The Review of Financial Studies

Memory Moves Markets

Constantin Charles
London School of Economics, United Kingdom

I show that memory-induced attention can distort prices in financial markets. I exploit rigid earnings announcement schedules to identify which firms are associated in investors' memory. Firms with randomly overlapping earnings announcements are associated in memory because many investors experience them in the same context. Months later, when only one of the two firms announces earnings, this context is cued, and triggers the recall of the other, associated firm. On such days, I find that memory-induced attention leads to buying pressure in the associated firm's stock. The strength of this effect varies as predicted by associative memory theory. (*JEL* G14, G41)

Received: August 10, 2022; Editorial decision: June 2, 2024
Editor: Lauren Cohen

A large literature in finance analyzes the role of attention in financial markets (for an overview, see Barber, Lin, and Odean 2019). This literature has uncovered various sources of investor attention, including media coverage, abnormal trading volume, extreme stock returns, and the display of information (e.g., Barber and Odean 2008; Hartzmark 2015).[1] A unifying theme of these sources is that they are *external* sources of attention. By contrast, there might also be *internal* sources of attention – like memory associations inside an investor's mind – that systematically direct attention.

I especially thank my advisor Cary Frydman for his invaluable advice and guidance. I am grateful to Lauren Cohen (Editor), two anonymous referees, Kenneth Ahern, Colin Camerer, Tom Chang, Joey Engelberg, Sam Hartzmark, David Hirshleifer, Xing Huang, Chad Kendall, Mete Kilic, Spencer Kwon, Chris Parsons, Vesa Pursiainen, Daniel Schmidt, Kelly Shue, Johnny Tang, Felix von Meyerinck, Nancy Xu, and Tiange Ye, as well as participants at the Miami Behavioral Finance Conference, the American Finance Association Meeting, the SFS Cavalcade, Nanyang Technological University, the University of Miami, HEC Paris, Vanderbilt University, the London School of Economics, the SAFE Institute, the Frankfurt School of Finance & Management, Oxford University, Emory University, Boston College, Dartmouth College, and the University of Washington for helpful comments and suggestions. Parts of this research were conducted while I was visiting the University of St. Gallen.
Send correspondence to Constantin Charles, c.charles@lse.ac.uk.

[1] Further external sources of attention identified in the literature include ownership of an asset (Hartzmark, Hirshman, and Imas 2021), earnings announcements (Hartzmark and Shue 2018; Hirshleifer, Lim, and Teoh 2009; Schmidt 2019), extraordinary events (Chen et al. 2019; Seasholes and Wu 2007), advertisement (Lou 2014), information display (Barber et al. 2022; Frydman and Wang 2020; Chen et al. 2023), social media (Jiang et al. 2022), and days of the week (DellaVigna and Pollet 2009).

The Review of Financial Studies 00 (2024) 1–46

https://doi.org/10.1093/rfs/hhae086 Advance Access publication December 23, 2024

So far, however, there is little evidence of internally generated attention in financial markets. One possible reason for this dearth of evidence is that researchers have lacked formal models of, and empirical proxies for, these internal sources. Fortunately, recent advances in memory theory provide the necessary structure to analyze memory recall as a source of internally generated attention (Bordalo, Gennaioli, and Shleifer 2020; Wachter and Kahana 2024). Memory models formalize which items will be associated in an investor's memory and therefore allow for targeted tests of memory-induced attention in financial markets. Put differently, memory models allow me to test for a different type of investor attention than has previously been investigated.

In this paper, I test whether an event that increases attention to one firm also channels attention to another firm if the two firms are associated in investors' memory. My tests build on the strong existing evidence showing that individual investors are net buyers of attention-grabbing stocks, to the point where they can create positive price pressure (Barber and Odean 2008; Da, Engelberg, and Gao 2011). Motivated by this evidence, I test and confirm the hypothesis that memory-induced attention creates buying pressure, and show that it leads to positive abnormal returns for memory-associated firms.

The two key empirical challenges in testing this hypothesis are (1) estimating which firms are associated in investors' memory and (2) identifying memory associations that are orthogonal to firm fundamentals. In addressing these challenges, I am guided by the model of Bordalo, Gennaioli, and Shleifer (2020), which builds heavily on associative memory theory (Kahana 2012). In associative memory theory, recall is shaped by two competing forces: similarity and interference. To see how these forces operate, assume that firm A announces earnings on day t. When an investor is cued by this event, she recalls past experiences that are similar to the cue. For instance, the investor might recall firm B, because firm B announced earnings on the same day as firm A in the previous quarter, and the two firms were both covered in the news on that day. In the terminology of the model, these two firms are encoded as more *similar* in memory, because they were experienced in a similar context by the investor in the previous quarter. In contrast, if many other firms announced earnings on the same day as firms A and B in the previous quarter, the investor may recall one of these other firms instead of firm B. Thus, the memories of these other firms *interfere* with the recall of firm B on day t.

My hypothesis is simple. When firm A announces earnings on day t, this event naturally attracts attention to firm A itself. But since firms A and B are associated in memory, some attention might also be directed to firm B, even though firm B does not announce earnings on day t. Thus, I hypothesize that firm A's earnings announcement on day t (the cue) creates memory-induced attention to firm B, leading to buying pressure in firm B's stock.[2]

[2] In a related study, Charles (2022) uses microdata to show that such a cue increases the probability that an individual investor trades the stock of firm B.

In my empirical tests, I follow the logic of this example and estimate which firms are associated in investors' memory using plausibly random overlaps of earnings announcements in the previous quarter. To capture memory associations that are orthogonal to firm fundamentals, I use only earnings announcements that are exogenously shifted by calendar rotations (Noh, So, and Verdi 2021). These are announcements of firms that follow a strict pattern in their announcement timing for several years in a row and do not deviate from this pattern by even a single day.

Examples of patterns that firms follow are to always announce on the first Thursday of a month, or to always announce on the fifth Thursday after the fiscal period end.[3] For such "Pattern firms," the day-of-week on which a calendar month begins determines the date of their earnings announcement. Crucially, the day-of-week on which a calendar month begins rotates from year to year, shifting the dates of earnings announcements, and creating exogenous overlaps of earnings announcements for Pattern firms. By exploiting these rigid schedules, my approach directly addresses the concern that some firms strategically advance or delay their earnings announcement depending on the earnings they plan to report (Penman 1987; Bagnoli, Kross, and Watts 2002; Johnson and So 2018).

