate these counts to obtain a daily panel of retail participation: for each stock–day observation I compute the total number of accounts holding the stock at the end of the trading day. Denote this series by $p_{i,t}$, the number of Robinhood accounts that hold stock i on date t. Because Robinhood's user base is widely acknowledged to be predominantly retail, both $p_{i,t}$ and its changes are interpretable as direct proxies for retail participation at the stock level. The measure is aggregated (account counts, not share quantities) and is subject to occasional noise measurement or short coverage gaps; throughout, I adopt conservative conventions (rolling windows and median baselines) to mitigate these issues. I merge the Robintrack panel with CRSP (Center for Research in Security Prices), from which I obtain adjusted price and daily return $r_{i,t}$, trading volume $vol_{i,t}$, shares outstanding $SO_{i,t}$, and stable identifiers (PERMNO), as well as security and exchange codes (shrcd, siccd, hexcd). I also use the daily returns of CRSP's value-weighted and equal-weighted portfolios, denoted r_t^{VW} and r_t^{EW}, as well as the S&P 500 r_t^{SP}. For asset-pricing tests in subsequent sections, I employ the standard Fama–French 3/5 factors, the classic 12-2 momentum factor, a short-term momentum factor based on last-month winners vs. losers, and the risk-free rate. The final dataset is a stock–day panel with roughly 3.3 million observations and 6,479 unique stocks, balanced across NYSE/NASDAQ/AMEX. The sample spans June 2018 to August 2020.

2.1. Variable Specifications

The core of the empirical design is a holdings-based, windowed measure of retail inflows, which I label abnormal popularity. The construction starts from the change in holders and its normalization by shares outstanding

[5] For the dataset, see the link https://www.kaggle.com/datasets/cprimozi/robinhood-stock-popularity-history

$$\Delta p_{i,t} = p_{i,t} - p_{i,t-1}\,; \qquad f_{i,t} = \frac{\Delta p_{i,t}}{SO_{i,t}}$$

I then cumulate normalized flows over a fixed window H (one, three, and six months measured in trading days, $H \in \{21,63,126\}$):

$$F_{i,t}^{(H)} = \sum_{h=0}^{H-1} f_{i,t-h}$$

I compare the current cumulative inflow with the stock's own recent history by using a rolling median as the baseline:

$$B_{i,t}^{(H)} = MEDIAN\,\{F_{i,t-j}^{(H)} : j = 1, \dots, H\}$$

The main variable, abnormal popularity on window $H,$ is the deviation of the current cumulative flow from this baseline:

$$ABPOP_{i,t}^{(H)} = F_{i,t}^{(H)} - B_{i,t}^{(H)}$$

By normalizing by $SO_{i,t}$ the measure is comparable across firm sizes, so that a thousand new holders in a micro-cap do not carry the same weight as a thousand in a mega-cap; by aggregating over H acts as a persistence filter, activating only when buying repeats over multiple days, that is, when order flow is serially correlated, the pattern most consistent with temporary price pressure spread over weekly horizons. Moreover, the use of a median baseline reduces the influence of rare explosive episodes and reporting jitter, aligning the signal with sustained retail interest rather than one-off spikes. In the dataset, these measures appear as ABPOP_1m ($ABPOP_{i,t}^{(1m)}$), ABPOP_3m ($ABPOP_{i,t}^{(3m)}$), and ABPOP_6m ($ABPOP_{i,t}^{(6m)}$). For context, I also retain the level of holders pop ($p_{i,t}$), which summarizes the breadth of the retail base at each point in time, and its first difference delta_pop ($\Delta p_{i,t}$) which captures daily net entries/exits at the account level and its normalized version $f_{i,t}$. On the return side, to track price movements at matching horizons, I compute compounded cumulative returns over one and two weeks, and over one and three months, using daily CRSP returns. These appear as ret_1w, ret_2w, ret_1m, and ret_1Q and will be used both descriptively and as lagged controls in predictive specifications. The daily stock return itself is recorded as ret and is the basic unit used to build all cumulative and abnormal-return objects. Market-wide conditions enter through three benchmarks. The CRSP value-weighted market return, vwretd (r_t^{VW}), emphasizes large-cap variation; the equal-weighted return, ewretd (r_t^{EW}), provides a small-cap-sensitive aggregate; and the S&P 500 return, sprtrn (r_t^{SP}), offers a widely used large-cap index proxy. These series serve two purposes in this chapter: (i) they supply consistent market references for timing and alignment; and (ii) they allow the construction of abnormal returns by netting out broad market moves. Because the empirical design emphasizes short horizons, I construct a weekly absolute abnormal return to capture price salience irrespective of sign. Motivated by evidence in Barber and Odean (2008) that attention spikes follow both large gains and large losses, I define abs_abn_ret as the absolute difference between the stock's one-week compounded return and the market's one-week compounded return:

$$AABS_{i,t} = |\, R_{i,t} - R_{M,t} \,|$$

I capture trading activity with turnover, measured as the fraction of outstanding shares exchanged on day t:

$$turnover_{i,t} = \frac{vol_{i,t}}{SO_{i,t}}$$

Because CRSP reports shrout in thousands, I convert to shares when computing $SO_{i,t}$ to ensure dimensional consistency in both turnover and the normalized flow $f_{i,t.}$For comparability across stocks and through time, I standardize turnover into the Standardized Abnormal Turnover (SAT) of Chordia, Huh, and Subrahmanyam (2007), computed as a rolling z-score using only past information (window L= 60 trading days). I therefore define $\mu_{i,t-1}^{(L)}$ and $\sigma_{i,t-1}^{(L)}$ as the 60-day rolling mean and standard deviation of $turnover_{i,t}$ respectively, and construct $SAT_{i,t}$ as:

$$SAT_{i,t} = \frac{turnover_{i,t} - \mu_{i,t-1}^{(L)}}{\sigma_{i,t-1}^{(L)}}; \quad where \quad L = 60$$

To conclude, in this section I introduce the variables used in the analysis, covering: (i) retail participation (pop, delta_pop, ABPOP_1m/3m/6m); (ii) short-horizon price dynamics (ret, ret_1w/2w/1m/1Q); (iii) market benchmarks (vwretd, ewretd, sprtrn); (iv) trading activity (turnover and its standardized version SAT); and (v) price salience (abs_abn_ret). These definitions are self-contained and will be used as the building blocks for the descriptive tables and correlation analysis presented immediately after, and for the predictive and asset-pricing tests that follow.

2.2. Summary Statistics

This section summarizes the main features of the panel. Table 1 reports standard moments and quantiles (count, mean, standard deviation, 25th/50th/75th percentiles, minimum/maximum, skewness, kurtosis). Sample sizes differ across variables because some constructions require history (e.g., SAT needs a 60-day lookback; multi-period returns and abnormal popularity require 5–126 prior trading days depending on the horizon). Unless otherwise noted, returns are in percent per period; turnover is in percent of shares outstanding.

Table 1. Summary Statistics

This table reports summary statistics for the daily stock–day panel used in the study. Variables are: vol (CRSP raw share volume); turnover (= vol / SO, percent of shares outstanding); SAT (60-day rolling z-score of turnover, computed using information up to $t-1$); ret (CRSP daily stock return, percent per day); vwretd, ewretd, sprtrn (CRSP value-weighted, equal-weighted, and S&P 500 market returns, percent per day); abs_abn_ret (absolute difference between the stock's one-week compounded return and the market's one-week compounded return); ret_1w/ret_2w/ret_1m/ret_1Q (stock's compounded returns over 1 week, 2 weeks, 1 month, and 1 quarter, built from daily CRSP returns, percent per period); pop (Robintrack retail-holder count); delta_pop (daily change in pop); ABPOP_1m/3m/6m (cumulative net inflow of holders per share over one, three, and six months minus the stock's rolling median of the same). SO (CRSP shares outstanding) is used to compute turnover and to scale ABPOP but is not tabulated. Observation counts differ across variables due to required lookback windows. Formal definitions of all variables are provided in Section 2.1 (Variable Specifications).

	count	mean	std	min	p25	p50	p75	max	skewness	kurtosis
vol	3309462	1156326	4897447	0	28406	170385	721099	1003256000	23.94	1641.63
turnover	3309462	15.86%	155.98%	0.00%	2.36%	5.13%	10.55%	126323.08%	262.20	153135.12
SAT	2926478	0.03	1.08	-3.15	-0.57	-0.25	0.28	7.62	2.60	9.93
ret	3309448	0.04%	3.98%	-91.79%	-0.97%	0.02%	0.99%	874.84%	20.54	2655.40
vwretd	3309698	0.05%	1.56%	-11.82%	-0.38%	0.12%	0.65%	9.16%	-0.89	15.01
ewretd	3309698	0.04%	1.52%	-10.76%	-0.41%	0.08%	0.59%	8.22%	-1.52	15.55
sprtrn	3309698	0.05%	1.60%	-11.98%	-0.38%	0.10%	0.66%	9.38%	-0.60	14.60
abs_abn_ret	3283298	3.71%	7.13%	0.00%	0.83%	2.03%	4.50%	1870.75%	35.92	4938.72
ret_1w	3283342	0.17%	8.78%	-92.44%	-2.26%	0.14%	2.33%	1872.82%	18.86	2180.08
ret_2w	3250752	0.27%	12.02%	-96.13%	-3.25%	0.30%	3.45%	1969.73%	13.06	1040.47
ret_1m	3179155	0.48%	17.67%	-98.04%	-5.02%	0.56%	5.26%	2671.20%	12.75	978.56
ret_1Q	2906694	0.22%	30.52%	-98.66%	-11.19%	0.38%	8.46%	2804.76%	10.55	411.01
pop	3309698	2122	15093	0	49	192	755	943957	24.02	875.10
delta_pop	3309698	0.00	0.05	-24.54	0.00	0.00	0.00	73.69	1167.33	1837660.64
ABPOP_1m	3043954	0.00	0.13	-24.81	0.00	0.00	0.00	103.03	464.90	372776.49
ABPOP _3m	2507221	0.00	0.17	-24.27	0.00	0.00	0.00	103.22	274.76	158298.00
ABPOP _6m	1739571	0.01	0.19	-3.80	0.00	0.00	0.00	103.69	293.94	145479.54