In my main specification, I regress the return of firm B over the window $[t,t+1]$ on a dummy variable for a memory cue on day t. This dummy is equal to one if at least one firm that announced earnings on the same day as firm B in the previous quarter (e.g., firm A) announces earnings on day t. In all my tests, I control for an own earnings announcement of firm B in the window $[t-10, t+10]$, to avoid confounding effects from a firm's own earnings announcement. I also control for other firm news events of firm B, including filings of Form 8-Ks, a whole host of firm events sourced from the Capital IQ Key Developments database, volume and sentiment of news coverage from RavenPack, and firm-year-quarter fixed effects. Finally, in my tests I only consider firm-pairs that operate in very different industries, using the text-based industries of Hoberg and Phillips (2010, 2016), firm-pairs that are not vertically related based on the measure of Frésard, Hoberg, and Phillips (2020), and firm-pairs that do not share any analyst following.

I find that a memory cue leads to an abnormal return of about 4 basis points (bps) for the cued firm over the window $[t,t+1]$. When I analyze the dynamics of this effect, I find that it materializes over the course of about 2 weeks, after which it fully reverses over the course of another 3 weeks. The effect size is about 8 bps over the window $[t,t+10]$, which is roughly on par with the earnings announcement premium (Frazzini and Lamont 2007; Barber et al. 2013). Risk-based return movements are unlikely to explain these results since I use characteristic-adjusted returns as in Daniel et al. (1997). It is also possible

[3] Such patterns are sometimes explicitly required by the firms' bylaws (Noh, So, and Verdi 2021).

to construct a trading strategy that takes advantage of predictable memory cues. The profitability of this trading strategy has increased over time and it generates an alpha of about 70 bps over the window $[t, t+10]$ from 2015 onward.

My hypothesis that memory-induced attention creates buying pressure also makes predictions about the underlying trading behavior. Therefore, I next examine whether memory cues lead to net buying in cued firms' stock. Since previous work shows that it is retail investors who buy attention-grabbing stocks, I focus on the trading behavior of retail investors, and calculate retail order imbalance using the algorithms proposed by Barber et al. (2024) and Boehmer et al. (2021). Consistent with my hypothesis, I find that memory cues lead to net retail buying. The dynamics of this behavior match the dynamics of the return effect.

As a deeper test of the mechanism, I next examine how the psychological properties of memory modulate the strength of the documented effect. My tests can be organized into three categories. First, I test whether firms that received more attention during the encoding are associated more strongly in investors' memory. Using three different proxies for investor attention suggested by Barber and Odean (2008) – high abnormal trading volume, high media coverage, and extreme returns – I find that attention during the encoding increases the strength of the underlying association. Firms with high abnormal trading volume during the encoding share especially strong memory associations.

Second, I test whether similarity, one of the two key forces of associative memory theory, modulates the strength of the documented effect. Similarity predicts that two experiences are more strongly associated if their features are more similar. Applying this intuition to my setting, similarity predicts that two firms are more associated if they had more similar earnings surprises during the encoding. My results suggest that higher similarity indeed increases the strength of the underlying association. Perhaps most strikingly, I find that similarity also determines the nature of the association. Firms that had extremely negative earnings surprises during the encoding share a negative association. In these cases, a memory cue leads to a negative return response on the day of the cue. However, these negative return responses only occur if the underlying memory association is encoded based on particularly negative tail events.

Third, I test whether interference weakens the documented effect. As discussed in the introductory example, if many firms announced earnings together with firms A and B, this should weaken the memory association between these two firms. Inherent in this line of reasoning is the assumption that there is strong attentional interaction between firms that have earnings announcements on the same day. In support of this key premise, Hirshleifer, Lim, and Teoh (2009) find that attention to one firm reduces attention to another firm on earnings announcement days. In my tests, I find suggestive evidence that memory associations that were encoded on days with many earnings announcements lead to weaker effects than associations that were encoded on days with few earnings announcements. While not conclusive, these findings are consistent with a signature prediction of associative memory theory.

In further tests, I explore whether firm events other than earnings announcements are effective at cueing the underlying associations. Associative memory theory is helpful at organizing these tests. It predicts that when two firms announce earnings on the same day, the shared context of this joint announcement generates the memory association. Thus, cues that bring to mind the context "earnings announcement" are effective at targeting the association. To test this prediction, I analyze the cueing effect of several types of firm events that differ in their similarity to the encoding context. I find that alternative firm events can act as memory cues, but only if the underlying memory association is particularly strong. Overall, earnings announcements are by far the strongest cues, highlighting the importance of contextual similarity between the cue and the underlying association during recall.

My paper contributes to the large literature in finance that studies the role of limited attention in financial markets (e.g., Barber and Odean 2008; Hirshleifer, Lim, and Teoh 2009; Da, Engelberg, and Gao 2011). Memory theory offers one explanation for why investors allocate their attention to certain firms: when cued with an event, investors retrieve associated firms from memory, and subsequently allocate more attention to these firms. This memory-induced attention can be strong enough to distort the stock prices of these firms. A related strand of the literature documents that recurring firm events are associated with predictably high returns (Hartzmark and Solomon 2018). In contrast to most studies in this literature, I do not focus on firms' own returns following a recurring event. Rather, I show that recurring firm events, such as earnings announcements, can serve as cues that trigger the recall of associated memories. Through these memory associations, events at the cueing firm can affect the returns of associated firms.

My results also relate to the literature on categorization (Barberis and Shleifer 2003; Barberis, Shleifer, and Wurgler 2005; Peng and Xiong 2006; Huang 2019). This literature has typically focused on categories derived from the characteristics of firms. Associative memory theory naturally generates such categorization, since items that share more similar features are associated more strongly in memory. But in associative memory theory, items can also become associated based on the context in which they were experienced. Consistent with this prediction, I find that experiencing two firms in the joint context of an earnings announcement leads to meaningful associations in investors' memory.