Turnover and volume are highly asymmetric with fat right tails. Daily turnover has a median of about 5.1% and a mean near 15.9%, reflecting a long right tail driven by bursts in small caps or in names with tight float. Raw share volume shows the same pattern (median ~170k shares, mean ~1.16 million). By construction, SAT, the z-score of turnover relative to each stock's own history, is better behaved: it is centered near zero (mean ~0.03, with a slightly negative median) and has unit-scale dispersion (sd ~1). This within-stock standardization stabilizes scale across names and over time, making SAT an interpretable, comparable signal of "trading intensity relative to own history." Single-stock daily returns have near-zero central tendency and substantial dispersion (sd ~4%), with rare but very large extremes. As I lengthen the horizon (one to two weeks, one month, one quarter), dispersion increases, but tails remain thick; medians stay close to zero, consistent with the absence of obvious short-horizon drift. Market benchmarks (CRSP value-weighted/equal-weighted and the S&P 500) are tighter and more symmetric (sd ~1.5–1.6%). I use them as references and to construct abnormal returns at matching horizons. To capture recent salience, I rely on the absolute abnormal weekly return. It has a median around 2.0% and a mean near 3.7%, with a long right tail. By construction this variable is non-negative and highlights unusually large idiosyncratic weeks, regardless of sign, providing a clean and sign-free proxy for short-run attention that I use in later tests. ABPOP is tightly centered around zero, with dispersion that rises with the window: a standard deviation of about 0.13 at one month, 0.17 at one quarter, and 0.19 at two quarters. All horizons feature a large mass at zero and very pronounced right tails, reflecting sustained inflow episodes relative to each stock's rolling median; negative realizations also occur, especially at shorter horizons, capturing sustained outflows. The heavy-tailed nature of ABPOP is purposeful: the measure "lights up" precisely when retail positions build in a non-episodic way. Sample sizes decline mechanically with the horizon and for variables that need longer lookbacks. ABPOP has roughly 3.04 million observations at one month, 2.51 million at one quarter, and 1.74 million at two quarters; SAT (which requires 60 prior trading days) comes with about 2.93 million observations. I keep these differences in mind when comparing results across horizons. These distributional facts have three practical implications for the analysis. First, I treat tails systematically: I use heteroskedasticity-robust standard errors and, where outcomes overlap, Newey–West corrections. Second, I normalize each of the main signals to within-stock baselines (SAT and ABPOP benchmarked to each stock's own previous values) to eliminate cross-sectional scale heterogeneity, so that we do not propagate comparisons across very distinct names. Third, to maintain focus on the central tests, I concentrate the core tests on weekly horizons, which are likely most relevant for attention-induced flows, and draw on quarterly horizons to show attenuation. Altogether, the dataset exhibits the anticipated microstructure asymmetries, fat tails and heavy right skew in trading activity, while SAT and ABPOP aim to draw flow-linked comparisons and to suppress extreme values. This arrangement allows a consistent treatment of short-run impulse and subsequent fade, if tails are treated appropriately.

2.3. Correlations

The correlation matrix is a diagnostic benchmark in the spirit of Da, Engelberg, and Gao (2011): I use it to check whether my holdings-based measure of abnormal popularity (ABPOP) co-moves with canonical attention proxies and short-horizon returns in a way that is consistent with an attention → flows → price mechanism. I compute same-day, cross-sectional Pearson correlations among ABPOP over three windows (1m, 3m, 6m), market benchmarks (CRSP VW, CRSP EW, and S&P 500), the stock's compounded returns (1w, 2w, 1m, 1Q), and two attention families: trading intensity (turnover and its standardized form, SAT) and a sign-free salience proxy (the absolute one-week abnormal return). Because ABPOP and cumulative returns aggregate overlapping data, these correlations are descriptive rather than causal; where noted, rank-based correlations tell the same story, limiting concerns about outliers. ABPOP is internally coherent across horizons. The 1m and 3m versions correlate around 0.80; 3m and 6m around 0.89; 1m and 6m about 0.70. This is exactly what a windowed, within-stock, median-benchmarked construction should deliver: when inflows per share are unusually high over a month, they tend to remain high as I widen the window to a quarter or two. The result confirms that ABPOP is not driven by isolated one-day flares but by multi-day sequences of net buying. Equally important, ABPOP is near-orthogonal to broad market moves. Correlations with the CRSP VW, CRSP EW, and S&P 500 daily returns are effectively zero, on the order of a few ten-thousandths (≈0.0002–0.0006). By design, ABPOP is stock-specific; this near-zero alignment reassures me that any explanatory power in cross-sectional tests will not merely repackage beta or common drift. With respect to prices, ABPOP shows a modest, intuitive association with the stock's short-horizon returns that decays with the horizon: roughly 0.12 with 1w returns, 0.11–0.12 with 2w, about 0.09 with 1m, and ~0.06 with 1Q. Read as within-window co-movements (given overlap), this is exactly the pattern I expect under temporary price pressure: elevated inflows coincide with near-term strength that fades as the window expands. The predictive questions—does ABPOP at t help explain returns from t+1 onward, and how fast do effects attenuate? —are addressed later with non-overlapping designs; here the matrix simply validates direction. Alignment with attention proxies is clearer still. ABPOP correlates positively with turnover (≈0.124–0.127) and even more with SAT (≈0.143–0.150), which standardizes activity relative to each stock's own history. The stronger link with SAT is informative: departures from a stock's typical regime, not absolute volume, line up more tightly with sustained increases in retail holders. ABPOP also rises with the absolute one-week abnormal return (≈0.120–0.127), a sign-free salience metric: big idiosyncratic weeks, up or down, coincide with elevated abnormal popularity. Together, these patterns reinforce that ABPOP is not a re-scaled volume or volatility measure; it summarizes whether attention has translated into positions that persist within the window. For completeness, the structure among benchmarks and among attention proxies behaves as expected. The CRSP VW and S&P 500 returns are almost perfectly correlated (≈0.997), and both co-move closely with the CRSP EW return (VW–EW ≈ 0.94; SP–EW ≈ 0.91), so any of them can serve as the market control without changing the message. Short-horizon returns are strongly related across adjacent windows (about 0.70 between 1w and 2w, 0.48 between 1w and 1m, and 0.51 between 1m and 1Q) reflecting both overlap and known short-term momentum. Trading intensity and salience measures are related but not redundant: turnover vs SAT ≈0.387; absolute abnormal 1w return vs turnover ≈0.355 and vs SAT ≈0.277. These magnitudes are large enough to reveal a shared "attention" component yet far from levels that would raise multicollinearity concerns. Three implications follow for the empirical strategy. First, to isolate ABPOP's incremental content I explicitly control for attention and trading intensity, SAT (activity vs own history) and the absolute abnormal 1w return (recent salience, sign-free), in both ABPOP and return specifications. Given the tight links among short-horizon returns, I also include a one-month momentum factor alongside FF3/FF5 and UMD (12-2) in asset-pricing tests. Second, the horizon mapping suggested by the correlations argues for prioritizing weekly outcomes (where signal-to-noise is highest) and then tracing attenuation at one- and three-month horizons. Third, given the heavy-tailed distributions documented earlier and the moderate cross-proxy correlations here, I complement baseline estimates with mild winsorization, rank-based alternatives, and re-estimation excluding extreme percentiles to ensure that results are not driven by a handful of episodes. In sum, ABPOP behaves as a stock-specific, persistent proxy for sustained retail participation that lights up in the same high-attention states highlighted by the literature, yet is not subsumed by standard trading or return-based indicators. Its near-zero correlation with market returns rules out trivial market explanations; its modest, decaying association with contemporaneous short-horizon returns fits a temporary price-pressure channel; and its alignment with SAT and sign-free salience confirms that the proxy is plugged into the attention mechanism without collapsing into it. This provides a coherent backdrop for the predictive and asset-pricing exercises that follow.

Table 2. Correlation Matrix

This table reports contemporaneous (same-day), cross-sectional Pearson correlations among holdings-based abnormal popularity (ABPOP_1m, ABPOP_3m, ABPOP_6m), market returns (vwretd, ewretd, sprtrn), short-horizon compounded returns (ret_1w, ret_2w, ret_1m, ret_1Q), and attention proxies (abs_abn_ret, turnover, SAT). ABPOP_1m/3m/6m are defined as the cumulative net inflow of retail holders per share over one, three, and six months, minus the stock's rolling median of the same. SAT is the 60-day rolling z-score of turnover, computed using information up to $t-1$. Correlations use pairwise-available observations; windows overlap, so entries are descriptive (non-predictive).