Overall, my results provide a strong empirical justification for incorporating aspects of human memory into economic models, an approach taken by a growing theoretical literature (Gilboa and Schmeidler 1995; Mullainathan 2002; Hirshleifer and Welch 2002; Bordalo, Gennaioli, and Shleifer 2020; Bodoh-Creed 2020; Nagel and Xu 2022; Bordalo et al. 2023; Wachter and Kahana 2024; Azeredo da Silveira, Sung, and Woodford 2024). My paper differs from previous empirical tests of memory models, as these tests largely focus on individual beliefs and decision-making (e.g., Goetzmann, Watanabe, and Watanabe 2022; Enke, Schwerter, and Zimmermann 2024; Colonnelli, Gormsen, and McQuade 2024; Jiang et al. 2024; Gödker, Jiao, and Smeets Forthcoming). In a related study, Charles (2022) shows that memory associations affect the trading behavior of individual investors. In contrast, the current study shows that memory effects can be powerful enough to affect asset prices.

1. Empirical Strategy

My goal is to identify exogenous associations of firms in investors' memory. In the ideal experiment, I would randomly associate firms in investors' memory, for example, by randomly exposing investors to different firms on different days. The resulting joint experience of two firms would create an association of those firms in investors' memory. I aim to approximate this ideal experiment using plausibly random overlaps of firms' quarterly earnings announcements. I use earnings announcements as building blocks for estimating memory associations, since these announcements naturally draw attention to announcing firms. In addition, there is evidence of attentional interaction between firms that announce on the same day (Hirshleifer, Lim, and Teoh 2009).

I consider two firms as associated in memory if they announced earnings on the same day in the previous fiscal quarter. This approach has the benefit of being simple while capturing the main forces of associative memory theory (Kahana 2012; Bordalo, Gennaioli, and Shleifer 2020). For instance, by comparing whether the memory association is stronger when both firms had similar earnings surprises on the day of the overlap, I can test for the effect of similarity in recall. Further, by comparing associations that were encoded on days with many announcing firms vs. on days with few announcing firms, I can test for the effect of interference in recall (Kahana 2012; Bordalo, Gennaioli, and Shleifer 2020; Bordalo et al. 2023).

In my baseline tests, I estimate regressions of the following type:

$$return_{i,[t,t+h]} = \beta \cdot cue_{i,t} + \gamma \cdot X_{i,t} + \delta_{i,q} + u_{i,t} \tag{1}$$

where $return_{i,[t,t+h]}$ is firm i's characteristic-adjusted cumulative return over the window $[t, t+h]$, $cue_{i,t}$ is a dummy variable that is equal to one if at least one firm that announced earnings on the same day as firm i in the previous fiscal quarter announces earnings on day t, $X_{i,t}$ is a vector of control variables, and $\delta_{i,q}$ is a firm-year-quarter fixed effect.

The coefficient β captures the effect of a memory cue for firm i on day t. My hypothesis is that such a cue creates memory-induced attention to firm i. Previous work suggests that attention, especially from retail investors, tends to lead to net buying, and ultimately to buying pressure (Barber and Odean 2008, Da, Engelberg, and Gao 2011). Therefore, I hypothesize that, unconditionally, β is positive. However, if the underlying memory association is particularly negative, it is possible that β is negative. In my tests, I will distinguish positive from negative associations to test whether the sentiment of the underlying association modulates the direction of the return effect. I cluster standard errors in all regressions by firm and day.

A natural worry with this approach is that firms with overlapping earnings announcements might be fundamentally more related than firms without overlapping earnings announcements. In this case, $cue_{i,t}$ could pick up fundamental relationships. Indeed, firms do announce earnings in clusters and their schedules are known to be correlated across industries and other firm characteristics. For instance, firms in the same industry tend to announce close in time to each other in a quarter. For this reason, in all my tests I exclude all firm-pairs that are in the same TNIC-2 industry (Hoberg and Phillips 2010, 2016). To rule out vertical relationships between firms, I exclude all firm-pairs that are vertically related based on the the 10% granularity version of the VTNIC database (Frésard, Hoberg, and Phillips 2020). Finally, to capture relationships that are not picked up by either of these classifications, I also exclude all firm-pairs with an overlap in analyst following.

While helpful, these restrictions may not fully capture cross-industry relationships, and they do not address concerns regarding endogenous timing of earnings announcements. There is a large literature showing that firms strategically advance or delay their earnings announcements depending on the news they plan to report (Penman 1987; Bagnoli, Kross, and Watts 2002; Johnson and So 2018). Firms that announce early in the quarter generally announce good news, while firms that announce late in the quarter generally announce bad news. As a result, the set of firms that announce early in a quarter is systematically different from the set of firms that announce late in a quarter.

To directly address these concerns, I exploit the fact that many firms follow rigid rules for their earnings announcement timing (Noh, So, and Verdi 2021). Two typical rules that firms follow are to always announce on the kth day-of-week of a calendar month, or to always announce on the kth day-of-week since the fiscal quarter-end.[4] Religiously following such rules results in overlaps for some earnings announcements, but not for others, depending only on how the calendar shakes out in each month.

To give a concrete illustration, consider two firms, A and B, whose fiscal quarters end on June 30th.[5] Further, assume that both firms are "Pattern firms"; that is, they both follow a strict quarter-specific pattern in their earnings announcement timing and have not deviated from this pattern by even one day

[4] Firms follow many other rules too. Noh, So, and Verdi (2021) provide a detailed list in appendix A of their paper.

[5] This example is adapted from Noh, So, and Verdi (2021).

June, July, August 2013						
Sun	Mon	Tue	Wed	Thu	Fri	Sat
30 Fiscal quarter-end (June 30th)	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31	1 A and B announce	2	3
4	5	6	7	8	9	10

June, July, August 2014						
Sun	Mon	Tue	Wed	Thu	Fri	Sat
29	30 Fiscal quarter-end (June 30th)	1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31 B announces	1	2
3	4	5	6	7 A announces	8	9

Figure 1
Calendar rotations: An example.

for at least 3 years. Specifically, firm A always announces on the first Thursday in August and firm B always announces on the fifth Thursday since the end of the fiscal quarter. As shown in Figure 1, the month of August began on a Thursday in 2013, but on a Friday in 2014. As a result, both firms announced on August 1st in 2013. However, in 2014, firm A announced on August 7th and firm B announced on July 31st. This is how calendar rotations – changes in the day-of-week on which a calendar month begins – create plausibly random overlaps of earnings announcements for Pattern firms. Unless explicitly noted otherwise, in my tests I only use earnings announcements that are exogenously shifted by calendar rotations to estimate memory associations and to proxy for memory cues.