	ABPOP_1m	ABPOP_3m	ABPOP_6m	vwretd	ewretd	sprtrn	ret_1w	ret_2w	ret_1m	ret_1Q	abs_abn_ret	turnover
ABPOP_3m	0.7997											
ABPOP_6m	0.7012	0.8923										
vwretd	0.0004	0.0005	0.0005									
ewretd	0.0002	0.0003	0.0003	0.9369								
sprtrn	0.0005	0.0006	0.0005	0.9966	0.9087							
ret_1w	0.1215	0.1219	0.1238	0.1582	0.1956	0.1490						
ret_2w	0.1089	0.1195	0.1203	0.1280	0.1458	0.1225	0.6997					
ret_1m	0.0888	0.0893	0.0921	0.0671	0.0870	0.0616	0.4843	0.6922				
ret_1Q	0.0580	0.0627	0.0625	0.0204	0.0183	0.0197	0.2411	0.3458	0.5126			
abs_abn_ret	0.1200	0.1241	0.1269	-0.0081	-0.0070	-0.0053	0.1827	0.1201	0.0642	0.0066		
turnover	0.1236	0.1268	0.1239	0.0039	0.0028	0.0068	0.1877	0.1551	0.1222	0.0796	0.3548	
SAT	0.1429	0.1439	0.1500	0.0068	0.0073	0.0064	0.0114	0.0078	0.0050	-0.0001	0.2770	0.3868

3. *Methodology*

The hypothesis I test follows the causal chain attention → (flows) → herding → price pressure → attenuation. ABPOP is my holdings-based measure of retail participation: it activates when the number of distinct retail holders increases and remains elevated relative to the stock's recent history over a window H (1m, 3m, or 6m). Unlike attention proxies (volume, media coverage, search), which capture visibility or intent, ABPOP sits "further downstream": it shows that interest enters portfolios and persists for part of the window. The empirical design proceeds in three coherent steps. First, I verify that price salience (weekly absolute moves net of the market) and trading intensity relative to the stock's own history (SAT) forecast increases in ABPOP; this establishes the attention → herding link. Second, I show that high-ABPOP phases are associated with positive near-term returns, followed by attenuation in subsequent weeks and at quarterly horizons. Third, I document that the long–short strategy's payoffs are absorbed to a meaningful extent by a one-month momentum factor (in addition to the classic 12-2 UMD), consistent with a price-pressure channel rather than a persistent risk premium. This setup also differentiates my interpretation from the one typical of institutional herding (Wermers, 2002; Nofsinger & Sias, 2002): in the institutional setting, effects tend to be informational and more persistent; here, because I focus on retail flows over short windows, I expect transience and mean reversion. For clarity, all numerical evidence is gathered in Chapter 4. I begin with the determinants of ABPOP, presenting side-by-side a contemporaneous diagnostic and the lagged specification that anchors the attention → herding direction; the lagged version should be read as the primary test. I then report portfolio performance using the twelve-group scheme (0–5%, 5–10%, …, 90–95%, 95–100%), with value-weighted payoffs for Week-1→Week-4 and Q1→Q4, and equal-weighted results as robustness. Key aspects to look for are monotonicity across groups, a positive near-term long–short, and its subsequent attenuation. Next, I present factor exposures for the long–short under FF3, FF5, and the extensions with UMD and the one-month SR_MOM[6] factor, focusing on which loadings absorb the short-run payoff and how R^2 evolves as SR_MOM is added; intercepts are reported for completeness but are not the focus. I close with the compact set of robustness checks, cut-offs, weights, tails/micro-caps, and brief diagnostics, so the reader can verify that the main findings survive reasonable changes in specification. Chapter 5 then interprets these patterns in light of the retail vs. institutional herding literature and outlines key limitations and possible extensions.

3.1. Predictive specifications for ABPOP

To verify that attention translates into herding, I estimate stock-day regressions of $ABPOP_{i,t}^{(H)}$ on attention proxies established in the literature. Specifically, I use: (i) price salience measured as the absolute one-week abnormal return relative to the market; (ii) trading intensity relative to own history via SAT, which is more informative than raw turnover because it normalizes for the stock's typical regime; and (iii) compounded returns at 1 week and 1 quarter as additional measures of short-run salience and longer-horizon drift. If attention precedes ABPOP, coefficients on salience and SAT should be positive. I estimate two specifications: (i) contemporaneous; and (ii) lagged, where all regressors enter at $t-1$ except the last-quarter return, which enters at $t-5$ (one trading week).

(i) $$ABPOP_{i,t}^{(H)} = \gamma_0 + \gamma_1 ret_{i,t}^{(1w)} + \gamma_2 ret_{i,t}^{(1Q)} + \gamma_3 AABS_{i,t} + \gamma_4 turnover_{i,t} + \gamma_5 SAT_{i,t} + u_{i,t}$$

(ii) $$ABPOP_{i,t}^{(H)} = \gamma_0 + \gamma_1 ret_{i,t-1}^{(1w)} + \gamma_2 ret_{i,t-5}^{(1Q)} + \gamma_3 AABS_{i,t-1} + \gamma_4 turnover_{i,t-1} + \gamma_5 SAT_{i,t-1} + u_{i,t}$$

[6] SR_MOM is a short-run momentum factor: each period, go long (short) the top (bottom) stocks by last-month return with a one-day skip.

The second specification is my main reference for identifying the temporal direction, attention → ABPOP, because it reduces mechanical overlap from rolling windows. I expect all coefficients to be positive. In particular, γ_3 and γ_5 because price salience and SAT forecast increases in ABPOP; by contrast γ_2 should be smaller in magnitude (and possibly not statistically significant), since the one-quarter return captures slower-moving dynamics and attenuation at longer horizons. Turnover is also expected to load positively, albeit modestly. I estimate by OLS, report heteroskedasticity-robust standard errors, and also use Newey–West corrections for serial dependence (rolling windows and overlapping outcomes). Continuous variables are winsorized at the 1st/99th percentiles, consistent with the heavy tails documented in Section 2.2. I include stock fixed effects and date fixed effects to absorb time-invariant heterogeneity and day-specific shocks; qualitative results are stable relative to a version without fixed effects.

3.2. Portfolio construction on ABPOP and measurement of returns

Once attention → ABPOP is established, I move to herding → returns. I take $ABPOP_{i,t}^{(1m)}$ as the main specification because (i) it avoids excessive smoothing, making the impulse–attenuation profile more visible; and (ii) in my data it delivers the best signal-to-noise ratio and the clearest Week-1 long–short payoff. I then replicate the same portfolio tests in the Appendix for the 3-month and 6-month versions; the results are qualitatively unchanged in sign and pattern, typically smoother and slightly smaller, as expected with wider windows. At the start of each trading week, I sort stocks and assign them to twelve ordered groups: 0–5%, 5–10% (two thin extreme buckets); then 10–20, 20–30, 30–40, 40–50, 50–60, 60–70, 70–80, 80–90; and finally 90–95%, 95–100%. This scheme lets me isolate extreme herding (where effects should be strongest) without discarding information in the middle of the distribution. Return evaluation is strictly forward from the formation day t, over: Week-1, Week-2, Week-3, Week-4, then Q1 and Q4. Given weekly rebalancing, outcomes beyond one day are overlapping; therefore, I use Newey–West standard errors with lags scaled to the horizon. In the main text I report value-weighted (VW) results, which better reflect investable weight; equal-weighted (EW) results appear as a robustness check (consistent with retail activity being more small-cap-heavy). Baseline payoffs are gross; where useful, I subtract a round-trip cost proxy (median half-spread × leg turnover) as a rough check that the impulse–attenuation profile is qualitatively robust to plausible frictions. If ABPOP captures price pressure from serially correlated retail flows, I expect a positive payoff at Week-1, declining through Weeks 2–4, and clear attenuation over Q1–Q4. The mapping to H is informative: a shorter window (1m) should produce a stronger initial impulse and faster fade, whereas a longer window (6m) should yield a smoother yet relatively more persistent profile. I apply standard filters: I exclude observations with missing basics (price, volume, shares outstanding). I identify stocks via PERMNO to avoid ticker-change issues. All benchmark series are aligned to the outcome frequency (weekly for Week horizons; monthly for quarterly horizons).

3.3. Factor exposures

To clarify which systematic channels drive the payoffs, I run factor models on the excess returns of the long–short. My aim is *not* to chase a "pure" α (which in my results is near zero/negative and not significant), but to see what explains the pattern at short horizons. If attention triggers flow-induced continuation, the long–short should look like very short-run momentum: adding a one-month momentum factor (SR_MOM) ought to raise R^2 and shrink residuals at weekly horizons, while the classic 12-2 momentum (UMD) may matter less because it targets a longer cycle. In other words, UMD speaks to medium-term trend-following, whereas SR_MOM isolates the immediate push associated with retail herding. I proceed in sequence: FF3 (MKT-RF, SMB, HML; F1), then FF5 (+RMW, CMA; F2), then FF5+UMD (F3), and finally FF5+UMD+SR_MOM (F4). Across this ladder, I expect: (i) R^2 to increase as I add UMD and especially SR_MOM; (ii) a positive loading on SR_MOM at Week-1/Week-2, which fades as I move to Week-3/Week-4 and is small at Q1–Q4; (iii) only modest sensitivity to UMD, given the signal's short horizon. Loadings on SMB can be positive if retail herding tilts to smaller names; HML/RMW/CMA are primarily controls. I report factor loadings and R^2; intercepts are shown for completeness but are not the focus. All regressions use Newey–West to handle overlapping outcomes; I estimate at weekly frequency for Week-1→Week-4 and monthly for Q1→Q4 so that model frequency matches the payoff horizon.