2. Data and Summary Statistics

2.1 Sample

My sample spans years 2005 to 2019 and consists of all firm-days for which I can calculate characteristic-adjusted returns over the two return windows that I use in most of my tests, $[t, t+1]$ and $[t, t+10]$. My sample is restricted to years 2005 - 2019, since this is the period for which I have data on Pattern firms, which is critical for my empirical design. Throughout the paper, the term "days" always refers to trading days and the term "returns" always refers to characteristic-adjusted returns. I calculate characteristic-adjusted returns as in Daniel et al. (1997) and Hartzmark and Shue (2018) using data from CRSP and Compustat. Specifically, I triple-sort stocks into quintiles of size, book-to-market, and momentum, and then match each individual stock to one of the resulting 125 portfolios. If I cannot match a stock to one of the portfolios because of missing data in one of the sorting variables, I drop it. The characteristic-adjusted return on day t is the stock's raw return on day t minus the value-weighted return of the portfolio on day t. I use portfolio stocks' market capitalization from day $t-3$ as the weights in this calculation.

To capture the aggregate trading behavior of retail investors, I identify and sign retail trades using the algorithm proposed by Barber et al. (2024). This algorithm classifies transactions in the Trade and Quote (TAQ) data with exchange code "D" as a retail sell if the execution price is less than the quote midpoint, and as a retail buy if the execution price is greater than the quote midpoint. Trades that execute between 40% and 60% of the National Best Bid or Offer are not signed.[6] I also identify and sign retail trades using the algorithm proposed by Boehmer et al. (2021). This algorithm classifies transactions in the TAQ data as a retail sell if the subpenny price improvement is in the interval (0, 0.4) and as a retail buy if it is in the interval (0.6, 1).[7]

For both methodologies, I calculate retail order imbalance for firm i on day t as the difference between retail buying and selling volume scaled by the sum of retail buying and selling volume.[8] For the Boehmer et al. (2021) algorithm, I also construct retail order imbalance using only odd lot trades, which are trades of fewer than 100 shares. Data on retail order imbalance is available for years 2007 - 2019 and for years 2013 - 2019 when using odd lots.

My empirical strategy relies on identifying overlaps of earnings announcements. It is therefore crucial that the earnings announcement dates are measured without error. Since I control for a firm's own earnings announcement in all of my tests, it is also important that I do not miss any earnings

---

6 The code to implement this algorithm can be found here. I execute the code exactly as-is, using the default filters, for years 2007 - 2019.

7 The code to implement this algorithm can be found here. I again execute the code exactly as-is, using the default filters, for years 2007 - 2019.

8 In a robustness test in Appendix Table C.1, I also use retail order imbalance calculated using the number of retail buy and sell trades. The results are very similar for this alternative measure.

announcements. To achieve both of these goals, I follow Noh, So, and Verdi (2021) and use data from I/B/E/S and Compustat. If earnings announcement dates are missing in one database, I use dates from the other one, to capture as many earnings announcements as possible. While the earnings announcement dates are mostly identical across the two databases, there are cases in which they differ. In these cases, I follow DellaVigna and Pollet (2009) and use the earlier of the two dates, because the later date often reflects the publication date in the *Wall Street Journal* rather than the date of the earnings announcement itself. Finally, using time stamps from I/B/E/S, I follow Patton and Verardo (2012) and treat earnings announcements that occur at or after 4pm (when the market closes) as effectively occurring on the following day.

I also calculate the earnings surprise associated with each earnings announcement. I identify each analyst's most recent forecast in I/B/E/S, and take the median of all analyst forecasts made between 2 and 45 days prior to the earnings announcement. Then, I calculate the surprise as the difference between the actual earnings announced by the firm and the median earnings forecast, scaled by the share price of the firm from three days prior to the announcement.

A key variable of interest in my analysis is the occurrence of a memory cue, which I capture with a simple dummy variable. A cueing event for firm i on day t occurs if at least one firm that announced earnings on the same day as firm i in the previous fiscal quarter announces earnings on day t. In my tests, I ensure that both the estimation of memory associations as well as the occurrence and timing of memory cues is exclusively driven by exogenous calendar rotations by focusing on earnings announcements of Pattern firms. I identify earnings announcements of Pattern firms using data provided by Noh, So, and Verdi (2021).[9] I focus on the *threshold3* data set provided by the authors, which classifies a firm as a Pattern firm if it followed a strict quarter-specific pattern in its earnings announcement timing for 3 or more years.

In my tests, I require the cueing firms to be in a different industry than firm i to avoid picking up within-industry information spillovers. I use the text-based network industry classifications of Hoberg and Phillips (2010, 2016) to ensure that the firms operate in dissimilar industries in the year of the cue.[10] I also require that cueing firms are not vertically related to firm i in the year of the cue based on the the 10% granularity version of the VTNIC database (Frésard, Hoberg, and Phillips 2020), and I ensure that the cueing firms and firm i do not share any overlap in analyst following in the days $t-45$ to t. Further, I require cueing firms to have a market capitalization (measured on

[9] I thank the authors for providing these data, which are available on Suzie Noh's website.

[10] This classification is more flexible than standard classifications (e.g., SIC or NAICS), as it changes over time and allows each firm to have a unique set of competitors. I use the broader TNIC-2 industries provided by the authors that have been calibrated to be as granular as two-digit SIC codes. Specifically, with this classification, 4.5% of randomly drawn firms are deemed to be peers.

day $t-3$) that is above the NYSE's 90th percentile of market capitalization in that month. I do so to focus on large and salient cues that attract the attention of many investors.[11] Later in the paper, I will analyze whether and under which circumstances smaller firms can act as cues.

I also collect data on other firm events. In particular, I source filing dates of Form 8-Ks from the WRDS SEC Analytics Suite, and collect dates of other firm news events from the Capital IQ Key Developments database. This database monitors over 100 different firm news events, including executive changes, M&A rumors, and changes in corporate guidance. In some tests, I follow Kwon and Tang (2023) and focus on the set of Key Developments that occur frequently.[12]

To capture detailed firm-level news coverage, I use data from RavenPack News Analytics Dow Jones Edition. I use the Aggregate Event Sentiment (AES) variable to gauge the overall tone of a firm's news coverage. This variable is a score that indicates the percentage of all non-neutral news events with a positive tone. AES is constructed on a daily level, using a 91-day rolling window, and considers all news sources from the Dow Jones universe (including Dow Jones Newswires, the *Wall Street Journal*, Barron's, and MarketWatch). I also use the Aggregate Event Volume variable to capture the amount of news coverage a firm received. This variable is also constructed daily, and is simply the count of all news articles that are used to construct the AES variable. In my sample, I winsorize all nonlogarithmized variables (except for dummy variables) at the 1st and the 99th percentiles.

2.2 Summary statistics

Table 1 presents summary statistics for my sample. I present cumulative returns for the two return windows that I use in most of my tests: $[t, t+1]$ and $[t, t+10]$. The average and median returns are slightly negative, but the sizable standard deviation for both windows indicates that there is large variation in returns.

I also present summary statistics for retail order imbalance, both for the Barber et al. (2024) (henceforth, BHJOS) measure and the Boehmer et al. (2021) (henceforth, BJZZ) measure. For the BJZZ measure, I show retail order imbalance constructed using all lots and using only odd lots. To match the return windows, I calculate the average retail order imbalance over $[t, t+1]$ and $[t, t+10]$.[13] Retail order imbalance is slightly negative on average, indicating slightly more selling than buying from retail investors, consistent with the summary statistics in Boehmer et al. (2021). Here, too, the considerable standard deviation indicates that there is much variation in retail order imbalance in the cross-section and over time.

---

11 My approach is motivated by Hartzmark and Shue (2018), who use the same cutoff to identify salient earnings surprises in their test of contrast effects in financial markets.

12 See Appendix Table D.1 for a list of these events.

13 In their main tests, Boehmer et al. (2021) also use firm-level average retail order imbalances over multiple days.

Table 1
Summary statistics

	Mean	p25	p50	p75	Std. Dev.	Min	Max	N
Return in % over $[t,t+1]$	−0.0310	−1.6360	−0.0844	1.4695	3.6471	−12.1098	13.7098	16,502,158
Return in % over $[t,t+10]$	−0.1121	−4.1097	−0.2875	3.5323	8.4877	−26.1222	31.6951	16,502,158
BHJOS retail OI in % over $[t,t+1]$	−4.4383	−19.6704	−2.4291	11.8836	30.9262	−100.0000	84.8624	9,159,341
BHJOS retail OI in % over $[t,t+10]$	−3.9370	−11.8761	−2.7535	4.9529	14.5156	−47.8068	32.3404	8,430,871
BJZZ retail OI in % over $[t,t+1]$	−2.6942	−18.4402	−1.9388	13.1443	26.4661	−71.7819	65.7579	8,038,145
BJZZ retail OI in % over $[t,t+10]$	−2.5887	−10.1575	−1.8349	5.4210	12.8316	−38.4445	29.1627	6,975,409
BJZZ retail OI in % (odd lots) over $[t,t+1]$	−0.3178	−15.8640	0.2656	15.5135	27.2288	−73.4037	70.9338	3,243,285
BJZZ retail OI in % (odd lots) over $[t,t+10]$	−0.1007	−8.2238	0.2125	8.2535	13.6882	−37.4915	34.7360	2,701,475
Cue (dummy)	0.0064	0.0000	0.0000	0.0000	0.0797	0.0000	1.0000	16,502,158
Number of cues	1.8932	1	1	2	1.7852	1	13	105,550
Surprise of cue (EW)	0.0005	0.0000	0.0003	0.0009	0.0014	−0.0040	0.0069	105,550
Surprise of cue (VW)	0.0005	0.0000	0.0002	0.0010	0.0014	−0.0042	0.0068	105,550

This table contains summary statistics of the sample used in the empirical analysis. The sample covers years 2005−2019. Return over the time window $[t,t+h]$ is the raw return of a stock over the time window $[t,t+h]$ minus the value-weighted return of a portfolio of stocks matched on size, book-to-market, and momentum. Retail order imbalance (OI) is calculated daily from Trade and Quote (TAQ) data using either the Barber et al. (2024) (BHJOS) or the Boehmer et al. (2021) (BJZZ) algorithm. Retail order imbalance over the time window $[t,t+h]$ is the average order imbalance over that time window. Data on retail order imbalance are available for years 2007 - 2019, and for years 2013 - 2019 when constructed using only odd lot trades (trades of fewer than 100 shares). Cue is a dummy variable equal to one if at least one firm that announced earnings on the same day as firm i in the previous fiscal quarter announces earnings on day t. Number of cues is the number of cueing firms announcing on firm-days with Cue equal to one. Surprise is the earnings surprise of the cueing firm(s) announcing on firm-days with Cue equal to one. On firm-days with multiple cues, I calculate both the equally weighted (EW) and value-weighted (VW) earnings surprise.

In my sample, there is a cueing event on 0.64% of firm-days, and of the 9,215 distinct firms in my sample, 5,390 firms (58.49%) have at least one cueing event. This statistic indicates that almost 60% of the firms in my sample are classified as a Pattern firm at some point in the sample period. While there is only one cue on the median firm-day with a cueing event, there are also firm-days with multiple cues. On such firm-days, I calculate the earnings surprise of the cue as either the equally weighted or value-weighted average surprise of the cueing firms, using each cueing firm's market capitalization three days prior to the announcement as value weights. Earnings surprises are typically close to zero.

3. Main Results

3.1 Return results

In my first test, I regress the return of firm i over the window $[t, t+1]$ on a dummy variable that is equal to one if there is a cueing event for firm i on day t. This dummy variable is equal to one if at least one firm that announced earnings on the same day as firm i in the previous fiscal quarter announces earnings on day t. The first column in panel A of Table 2 shows that the coefficient on this dummy variable is positive and highly significant. In terms of magnitude, the estimate implies that such a cue leads to an abnormal return of 6.2 bps over the return window $[t, t+1]$. However, this regression does not include any controls. Thus, one concern is that the cue dummy is simply correlated with firm i's own earnings announcement, and might therefore be picking up the well-documented earnings announcement premium (Frazzini and Lamont 2007, Barber et al. 2013). Therefore, I augment this regression with 21 dummy variables that capture a potential own earnings announcement of firm i on any day in the window $[t-10, t+10]$. The second column shows that including these control dummies does not affect the coefficient on the cue dummy, ruling out that the cue dummy is merely proxying for the earnings announcement premium.