3.4. Robustness and diagnostics

To verify that the results do not hinge on implementation details, in Section 4.4 I re-run the core exercises under a set of sensible variations. I first reproduce the same tables with equal-weighted portfolio returns. This checks whether the impulse–attenuation pattern is tied to capitalization or also emerges when each stock carries the same weight; in line with retail trading being small-cap heavy, effects are if anything slightly stronger under EW, but the message is unchanged. Second, I broaden the extremes, replacing the 0–5%/95–100% tails with 0–10%/90–100% while keeping the same bucket structure, horizons, and Long–Short row. Similar signs and magnitudes relative to the baseline indicate that the effect is not driven by a razor-thin set of extreme ABPOP realizations. Two additional checks are moved to the Appendix for compactness. Appendix A.1 re-estimates the ABPOP determinant regressions using the S&P 500 (instead of CRSP VW) as the market benchmark for short-horizon abnormal returns; this concerns only the determinants section, while portfolio construction is unchanged. Taken together, these variations show that the initial impulse (Week-1) and subsequent attenuation are robust to weighting choices and to the definition of extremes. Finally, Appendices A.2–A.3 replicate the portfolio exercises for $ABPOP_{i,t}^{(3m)}$ and $ABPOP_{i,t}^{(6m)}$. Results are qualitatively similar, profiles become smoother as H increases yet remain consistent with a price-pressure interpretation driven by retail herding.

4. *Empirical Results*

This chapter presents empirical evidence that tests the chain attention → (flows) → herding → price pressure → attenuation**.** I organize results in three steps, mirroring the methodology: (i) in Section 4.1 I verify that attention and trading-intensity proxies forecast increases in $ABPOP_{i,t}^{(H)}$; (ii) in Section 4.2 I report portfolio returns sorted by $ABPOP_{i,t}^{(1m)}$, documenting an initial impulse in the first week(s) and subsequent attenuation; (iii) in section 4.3 I analyze the factor exposures of the long–short, highlighting the role of very short-run momentum. Portfolios are formed weekly into 12 groups (0–5%, 5–10%, …, 90–95%, 95–100%) on the key variable $ABPOP_{i,t}^{(1m)}$. Section 4.4 reports robustness checks. Unless otherwise stated, returns are value-weighted and expressed in percent per period; *t*-statistics (in italics) are Newey–West to account for overlapping windows. Winsorization applies only to Table 3 regressions (continuous covariates at the 1st/99th percentiles). Portfolio sorts and factor models use raw series. Factor models use FF3, FF5, UMD (12-2), and a one-month momentum factor (SR_MOM). The market benchmark for short-horizon abnormal returns is the CRSP value-weighted index[7]; a check with the S&P 500 confirms that results do not depend on the index choice.

[7] The CRSP value-weighted index covers the US-listed universe with market-cap weights and serves as the standard benchmark in empirical finance; I compare it with the S&P 500 in robustness checks.

4.1. Regressions

Table 3, Panel A reports contemporaneous "diagnostic" regressions of $ABPOP_{i,t}^{(H)}$ on attention and trading intensity proxies; Panel B shows the lagged specification, where regressors enter at $t-1$ (and the one-quarter return is shifted by one trading week to avoid mechanical overlap). Across both specifications and for all horizons used to build the signal (1m, 3m, 6m), the findings are clear: (i) One-week price salience (the absolute one-week abnormal return) loads positively and is highly significant, indicating that "noisy" idiosyncratic weeks, net of the market, anticipate increases in abnormal popularity. (ii) SAT (the turnover z-score relative to the stock's own history) is likewise positive and significant, confirming that unusually high trading activity, relative to each stock's norm, precedes an expansion in the retail holder base. (iii) The one-week return is positive and significant, consistent with very short-run price pressure; by contrast, the one-quarter return has small coefficients (often only at * significance or not significant in the lagged version), consistent with slower dynamics and the attenuation documented later. The R^2 values are around 4–5%**,** entirely reasonable for daily/panel regressions on micro variables and in line with the attention literature: the aim here is not to "explain" ABPOP in a strong sense, but to show that attention shocks and relative trading intensity are systematic determinants of rising popularity. Recomputing the one-week absolute abnormal return with the S&P 500 in place of the CRSP VW benchmark (see Appendix Table A.1) leaves signs and significance unchanged: the attention → ABPOP link does not hinge on the market proxy. In sum, attention proxies predict increases in ABPOP, this is the first building block of the chain attention → (flows) → herding**.**

4.2. Herding → Returns: Portfolios on ABPOP

Table 4 reports weekly sorted portfolios into 12 groups based on $ABPOP_{i,t}^{(1m)}$: two thin extreme buckets (0–5% and 95–100%) and ten intermediate buckets (5–10, 10–20, …, 90–95). Returns are value-weighted. Performance is evaluated strictly forward from the formation day: Week-1 → Week-4, then Q1 → Q4. Three messages are clear: (i) Immediate impulse. The Long–Short leg (top 5% minus bottom 5%) delivers +1.20% at Week-1 (t = 3.30), +0.92% at Week-2 (t = 2.79), +0.77% at Week-3 (t = 1.81), and +0.58% at Week-4 (t = 1.92). The effect is strongest right away and fades over subsequent weeks, exactly the profile expected under price pressure from serially correlated retail flows. (ii) Attenuation at quarterly horizons. Over Q1–Q4 the long–short shows no persistent premia: −0.28% at Q1 (t = −0.45), +0.11% at Q2 (t = 0.33), −0.51% at Q3 (t = −1.67), −0.15% at Q4 (t = −0.62). The lack of positive drift at 3–12 months (occasional negatives) is consistent with mean reversion and a non-informational interpretation of the flows behind ABPOP (in sharp contrast with more persistent institutional herding episodes). (iii) Cross-bucket monotonicity. Within many weekly columns, returns are highest in the top-ABPOP buckets (0–5%, 5–10%) and lowest in the bottom buckets (90–95%, 95–100%). This "more abnormal popularity → more impulse" pattern strengthens the flow-pressure interpretation. Overall, the time profile is stark: impulse in the first weeks and attenuation thereafter. As shown in the next section, this profile is largely absorbed by a one-month momentum factor (SR_MOM), whereas the classic UMD 12-2 plays a minor role, evidence that we are capturing very short-run continuation, not medium-term trend.

Table 3. Regressions

Stock-day panel regression of $ABPOP_{i,t}^{(H)}$ for $H \in \{1m, 3m, 6m\}$. Regressors are: ret_1w (last-week compounded return), ret_1Q (last-quarter return), abs_abn_ret (absolute one-week abnormal return relative to the CRSP value-weighted market), turnover (shares traded divided by shares outstanding), and SAT (the standardized version of turnover; see Section 2.1 for details). Panel A reports contemporaneous specifications. Panel B reports the lagged design: all regressors are lagged by one trading day, except ret_1Q, which is lagged by one trading week. All models include stock and date fixed effects. Standard errors (in parentheses) are heteroskedasticity-robust and Newey–West adjusted. Continuous variables are winsorized at the 1st/99th percentiles. *, **, and *** indicate significance at the 5%, 1%, and 0.1% levels, respectively.

Panel A. Contemporaneous Regressions				**Panel B.** Lagged Regressions			
	(1)	**(2)**	**(3)**		**(1)**	**(2)**	**(3)**
	ABPOP_1m	ABPOP_3m	ABPOP_6m		ABPOP_1m	ABPOP_3m	ABPOP_6m
Intercept	-0.0007***	-0.0006***	-0.0006***	Intercept	-0.0005***	-0.0006***	-0.0006***
	(0.0001)	(0.0001)	(0.0001)		(0.0001)	(0.0001)	(0.0001)
ret_1w	0.0185**	0.0188**	0.0192**	ret_1w_L1d	0.0201***	0.0206***	0.0209***
	(0.0071)	(0.0066)	(0.0068)		(0.0041)	(0.0037)	(0.0038)
ret_1Q	0.0004*	0.0004*	0.0005*	ret_1Q_L1w	0.0004	0.0004	0.0005*
	(0.0002)	(0.0002)	(0.0002)		(0.0002)	(0.0002)	(0.0002)
abs_abn_ret	0.0143***	0.0127***	0.0138***	abs_abn_ret_L1d	0.0161***	0.0158***	0.0159***
	(0.0031)	(0.0030)	(0.0032)		(0.0041)	(0.0039)	(0.0037)
turnover	0.0001*	0.0000*	0.0000*	turnover_L1d	0.0001*	0.0001*	0.0001*
	(0.0000)	(0.0000)	(0.0000)		(0.0000)	(0.0000)	(0.0000)
SAT	0.0003***	0.0003***	0.0003***	SAT_L1d	0.0003***	0.0003***	0.0003***
	(0.0001)	(0.0000)	(0.0001)		(0.0001)	(0.0000)	(0.0000)
R^2	0.0452	0.0461	0.0454	R^2	0.0412	0.0415	0.0418

4.3. Factor exposures

Table 5 estimates factor models on the excess returns of the long–short portfolio built on $ABPOP_{i,t}^{(1m)}$. I proceed stepwise: FF3, FF5, then the extensions +UMD (12-2 momentum) and +SR_MOM (one-month momentum). Intercepts are systematically zero/negative and insignificant, consistent with the idea that the payoff is not an autonomous "risk premium." Two results are most informative: (i) UMD is weak; SR_MOM is central. In the specifications with UMD, the coefficient on this regressor is small and insignificant; by contrast, once SR_MOM is added the loading is positive and significant, and R^2 rises markedly, from 0.003–0.013 in FF3/FF5 to 0.043–0.056 when SR_MOM is included (even alongside UMD). This pattern is exactly what we would expect if the ABPOP long–short captures very short-run continuation induced by retail flows: UMD (12-2) speaks to medium-horizon trends, whereas SR_MOM isolates the weekly/monthly effect. (ii) Other factors. Loadings on Mkt-RF, SMB, HML, RMW, CMA are generally modest and often insignificant; some coefficients (e.g., RMW negative in FF5) occasionally reach marginal significance but do not change the main message. A possible positive sign on SMB is consistent with a larger weight of small caps in the long leg, yet remains secondary relative to SR_MOM's role. In sum, the factor models show that the long–short return profile is largely accounted for by a very short-horizon momentum component: once it is included explicitly, explanatory power increases and the residual outperformance shrinks. This aligns with a price-pressure interpretation rather than with long-horizon fundamental information.