In the third column, I include additional controls for other firm news events. In particular, I include a dummy that is equal to one if firm i filed a Form 8-K on day t, a dummy that is equal to one if there was a Key Development (other than an earnings announcement) for firm i on day t, two dummies that are equal to one if the change in aggregate news sentiment about firm i on day t is positive or negative, respectively, and the logarithm of 1 plus the absolute increase or decrease in aggregate news volume of firm i on day t. Including these controls does not affect the size or significance of the coefficient on the cue dummy. In the fourth column, I augment this regression with day fixed effects. These fixed effects account for the possibility that my results might be driven by days on which big and famous firms announce (Chen, Cohen, and Wang 2022). I find that the magnitude of the coefficient drops somewhat when I control for day

Table 2
Return results
A. Main sample

Dependent variable:	Return $[t,t+1]$ in %						Return $[t,t+10]$ in %					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Cue $[t]$	0.062***	0.065***	0.062***	0.043**	0.038**	0.018	0.171***	0.180***	0.178***	0.146***	0.084**	0.046
	(0.016)	(0.016)	(0.016)	(0.017)	(0.016)	(0.016)	(0.040)	(0.040)	(0.040)	(0.040)	(0.034)	(0.034)
Own EA $[t-10]$		0.009	0.010	0.009	0.009	0.006		0.013	0.014	0.079***	0.002	0.062***
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.024)	(0.024)	(0.024)	(0.022)	(0.021)
Own EA $[t-9]$		0.014	0.013	0.010	0.012	0.008		0.023	0.022	0.090***	0.009	0.072***
		(0.011)	(0.011)	(0.011)	(0.011)	(0.011)		(0.024)	(0.024)	(0.025)	(0.022)	(0.022)
Own EA $[t-8]$		0.016	0.016	0.016	0.015	0.014		0.032	0.032	0.097***	0.020	0.079***
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.024)	(0.024)	(0.025)	(0.022)	(0.023)
Own EA $[t-7]$		0.021**	0.021**	0.025**	0.019*	0.022**		0.035	0.035	0.101***	0.022	0.081***
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.024)	(0.024)	(0.026)	(0.022)	(0.024)
Own EA $[t-6]$		0.012	0.012	0.014	0.011	0.011		0.043*	0.044*	0.103***	0.030	0.081***
		(0.010)	(0.010)	(0.011)	(0.010)	(0.011)		(0.025)	(0.025)	(0.027)	(0.023)	(0.025)
Own EA $[t-5]$		0.023**	0.022**	0.017*	0.021**	0.015		0.054**	0.052**	0.111***	0.042*	0.092***
		(0.011)	(0.011)	(0.010)	(0.011)	(0.011)		(0.025)	(0.025)	(0.028)	(0.024)	(0.026)
Own EA $[t-4]$		0.037***	0.035***	0.031***	0.034***	0.027***		0.084***	0.082***	0.140***	0.069***	0.117***
		(0.010)	(0.010)	(0.010)	(0.011)	(0.010)		(0.024)	(0.024)	(0.027)	(0.024)	(0.026)
Own EA $[t-3]$		0.015	0.013	0.011	0.012	0.007		0.079***	0.078***	0.131***	0.063**	0.106***
		(0.011)	(0.011)	(0.011)	(0.011)	(0.011)		(0.025)	(0.025)	(0.028)	(0.025)	(0.026)
Own EA $[t-2]$		−0.057***	−0.059***	−0.057***	−0.061***	−0.061***		0.023	0.021	0.070**	0.005	0.045
		(0.012)	(0.012)	(0.012)	(0.012)	(0.012)		(0.026)	(0.026)	(0.029)	(0.026)	(0.028)
Own EA $[t-1]$		−0.159***	−0.163***	−0.164***	−0.165***	−0.169***		−0.094***	−0.099***	−0.062*	−0.112***	−0.086***
		(0.015)	(0.015)	(0.015)	(0.014)	(0.015)		(0.029)	(0.029)	(0.032)	(0.028)	(0.031)
Own EA $[t]$		−0.048**	−0.124***	−0.132***	−0.130***	−0.141***		−0.087***	−0.179***	−0.149***	−0.179***	−0.161***
		(0.019)	(0.020)	(0.020)	(0.019)	(0.019)		(0.033)	(0.035)	(0.038)	(0.033)	(0.035)
Own EA $[t+1]$		0.146***	0.093***	0.087***	0.088***	0.079***		0.015	−0.050	−0.021	−0.057*	−0.041
		(0.017)	(0.017)	(0.017)	(0.016)	(0.017)		(0.032)	(0.033)	(0.036)	(0.031)	(0.033)
Own EA $[t+2]$		0.148***	0.143***	0.141***	0.141***	0.136***		0.051	0.043	0.070**	0.028	0.044
		(0.012)	(0.012)	(0.013)	(0.012)	(0.013)		(0.032)	(0.032)	(0.035)	(0.030)	(0.032)
Own EA $[t+3]$		0.073***	0.069***	0.065***	0.067***	0.060***		0.063**	0.057*	0.078**	0.042	0.051
		(0.011)	(0.011)	(0.012)	(0.011)	(0.012)		(0.032)	(0.032)	(0.035)	(0.029)	(0.032)

(Continued)