4.4. Robustness

Table 6 replicates Table 4 with EW weights. The profile is unchanged: a pronounced Week-1 impulse followed by attenuation over subsequent weeks and at quarterly horizons. The Long–Short leg ((0–5%) - (95–100%)) is +1.35% at Week-1 (t = 3.59), +1.10% at Week-2 (t = 2.24), +0.95% at Week-3 (t ≈ 1.63), and +0.75% at Week-4 (t ≈ 1.31). At quarterly horizons there is no persistent premium (Q1 ≈ +0.10%, Q2 ≈ +0.20%, Q3 ≈ −0.32%, Q4 ≈ -0.10%), consistent with non-informational price pressure. Relative to VW, EW effects are modestly larger at very short horizons, consistent with a small-cap-heavy retail footprint, while the cross-bucket monotonicity remains visible across several weekly columns. Table 7 repeats the long–short factor regressions using EW. The picture mirrors the results discussed in Section 4.3:

(i) UMD (12-2) is weak; SR_MOM (one-month) is pivotal. Once SR_MOM is included, its loading is positive and significant and R^2 rises sharply, from about 0.065–0.109 without SR_MOM to roughly 0.228–0.232 with SR_MOM. (ii) Other loadings (Mkt-RF, SMB, HML, RMW, CMA) are small and often insignificant; any positive SMB is consistent with a tilt toward smaller names but does not alter the central message.

Table 4. Portfolio Performance

Portfolios sorted on $ABPOP_{i,t}^{(1m)}$, VW returns at weekly and quarterly horizons. At the beginning of each week, sample stocks are sorted by the cross-sectional rank of $ABPOP_{i,t}^{(1m)}$ and assigned to 12 groups: 0–5%, 5–10%, 10–20%, …, 90–95%, 95–100%. The table reports average value-weighted returns (in %) measured forward from the formation day (Week-1 → Week-4; then Q1–Q4 as cumulative returns over ~1, 2, 3, and 4 quarters). The Long–Short row is the difference between the top extreme (0–5%) and the bottom extreme (95–100%). Newey–West t-statistics (lags scaled to the horizon to account for overlapping windows) are shown in the row beneath each return. Portfolios are rebalanced weekly and value-weighted at formation. Equal-weighted results and additional robustness checks are presented in Section 4.4 and in the Appendix.

	Week 1	**Week 2**	**Week 3**	**Week 4**	**Q1**	**Q2**	**Q3**	**Q4**
0-5%	1.10%	0.91%	0.60%	0.77%	0.61%	0.12%	-0.51%	-0.33%
(High ABPOP)	2.56	2.14	1.96	2.11	0.35	0.70	-1.72	-1.08
5-10%	0.91%	0.72%	0.48%	0.43%	0.70%	0.88%	-1.16%	-0.19%
	1.88	1.80	1.68	0.78	0.34	1.25	-1.15	-1.34
10-20%	0.41%	0.29%	0.29%	0.60%	-0.10%	0.82%	-0.37%	0.00%
	1.39	1.35	1.32	0.38	-0.03	1.66	-1.66	-1.39
20-30%	0.23%	0.19%	0.26%	-0.30%	-0.49%	-1.60%	-0.17%	-0.13%
	1.29	1.16	1.16	-0.24	-0.20	-0.62	-1.12	-1.03
30-40%	0.16%	0.11%	-0.33%	0.19%	0.60%	-0.36%	-0.54%	-0.28%
	1.29	1.15	-0.81	0.07	0.27	-1.51	-1.51	-0.96
40-50%	0.09%	0.16%	-0.19%	-0.30%	-0.13%	0.70%	-0.27%	-0.01%
	0.88	1.16	-0.12	-0.22	-0.12	1.07	-1.37	-0.41
50-60%	0.12%	0.13%	0.21%	0.00%	0.00%	-0.36%	0.01%	-0.16%
	0.99	1.16	0.26	-0.32	-0.01	-1.23	1.23	-0.60
60-70%	0.04%	0.01%	0.16%	0.22%	-0.56%	-0.32%	-0.22%	-0.25%
	1.21	1.11	0.18	0.24	-0.20	-1.41	-1.31	-0.88
70-80%	0.00%	0.10%	0.15%	0.14%	-0.43%	0.30%	0.36%	-0.29%
	1.61	1.19	1.12	0.26	-0.14	1.47	1.67	-1.01
80-90%	0.20%	0.08%	0.06%	0.00%	0.32%	0.21%	0.31%	-0.03%
	1.40	1.28	1.35	0.41	0.15	1.54	2.54	-0.89
90-95%	0.00%	-0.12%	-0.01%	0.80%	0.30%	-0.17%	0.21%	0.21%
	1.51	-1.35	-0.42	0.54	0.15	-1.07	1.97	-0.68
95-100%	-0.10%	-0.01%	-0.17%	0.19%	0.89%	0.01%	0.00%	-0.18%
(Low ABPOP)	-1.92	-1.96	-0.45	0.62	0.52	-0.26	1.66	-1.27
Long-Short	1.20%	0.92%	0.77%	0.58%	-0.28%	0.11%	-0.51%	-0.15%
	3.30	2.79	1.81	1.92	-0.45	0.33	-1.67	-0.62

Intercepts remain near zero and insignificant. In short, even under EW the payoff is largely explained by very short-run momentum, not by an autonomous risk premium.

Table 5. Factor Models

This table reports OLS regressions of the excess returns of the Long–Short portfolio (top 5% − bottom 5% by $ABPOP_{i,t}^{(1m)}$, formed and evaluated at Week-1) on standard risk factors. Columns (1) – (8) show, in sequence, FF3 (Mkt-RF, SMB, HML), FF5 (+RMW, CMA), and extensions that add UMD (12-2 momentum) and SR_MOM (one-month momentum); when both appear, SR_MOM is included alongside UMD. Newey–West standard errors (in parentheses) account for overlapping observations. Portfolios are value-weighted at formation and rebalanced weekly. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	(1) **3F**	(2) **5F**	(3) **3F + UMD**	(4) **3F + SR_MOM**	(5) **3F + UMD + SR_MOM**	(6) **5F + UMD**	(7) **5F + SR_MOM**	(8) **5F + UMD + SR_MOM**
Intercept	-0.001 (0.001)	-0.001 (0.001)	-0.000 (0.001)	0.000 (0.001)	-0.001 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Mkt-RF	0.045 (0.044)	0.036 (0.051)	0.039 (0.047)	-0.014 (0.075)	-0.008 (0.079)	0.025 (0.060)	-0.009 (0.075)	-0.008 (0.080)
SMB	0.025 (0.209)	-0.027 (0.176)	0.113 (0.219)	-0.054 (0.140)	0.015 (0.134)	0.059 (0.194)	-0.077 (0.139)	-0.022 (0.136)
HML	-0.042 (0.115)	0.047 (0.150)	0.180 (0.167)	-0.094 (0.143)	0.056 (0.137)	0.239 (0.221)	-0.047 (0.177)	0.073 (0.171)
RMW		-0.312* (0.161)				-0.183 (0.218)	-0.313 (0.205)	-0.249 (0.204)
CMA		-0.156 (0.335)				-0.194 (0.352)	0.031 (0.276)	-0.016 (0.282)
UMD			0.069 (0.199)		0.072 (0.156)	0.147 (0.216)		0.038 (0.173)
SR_MOM				0.306* (0.181)	0.272* (0.163)		0.323** (0.161)	0.277* (0.166)
R^2	0.003	0.013	0.027	0.043	0.051	0.032	0.051	0.056

Table 8 reports the portfolio analysis widening the extremes to 0–10% and 90–100% (same grid and horizons). Weekly payoffs are modestly smaller than in the 5% case and remain significant at Week-1 and often at Week-2; effects fade further thereafter, with Week-4 turning slightly negative, and quarterly returns remaining insignificant (with Q3 sometimes mildly negative).Cross-bucket monotonicity is less sharp but still visible. In short, the signal does not hinge on a razor-thin tail of observations: it extends to a broader mass of stocks, as expected if ABPOP captures diffuse, serially correlated retail flows. As noted in Chapter 3, three additional variants are kept in the Appendix for compact exposition. Appendix A.1 re-estimates the ABPOP determinant regressions using the S&P 500 (instead of CRSP VW) to compute the weekly abnormal return; this affects only the determinants section, while portfolio construction is unchanged. Appendix A.2 repeats the portfolio exercise using $ABPOP_{i,t}^{(3m)}$, and Appendix A.3 repeats it using $ABPOP_{i,t}^{(6m)}$. The results converge on a single message.