Table 2
Continued
A. Main sample

Dependent variable:	Return $[t, t+1]$ in %						Return $[t, t+10]$ in %					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Own EA $[t+4]$		0.045***	0.042***	0.031***	0.039***	0.026**		0.065**	0.060*	0.066*	0.046	0.040
		(0.011)	(0.011)	(0.012)	(0.011)	(0.012)		(0.032)	(0.032)	(0.035)	(0.030)	(0.032)
Own EA $[t+5]$		0.041***	0.035***	0.024**	0.033***	0.020*		0.078**	0.068**	0.069**	0.057*	0.046
		(0.010)	(0.010)	(0.012)	(0.010)	(0.011)		(0.033)	(0.033)	(0.035)	(0.030)	(0.032)
Own EA $[t+6]$		0.043***	0.040***	0.035***	0.037***	0.029***		0.080**	0.074**	0.072**	0.058**	0.045
		(0.010)	(0.010)	(0.011)	(0.010)	(0.011)		(0.033)	(0.033)	(0.034)	(0.030)	(0.032)
Own EA $[t+7]$		0.029***	0.028***	0.025**	0.025**	0.020*		0.065**	0.063*	0.057*	0.045	0.028
		(0.010)	(0.010)	(0.011)	(0.010)	(0.011)		(0.032)	(0.032)	(0.034)	(0.029)	(0.031)
Own EA $[t+8]$		0.010	0.009	0.004	0.006	−0.001		0.064**	0.062*	0.051	0.045	0.022
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.032)	(0.032)	(0.033)	(0.029)	(0.030)
Own EA $[t+9]$		−0.011	−0.013	−0.024**	−0.016*	−0.029***		0.102***	0.099***	0.075**	0.082***	0.046
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.032)	(0.032)	(0.032)	(0.029)	(0.029)
Own EA $[t+10]$		−0.008	−0.013	−0.024**	−0.016*	−0.029***		0.221***	0.214***	0.186***	0.198***	0.159***
		(0.010)	(0.010)	(0.010)	(0.010)	(0.010)		(0.031)	(0.030)	(0.030)	(0.028)	(0.028)
Own 8-K $[t]$			−0.037***	−0.037***	−0.026***	−0.026***			−0.124***	−0.125***	−0.042***	−0.044***
			(0.007)	(0.007)	(0.007)	(0.007)			(0.016)	(0.016)	(0.012)	(0.012)
Own key development $[t]$			0.093***	0.094***	0.092***	0.093***			0.131***	0.138***	0.089***	0.096***
			(0.005)	(0.005)	(0.004)	(0.004)			(0.011)	(0.011)	(0.009)	(0.009)
Positive news coverage $[t]$			0.240***	0.241***	0.235***	0.235***			0.298***	0.306***	0.302***	0.311***
			(0.009)	(0.009)	(0.007)	(0.008)			(0.018)	(0.018)	(0.013)	(0.013)
Negative news coverage $[t]$			−0.434***	−0.433***	−0.432***	−0.432***			−0.519***	−0.517***	−0.417***	−0.412***
			(0.011)	(0.011)	(0.009)	(0.010)			(0.020)	(0.020)	(0.014)	(0.015)
log(1 + Increase in news volume)			0.127***	0.128***	0.130***	0.132***			0.188***	0.188***	0.120***	0.121***
			(0.007)	(0.007)	(0.007)	(0.007)			(0.014)	(0.014)	(0.011)	(0.011)
log(1 + Decrease in news volume)			0.135***	0.136***	0.134***	0.136***			0.182***	0.181***	0.095***	0.093***
			(0.006)	(0.006)	(0.006)	(0.006)			(0.013)	(0.012)	(0.011)	(0.010)
Day FE	no	no	no	yes	no	yes	no	no	no	yes	no	yes
Firm x Year-quarter FE	no	no	no	no	yes	yes	no	no	no	no	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.000	.000	.001	.005	.029	.033	.000	.000	.000	.004	.156	.158

Table 2
Continued
B. Using all EAs

Dependent variable:	Return $[t, t+1]$ in %						Return $[t, t+10]$ in %					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Cue $[t]$	0.050***	0.028***	0.022***	0.016***	0.007	−0.002	0.089***	0.050***	0.040**	0.055***	−0.018	−0.011
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.018)	(0.017)	(0.017)	(0.016)	(0.015)	(0.014)
Own EA $[t-10]$		0.021***	0.021***	0.023***	0.019**	0.018**		0.095***	0.095***	0.149***	0.064***	0.108***
		(0.008)	(0.008)	(0.008)	(0.008)	(0.008)		(0.020)	(0.020)	(0.019)	(0.018)	(0.017)
Own EA $[t-9]$		0.012	0.011	0.015*	0.009	0.010		0.095***	0.094***	0.154***	0.063***	0.111***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.021)	(0.021)	(0.020)	(0.019)	(0.018)
Own EA $[t-8]$		0.018**	0.018**	0.021***	0.015*	0.016**		0.094***	0.093***	0.152***	0.063***	0.109***
		(0.008)	(0.008)	(0.008)	(0.008)	(0.008)		(0.021)	(0.021)	(0.020)	(0.019)	(0.019)
Own EA $[t-7]$		0.024***	0.024***	0.027***	0.021***	0.022***		0.085***	0.085***	0.143***	0.054***	0.099***
		(0.008)	(0.008)	(0.008)	(0.008)	(0.008)		(0.021)	(0.021)	(0.021)	(0.019)	(0.019)
Own EA $[t-6]$		0.007	0.007	0.009	0.005	0.003		0.078***	0.078***	0.132***	0.049**	0.086***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.021)	(0.021)	(0.021)	(0.019)	(0.020)
Own EA $[t-5]$		0.006	0.005	0.007	0.004	0.003		0.081***	0.080***	0.135***	0.056***	0.093***
		(0.009)	(0.009)	(0.008)	(0.009)	(0.008)		(0.021)	(0.021)	(0.022)	(0.020)	(0.020)
Own EA $[t-4]$		0.001	−0.000	0.006	−0.002	0.001		0.077***	0.076***	0.133***	0.050**	0.088***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.021)	(0.021)	(0.022)	(0.020)	(0.020)
Own EA $[t-3]$		−0.024***	−0.025***	−0.017*	−0.026***	−0.022**		0.053**	0.053**	0.105***	0.026	0.059***
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)		(0.021)	(0.021)	(0.022)	(0.020)	(0.021)
Own EA $[t-2]$		−0.083***	−0.083***	−0.074***	−0.084***	−0.079***		−0.016	−0.016	0.033	−0.042**	−0.013
		(0.010)	(0.010)	(0.010)	(0.010)	(0.009)		(0.022)	(0.022)	(0.023)	(0.021)	(0.022)
Own EA $[t-1]$		−0.162***	−0.162***	−0.158***	−0.163***	−0.163***		−0.124***	−0.126***	−0.086***	−0.146***	−0.130***
		(0.012)	(0.012)	(0.012)	(0.011)	(0.011)		(0.024)	(0.024)	(0.026)	(0.023)	(0.024)
Own EA $[t]$		−0.002	−0.048***	−0.047***	−0.050***	−0.054***		−0.060**	−0.130***	−0.084***	−0.117***	−0.110***
		(0.015)	(0.015)	(0.015)	(0.015)	(0.015)		(0.028)	(0.029)	(0.030)	(0.026)	(0.027)
Own EA $[t+1]$		0.237***	0.206***	0.212***	0.204***	0.205***		0.082***	0.036	0.079***	0.031	0.043
		(0.013)	(0.013)	(0.014)	(0.013)	(0.013)		(0.027)	(0.028)	(0.029)	(0.025)	(0.026)
Own EA $[t+2]$		0.238***	0.235***	0.242***	0.233***	0.236***		0.160***	0.155***	0.189***	0.131***	0.144***
		(0.010)	(0.010)	(0.011)	(0.010)	(0.010)		(0.027)	(0.027)	(0.029)	(0.025)	(0.026)
Own EA $[t+3]$		0.144***	0.142***	0.141***	0.139***	0.135***		0.203***	0.200***	0.226***	0.172***	0.177***
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)		(0.027)	(0.027)	(0.028)	(0.024)	(0.026)
Own EA $[t+4]$		0.093***	0.091***	0.081***	0.088***	0.075***		0.223***	0.220***	0.233***	0.194***	0.185***
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)		(0.027)	(0.027)	(0.028)	(0.025)	(0.026)