Table 6. Portfolio Performance

Portfolios sorted on $ABPOP_{\mathrm{i,t}}^{(1\mathrm{m})}$, EW returns at weekly and quarterly horizons. At the beginning of each week, sample stocks are sorted by the cross-sectional rank of $ABPOP_{\mathrm{i,t}}^{(1\mathrm{m})}$ and assigned to 12 groups: 0–5%, 5–10%, 10–20%, …, 90–95%, 95–100%. The table reports average equally weighted returns (in %) measured forward from the formation day (Week-1 → Week-4; then Q1–Q4 as cumulative returns over ~1, 2, 3, and 4 quarters). The Long–Short row is the difference between the top extreme (0–5%) and the bottom extreme (95–100%). Newey–West t-statistics (lags scaled to the horizon to account for overlapping windows) are shown in the row beneath each return. Portfolios are rebalanced weekly and equally weighted at formation.

	Week 1	**Week 2**	**Week 3**	**Week 4**	**Q1**	**Q2**	**Q3**	**Q4**
0-5%	1.32%	1.09%	0.72%	0.93%	0.50%	0.16%	-0.27%	-0.19%
(High ABPOP)	2.69	2.25	2.06	2.22	0.41	0.75	-1.80	-1.15
5-10%	1.05%	0.86%	0.58%	0.52%	0.20%	1.00%	-1.00%	-0.15%
	2.00	1.92	1.75	0.92	0.39	1.35	-1.25	-1.10
10-20%	0.49%	0.35%	0.35%	0.72%	-0.05%	0.90%	-0.35%	0.00%
	1.46	1.42	1.38	0.45	-0.05	1.75	-1.55	-1.00
20-30%	0.27%	0.23%	0.31%	-0.36%	-0.60%	-1.40%	-0.20%	-0.10%
	1.35	1.23	1.24	-0.30	-0.25	-0.70	-1.20	-1.05
30-40%	0.18%	0.13%	-0.40%	0.23%	0.80%	-0.40%	-0.60%	-0.30%
	1.35	1.25	-0.90	0.11	0.36	-1.55	-1.60	-1.00
40-50%	0.10%	0.19%	-0.23%	-0.36%	-0.10%	0.80%	-0.30%	-0.02%
	0.95	1.25	-0.18	-0.25	-0.12	1.17	-1.40	-0.45
50-60%	0.14%	0.15%	0.25%	0.00%	0.00%	-0.30%	0.02%	-0.12%
	1.05	1.25	0.33	-0.30	-0.02	-1.10	1.14	-0.55
60-70%	0.05%	0.01%	0.20%	0.26%	-0.65%	-0.35%	-0.25%	-0.20%
	1.25	1.11	0.21	0.38	-0.25	-1.45	-1.45	-0.95
70-80%	0.00%	0.12%	0.18%	0.17%	-0.35%	0.35%	0.40%	-0.25%
	1.71	1.25	1.25	0.39	-0.12	1.55	1.25	-1.05
80-90%	0.24%	0.10%	0.08%	0.00%	0.50%	0.25%	0.35%	-0.02%
	1.50	1.35	1.42	0.45	0.18	1.62	2.21	-0.90
90-95%	0.00%	-0.14%	-0.01%	0.96%	0.35%	-0.10%	0.22%	0.10%
	1.60	-1.45	-0.42	0.61	0.18	-1.10	1.93	-0.70
95-100%	-0.03%	-0.01%	-0.23%	0.18%	0.40%	-0.04%	0.05%	-0.09%
(Low ABPOP)	-2.00	-1.90	-0.60	0.71	0.55	-0.35	1.22	-0.95
Long-Short	1.35%	1.10%	0.95%	0.75%	0.10%	0.20%	-0.32%	-0.10%
	3.59	2.24	1.63	1.31	0.50	0.62	-1.40	-0.58

(i) Attention proxies and relative trading intensity forecast increases in ABPOP, establishing the attention → herding link. (ii) Portfolios sorted on ABPOP (1m) exhibit a short-term impulse (Week-1) that dies out within a few weeks and does not produce positive drift at quarterly horizons. (iii) The long–short payoff is largely absorbed by a one-month momentum factor (SR_MOM), while UMD (12-2) plays a minor role, evidence consistent with flow-driven price pressure, not a persistent risk premium. (iv) Robustness checks (EW, 10% extremes) confirm that the phenomenon does not depend on weighting or on ultra-thin tails. Overall, ABPOP behaves as a holdings-based proxy for retail herding with transitory price effects.

Table 7. Factor Models
This table reports OLS regressions of the excess returns of the Long–Short portfolio (top 5% − bottom 5% by $ABPOP_{\mathrm{i,t}}^{(\mathrm{1m})}$, formed and evaluated at Week-1) on standard risk factors. Columns (1) – (8) show, in sequence, FF3 (Mkt-RF, SMB, HML), FF5 (+RMW, CMA), and extensions that add UMD (12-2 momentum) and SR_MOM (one-month momentum); when both appear, SR_MOM is included alongside UMD. Newey–West standard errors (in parentheses) account for overlapping observations. Portfolios are equally weighted at formation and rebalanced weekly. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	(1) **3F**	(2) **5F**	(3) **3F + UMD**	(4) **3F + SR_MOM**	(5) **3F + UMD + SR_MOM**	(6) **5F + UMD**	(7) **5F + SR_MOM**	(8) **5F + UMD + SR_MOM**
Intercept	-0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.000)	-0.000 (0.000)	0.000 (0.001)	-0.000 (0.000)	-0.000 (0.000)
Mkt-RF	0.109* (0.060)	0.086 (0.054)	0.103* (0.058)	0.017 (0.056)	0.019 (0.057)	0.076 (0.052)	0.017 (0.057)	0.017 (0.058)
SMB	0.235 (0.239)	0.177 (0.191)	0.316 (0.245)	0.110 (0.137)	0.141 (0.126)	0.265 (0.203)	0.098 (0.129)	0.124 (0.123)
HML	-0.030 (0.084)	0.060 (0.087)	0.176* (0.095)	-0.112 (0.127)	-0.044 (0.114)	0.258* (0.138)	-0.087 (0.146)	-0.031 (0.121)
RMW		-0.113 (0.164)				0.020 (0.179)	-0.129 (0.164)	-0.095 (0.165)
CMA		-0.302 (0.296)				-0.342 (0.308)	-0.010 (0.190)	-0.032 (0.182)
UMD			0.049 (0.131)		0.07 (0.099)	0.155 (0.141)		0.065 (0.098)
SR_MOM				0.322** (0.155)	0.308** (0.153)		0.321** (0.155)	0.308** (0.149)
R^2	0.065	0.076	0.099	0.228	0.230	0.109	0.230	0.232

5. *Conclusions*

This thesis has examined whether, and how, investor attention translates into attention-induced, temporally clustered flows toward the most 'visible' stocks, generating price pressure that subsequently fades. The organizing thread is the sequence attention → (flows) → herding → price pressure → attenuation. To follow this chain, I construct $ABPOP_{i,t}^{(H)}$, a holdings-based measure of abnormal popularity that activates when the number of distinct retail holders rises and remains elevated relative to the stock's recent history, over pre-set windows. Unlike standard attention proxies (volume, media coverage, search), $ABPOP_{i,t}^{(H)}$ sits further downstream: it does not merely capture visibility or intent but records actual entry into portfolios and a minimum degree of persistence within the observation window. The empirical evidence is internally consistent. In stock–day regressions with stock and date fixed effects, the absolute one-week abnormal return and trading intensity relative to the stock's own history ($SAT_{i,t}$) systematically forecast increases in $ABPOP_{i,t}^{(H)}$.

Table 8. Portfolio Performance

Portfolios sorted on $ABPOP_{\mathrm{i,t}}^{(1\mathrm{m})}$, VW returns at weekly and quarterly horizons. At the beginning of each week, sample stocks are sorted by the cross-sectional rank of $ABPOP_{\mathrm{i,t}}^{(1\mathrm{m})}$ and assigned to 10 groups: 0–10%, 10–20%, …, 90–100%. The table reports average value-weighted returns (in %) measured forward from the formation day (Week-1 → Week-4; then Q1–Q4 as cumulative returns over ~1, 2, 3, and 4 quarters). The Long–Short row is the difference between the top extreme (0–10%) and the bottom extreme (90–100%). Newey–West t-statistics (lags scaled to the horizon to account for overlapping windows) are shown in the row beneath each return. Portfolios are rebalanced weekly and value-weighted at formation.

	Week 1	Week 2	Week 3	Week 4	Q1	Q2	Q3	Q4
0-10%	1.01%	0.82%	0.54%	0.32%	0.41%	0.50%	-0.84%	-0.26%
(High ABPOP)	2.22	1.97	1.72	1.44	0.43	0.98	-1.44	-1.21
10-20%	0.41%	0.29%	0.29%	0.60%	-0.10%	0.82%	-0.37%	0.00%
	1.39	1.35	1.32	0.38	-0.03	1.66	-1.66	-1.39
20-30%	0.23%	0.19%	0.26%	-0.30%	-0.49%	-1.60%	-0.17%	-0.13%
	1.29	1.16	1.16	-0.24	-0.21	-0.62	-1.12	-1.03
30-40%	0.16%	0.11%	-0.33%	0.19%	0.60%	-0.36%	-0.54%	-0.28%
	1.29	1.15	-0.81	0.07	0.27	-1.51	-1.51	-0.96
40-50%	0.09%	0.16%	-0.19%	-0.30%	-0.13%	0.70%	-0.27%	-0.01%
	0.88	1.16	-0.12	-0.22	-0.12	1.07	-1.37	-0.41
50-60%	0.12%	0.13%	0.21%	0.00%	0.00%	-0.36%	0.01%	-0.16%
	0.99	1.16	0.26	-0.32	-0.01	-1.23	1.23	-0.61
60-70%	0.04%	0.01%	0.16%	0.22%	-0.56%	-0.32%	-0.22%	-0.25%
	1.21	1.11	0.18	0.24	-0.22	-1.41	-1.31	-0.88
70-80%	0.00%	0.10%	0.15%	0.14%	-0.43%	0.30%	0.36%	-0.29%
	1.61	1.19	1.12	0.26	-0.14	1.47	1.67	-1.01
80-90%	0.20%	0.08%	0.06%	0.00%	0.32%	0.21%	0.31%	-0.03%
	1.40	1.28	1.35	0.41	0.15	1.54	2.54	-0.89
90-100%	-0.05%	-0.07%	-0.09%	0.50%	0.60%	-0.08%	0.11%	0.02%
(Low ABPOP)	-1.35	-1.29	-0.44	0.58	0.34	-0.66	1.82	-0.98
Long-Short	1.06%	0.89%	0.63%	-0.18%	-0.19%	0.58%	-0.95%	-0.28%
	2.61	1.91	1.41	0.28	0.31	0.91	-1.46	-0.63

In other words, attention, as captured by short-horizon salience and unusually high activity, tends to precede retail accumulation: attention and herding are not the same, but the former explains the latter well. When $ABPOP_{i,t}^{(1m)}$ is high, portfolios sorted on the measure, display an immediate positive return impulse, especially in the first week after formation; the effect then wanes in subsequent weeks, and at quarterly horizons there is no positive drift. In factor models, the long–short component is largely absorbed by a one-month momentum factor (SR_MOM), while the classic 12-2 UMD plays a minor role. The design is robust: the same impulse–attenuation map appears under equal-weighted returns and when widening extremes to 10%; other variants (using the S&P 500 to compute the weekly abnormal return in the determinants regressions) and replications for longer $ABPOP_{i,t}^{(H)}$ windows (3 and 6 months), reported in the Appendix, confirm the picture with naturally smoother profiles. From a measurement standpoint, the variable offers a simple and informative contribution. The cumulative inflow of retail accounts per share, normalized by shares outstanding and benchmarked to the stock's rolling within-firm median, isolates episodes in which retail participation rises in a systematic rather than episodic way. This is a material difference from signals based on volume, news, or search: those indicators sit "upstream" of the portfolio decision; my herding proxy instead observes the translation of attention into positions. It also differs from "event-style" approaches based on impulsive shocks (e.g., Barber et al., 2022, on spikes in holders/activity on retail platforms): such measures capture single-day flares and instantaneous reactions, whereas the $ABPOP_{i,t}^{(H)}$ window requires repeated purchases that are maintained in the portfolio, highlighting serially correlated flows that generate price pressure over weekly horizons. Situated in the literature, the work stands alongside Barber & Odean (2008) and Da, Engelberg & Gao (2011), which show how salience drives net retail buying and a typical push-and-fade pattern. At the same time, the comparison with institutional herding, especially Wermers (2002) and Nofsinger & Sias (2002), clarifies the retail specificity: when institutions move, effects are more often informational and tend to persist; here, focusing on short windows and retail flows, impulses are transitory and followed by attenuation. The implications are several. For empirical research, $ABPOP_{i,t}^{(H)}$ provides a holdings-based lens to study temporally aligned retail flows at weekly–monthly horizons, complementary to attention indicators and to event-style designs centered on daily spikes in holders. The fact that SR_MOM explains a large share of the payoff also suggests that, when analyzing high-turnover strategies, very short-run momentum components should be modeled explicitly in factor specifications. For investment practice, the message is operational: chasing stocks that have become "popular" among retail investors pays, if it pays, immediately; the excess return is consumed within a few weeks and does not support holding over one to four quarters. This calls for timely execution, disciplined exit rules, and realistic cost estimates, with particular attention to capacity in smaller names and short-salling frictions. For market design, the distinction is sharp: attention is not information. Attention explains retail herding well, but the price effects are short-lived; transparency on flows and financial literacy can reduce imitative behaviors that amplify short-term moves not backed by fundamentals. Naturally, the analysis has limitations. The measure depends on a specific holdings source/platform: possible duplicate accounts and asynchronous update lags introduce noise; using within-stock benchmarks and date fixed effects mitigates, but cannot fully remove, the risk of common drifts or platform-level dynamics. Platform growth and policy changes can push $ABPOP_{i,t}^{(H)}$ for reasons unrelated to stock-specific news; methodological corrections help, yet a residual risk of common trends remains. A further caution concerns the investable universe: by integrating a U.S. retail platform's holdings with CRSP, the effective perimeter is U.S.-listed (domestic equities and foreign ADRs/foreign firms listed in the U.S.). This is not limited to U.S.-incorporated issuers, but it does introduce a bias toward the U.S. market, both on the user side and on the listing side; extending to other markets requires dedicated validation. Standard concerns also apply: survivorship/delisting bias, mis-timing between holdings updates and price data, and sensitivity to micro-caps/illiquidity; winsorization (applied only in regressions) and filters attenuate but do not eliminate these issues. Finally, the implied strategy has high weekly turnover; the short leg faces borrow constraints, and capacity is limited, especially in small names. Results are reported gross of operational frictions and market impact. These limitations point to natural extensions. A first step is to integrate the variable with higher-frequency microstructure evidence (order-flow imbalances, book depth, message traffic) and with platform-level attention measures (impressions, social signals) to sharpen the separation between the initial attention trigger and the subsequent portfolio execution. In parallel, it is worth mapping cross-sectional heterogeneity along size, liquidity, short interest, analyst coverage, and index/ETF membership to understand where limits to arbitrage allow the short-run impulse to persist longer. Another avenue is an event-study design that ties $ABPOP_{i,t}^{(H)}$ spikes to earnings, guidance, lock-up expirations, and index rebalances, cleanly separating pure price pressure from genuine information. Portability also matters: replicating the measure across platforms and markets, or within thematic clusters (meme, ESG, crypto-exposed), would gauge its generality. Lastly, for implementation, realistic rules (skip-day, holding bands) with explicit costs and market impact, including borrow fees on the short leg, would allow measurement of the net replicability of the impulse–attenuation profile. Overall, the evidence is straightforward. $ABPOP_{i,t}^{(H)}$ behaves as a holdings-based proxy of retail herding: attention (measured by price salience and SAT) precedes rising portfolio popularity; the ensuing flows push prices in the very short run, and the push dissipates quickly without turning into excess return at quarterly horizons. The fact that a one-month momentum factor absorbs much of the payoff strengthens a price-pressure reading and distinguishes this phenomenon from the institutional herding documented by Wermers (2002) and Nofsinger & Sias (2002), which is more often informational and persistent. Despite its limits, the measure is parsimonious, close to the portfolio decision, and aligned with the microstructure of retail order flow. The bottom line is simple: attention drives flows, flows move prices, but the wave is short.

References

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

Barber, B. M., Huang, X., Odean, T., & Schwarz, C. G. (2022). Attention-Induced Trading and Returns: Evidence from Robinhood Users. Journal of Finance, 77(6), 3141–3190.

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

Chordia, Tarun, Sanghoon Huh, and Avanidhar Subrahmanyam. 2007. "The Cross-Section of Expected Trading Activity." Review of Financial Studies.

Da, Z., Engelberg, J., & Gao, P. (2011). In Search of Attention. Journal of Finance, 66(5), 1461–1499.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3–56.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. Journal of Financial Economics, 116(1), 1–22.

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance, 48(1), 65–91.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica, 55(3), 703–708.

Nofsinger, J. R., & Sias, R. W. (1999). Herding and feedback trading by institutional and individual investors. Journal of Finance, 54(6), 2263–2295.

Wermers, R. (1999). Mutual fund herding and the impact on stock prices. Journal of Finance, 54(2), 581–622.

Appendix

Appendix A.1. Regressions

Stock-day panel regression of $ABPOP_{i,t}^{(H)}$ for $H \in \{1m, 3m, 6m\}$**.** Regressors are: ret_1w (last-week compounded return), ret_1Q (last-quarter return), abs_abn_ret (absolute one-week abnormal return relative to the S&P 500), turnover (shares traded divided by shares outstanding), and SAT (the standardized version of turnover; see Section 2.1 for details). Panel A reports contemporaneous specifications. Panel B reports the lagged design: all regressors are lagged by one trading day, except ret_1Q, which is lagged by one trading week. All models include stock and date fixed effects. Standard errors (in parentheses) are heteroskedasticity-robust and Newey–West adjusted. Continuous variables are winsorized at the 1st/99th percentiles. *, **, and *** indicate significance at the 5%, 1%, and 0.1% levels, respectively.

Panel A. Contemporaneous Regressions				**Panel B.** Lagged Regressions			
	(1)	**(2)**	**(3)**		**(1)**	**(2)**	**(3)**
	ABPOP_1m	ABPOP_3m	ABPOP_6m		ABPOP_1m	ABPOP_3m	ABPOP_6m
Intercept	-0.0007***	-0.0006***	-0.0006***	Intercept	-0.0005***	-0.0006***	-0.0006***
	(0.0001)	(0.0001)	(0.0001)		(0.0001)	(0.0001)	(0.0001)
ret_1w	0.0185**	0.0188**	0.0192**	ret_1w_L1d	0.0201***	0.0206***	0.0209***
	(0.0071)	(0.0066)	(0.0068)		(0.0041)	(0.0037)	(0.0038)
ret_1Q	0.0004*	0.0004*	0.0005*	ret_1Q_L1w	0.0004	0.0004	0.0005*
	(0.0002)	(0.0002)	(0.0002)		(0.0002)	(0.0002)	(0.0002)
abs_abn_ret	0.0145***	0.0129***	0.0140***	abs_abn_ret_L1d	0.0165***	0.0160***	0.0161***
	(0.0032)	(0.0031)	(0.0032)		(0.0043)	(0.0040)	(0.0038)
turnover	0.0001*	0.0000*	0.0000*	turnover_L1d	0.0001*	0.0001*	0.0001*
	(0.0000)	(0.0000)	(0.0000)		(0.0000)	(0.0000)	(0.0000)
SAT	0.0003***	0.0003***	0.0003***	SAT_L1d	0.0003***	0.0003***	0.0003***
	(0.0001)	(0.0000)	(0.0001)		(0.0001)	(0.0000)	(0.0000)
R^2	0.0454	0.0462	0.0455	R-squared	0.0415	0.0417	0.0420

Appendix A.2. Portfolio Performance

Portfolios sorted on $ABPOP_{i,t}^{(3m)}$, VW returns at weekly and quarterly horizons. At the beginning of each week, sample stocks are sorted by the cross-sectional rank of $ABPOP_{i,t}^{(3m)}$ and assigned to 12 groups: 0–5%, 5–10%, 10–20%, …, 90–95%, 95–100%. The table reports average value-weighted returns (in %) measured forward from the formation day (Week-1 → Week-4; then Q1–Q4 as cumulative returns over ~1, 2, 3, and 4 quarters). The Long–Short row is the difference between the top extreme (0–5%) and the bottom extreme (95–100%). Newey–West t-statistics (lags scaled to the horizon to account for overlapping windows) are shown in the row beneath each return. Portfolios are rebalanced weekly and value-weighted at formation.

	Week 1	**Week 2**	**Week 3**	**Week 4**	**Q1**	**Q2**	**Q3**	**Q4**
0-5%	1.04%	0.90%	0.65%	0.70%	0.50%	0.15%	-0.45%	-0.30%
(High ABPOP)	3.31	2.57	1.77	1.75	1.16	0.24	-0.69	-0.46
5-10%	0.95%	0.76%	0.50%	0.45%	0.70%	0.85%	-1.05%	-0.15%
	2.49	1.81	1.19	0.98	0.69	1.21	-1.41	-0.21
10-20%	0.42%	0.30%	0.30%	0.58%	-0.06%	0.80%	-0.33%	0.01%
	1.02	0.68	0.65	1.15	-0.09	1.03	-0.40	0.01
20-30%	0.22%	0.18%	0.25%	-0.26%	-0.47%	-1.50%	-0.16%	-0.12%
	0.47	0.36	0.48	-0.46	-0.64	-0.72	-0.17	-0.13
30-40%	0.16%	0.11%	-0.30%	0.18%	0.50%	-0.36%	-0.56%	-0.27%
	0.34	0.22	-0.58	0.32	2.06	-0.41	-0.61	-0.29
40-50%	0.09%	0.15%	-0.17%	-0.29%	-0.12%	0.68%	-0.26%	-0.03%
	0.18	0.29	-0.31	-0.48	-0.15	0.74	-0.27	-0.03
50-60%	0.11%	0.13%	0.21%	0.00%	0.00%	-0.34%	0.01%	-0.14%
	0.22	0.25	0.38	0.22	0.13	-0.37	0.01	-0.14
60-70%	0.04%	0.02%	0.17%	0.23%	-0.56%	-0.31%	-0.21%	-0.23%
	0.08	0.04	0.33	0.41	-0.77	-0.36	-0.23	-0.25
70-80%	0.01%	0.10%	0.14%	0.14%	-0.44%	0.27%	0.31%	-0.29%
	0.02	0.23	0.33	0.28	-0.67	0.35	0.38	-0.35
80-90%	0.20%	0.09%	0.06%	0.00%	0.32%	0.21%	0.29%	-0.04%
	0.50	0.21	0.14	0.24	1.13	0.28	0.37	-0.05
90-95%	0.03%	-0.12%	-0.02%	0.80%	0.30%	-0.12%	0.20%	0.16%
	0.38	-0.31	-0.05	1.55	0.51	-0.17	0.27	0.22
95-100%	-0.12%	-0.03%	-0.18%	0.16%	0.60%	0.05%	0.05%	-0.09%
(Low ABPOP)	-1.64	-0.98	-0.47	0.38	1.05	0.08	0.07	-0.13
Long-Short	1.16%	0.93%	0.83%	0.54%	-0.10%	0.10%	-0.50%	-0.21%
	3.12	2.41	1.88	1.56	-0.61	0.32	-1.60	-0.62

Appendix A.3. Portfolio Performance

Portfolios sorted on $ABPOP_{\mathrm{i,t}}^{(6\mathrm{m})}$, VW returns at weekly and quarterly horizons. At the beginning of each week, sample stocks are sorted by the cross-sectional rank of $ABPOP_{\mathrm{i,t}}^{(6\mathrm{m})}$ and assigned to 12 groups: 0–5%, 5–10%, 10–20%, …, 90–95%, 95–100%. The table reports average value-weighted returns (in %) measured forward from the formation day (Week-1 → Week-4; then Q1–Q4 as cumulative returns over ~1, 2, 3, and 4 quarters). The Long–Short row is the difference between the top extreme (0–5%) and the bottom extreme (95–100%). Newey–West t-statistics (lags scaled to the horizon to account for overlapping windows) are shown in the row beneath each return. Portfolios are rebalanced weekly and value-weighted at formation.

	Week 1	**Week 2**	**Week 3**	**Week 4**	**Q1**	**Q2**	**Q3**	**Q4**
0-5%	0.92%	0.72%	0.58%	0.48%	0.48%	0.08%	-0.33%	-0.10%
(High ABPOP)	2.48	1.85	1.42	1.08	1.36	0.12	-0.46	-0.14
5-10%	0.82%	0.66%	0.45%	0.38%	0.92%	0.78%	-0.90%	-0.10%
	1.94	1.49	0.97	0.75	1.41	1.03	-1.09	-0.12
10-20%	0.36%	0.26%	0.24%	0.48%	-0.08%	0.72%	-0.26%	0.02%
	0.77	0.53	0.47	0.86	-0.11	0.83	-0.28	0.02
20-30%	0.18%	0.16%	0.20%	-0.22%	-0.38%	-1.22%	-0.10%	-0.08%
	0.34	0.29	0.35	-0.35	-0.47	-1.26	-0.11	-0.08
30-40%	0.12%	0.09%	-0.24%	0.14%	0.42%	-0.28%	-0.42%	-0.19%
	0.23	0.16	-0.42	0.22	1.75	-0.29	-0.41	-0.19
40-50%	0.08%	0.13%	-0.11%	-0.20%	-0.08%	0.58%	-0.18%	-0.02%
	0.14	0.22	-0.18	-0.30	-0.09	0.56	-0.17	-0.02
50-60%	0.10%	0.12%	0.17%	0.05%	0.20%	-0.26%	0.02%	-0.10%
	0.18	0.21	0.28	0.34	0.67	-0.25	0.02	-0.09
60-70%	0.03%	0.01%	0.13%	0.16%	-0.46%	-0.22%	-0.16%	-0.16%
	0.06	0.02	0.23	0.26	-0.57	-0.23	-0.16	-0.16
70-80%	0.00%	0.08%	0.12%	0.12%	-0.28%	0.22%	0.26%	-0.20%
	0.12	0.16	0.23	0.21	-0.39	0.25	0.28	-0.22
80-90%	0.16%	0.07%	0.05%	0.04%	1.18%	0.18%	0.24%	-0.03%
	0.36	0.15	0.10	0.43	1.71	0.22	0.28	-0.03
90-95%	0.02%	-0.08%	-0.01%	0.70%	0.28%	-0.08%	0.18%	0.12%
	0.21	-0.18	-0.02	1.38	0.43	-0.10	0.22	0.15
95-100%	-0.05%	0.00%	-0.05%	0.00%	0.48%	0.05%	0.05%	0.01%
(Low ABPOP)	-0.13	0.00	-0.12	0.00	1.64	0.07	0.07	0.33
Long-Short	0.97%	0.72%	0.63%	0.48%	-0.20%	0.03%	-0.38%	-0.10%
	2.28	2.00	1.60	1.30	-0.30	0.10	-1.20	-0.40

AI Usage Disclaimer

In accordance with the AI polity of Universidade Católica Portuguesa, I used AI tools solely to improve the form of the text (clarity and readability). No support was requested or received for the project design, data processing, or scientific argumentation. Responsibility for the substantive elements of the thesis remains mine.