(Continued)

Table 2
Continued
B. Using all EAs

Dependent variable:	Return $[t,t+1]$ in %						Return $[t,t+10]$ in %					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Own EA $[t+5]$		0.062***	0.057***	0.049***	0.055***	0.043***		0.244***	0.237***	0.246***	0.214***	0.200***
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)		(0.027)	(0.027)	(0.028)	(0.025)	(0.025)
Own EA $[t+6]$		0.054***	0.051***	0.049***	0.048***	0.041***		0.261***	0.258***	0.260***	0.228***	0.209***
		(0.009)	(0.009)	(0.009)	(0.008)	(0.008)		(0.027)	(0.027)	(0.027)	(0.024)	(0.025)
Own EA $[t+7]$		0.039***	0.039***	0.036***	0.035***	0.029***		0.275***	0.273***	0.266***	0.240***	0.213***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.027)	(0.027)	(0.027)	(0.024)	(0.024)
Own EA $[t+8]$		0.029***	0.027***	0.020**	0.024***	0.013		0.305***	0.304***	0.285***	0.272***	0.233***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.027)	(0.027)	(0.026)	(0.024)	(0.024)
Own EA $[t+9]$		0.016*	0.014*	−0.000	0.010	−0.008		0.360***	0.358***	0.324***	0.324***	0.271***
		(0.008)	(0.008)	(0.008)	(0.008)	(0.008)		(0.027)	(0.027)	(0.026)	(0.024)	(0.023)
Own EA $[t+10]$		0.016*	0.012	−0.001	0.007	−0.009		0.478***	0.473***	0.434***	0.440***	0.382***
		(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		(0.025)	(0.025)	(0.024)	(0.023)	(0.022)
Own 8-K $[t]$			−0.074***	−0.074***	−0.067***	−0.066***			−0.134***	−0.151***	−0.075***	−0.076***
			(0.007)	(0.007)	(0.007)	(0.007)			(0.016)	(0.015)	(0.011)	(0.011)
Own key development $[t]$			0.121***	0.122***	0.122***	0.125***			0.146***	0.134***	0.091***	0.100***
			(0.005)	(0.005)	(0.004)	(0.004)			(0.012)	(0.010)	(0.009)	(0.008)
Positive news coverage $[t]$			0.313***	0.315***	0.305***	0.307***			0.382***	0.377***	0.341***	0.351***
			(0.009)	(0.009)	(0.008)	(0.008)			(0.017)	(0.017)	(0.013)	(0.013)
Negative news coverage $[t]$			−0.513***	−0.513***	−0.505***	−0.504***			−0.584***	−0.601***	−0.467***	−0.465***
			(0.011)	(0.011)	(0.010)	(0.010)			(0.020)	(0.020)	(0.014)	(0.014)
log(1 + Increase in news volume)			0.122***	0.123***	0.123***	0.125***			0.193***	0.195***	0.106***	0.117***
			(0.008)	(0.008)	(0.007)	(0.007)			(0.014)	(0.014)	(0.011)	(0.011)
log(1 + Decrease in news volume)			0.157***	0.157***	0.154***	0.155***			0.219***	0.210***	0.112***	0.113***
			(0.006)	(0.006)	(0.006)	(0.006)			(0.013)	(0.012)	(0.010)	(0.010)
Day FE	no	no	no	yes	no	yes	no	no	no	yes	no	yes
Firm x Year-quarter FE	no	no	no	no	yes	yes	no	no	no	no	yes	yes
Observations	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572	32,631,572
R-squared	.000	.000	.001	.005	.026	.029	.000	.000	.000	.005	.151	.154

This table shows the return effect of a memory cue. In column 1, the return over $[t,t+1]$ is regressed on the cue dummy without any further controls. Column 2 adds 21 dummy variables that capture a potential own earnings announcement in the window $[t-10,t+10]$. Column 3 adds dummy variables that capture the filing of an own 8-K, an own Key Development, and positive/negative news coverage. This column also controls for the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. Columns 4 and 5 add day fixed effects and firm x year-quarter fixed effects, respectively. Column 6 adds both day and firm x year-quarter fixed effects simultaneously. Columns 7 - 12 mirror columns 1 - 6, except that the dependent variable is the return over $[t,t+10]$. Panel A shows results for the main sample, in which memory associations are estimated using only earnings announcements that are exogenously shifted by calendar rotations. Panel B shows results for a broader sample, in which memory associations are estimated using all earnings announcements of all firms. The sample in panel A covers years 2005 - 2019 and the sample in panel B covers years 1995 - 2020. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

fixed effects, suggesting that announcements of large and salient firms are a strong cue for many firms.[14]

In the fifth column, I replace the day fixed effects with firm-year-quarter fixed effects. These fixed effects control for a whole host of potentially time-varying unobservable characteristics of firm i and ensure that the effect is estimated using only variation between cue and noncue days within a firm-quarter. The effect remains significant with these fixed effects, but its magnitude reduces somewhat to 3.8 bps, which corresponds to 4.8% in annual terms. Finally, in the sixth column, I include both day fixed effects and firm-year-quarter fixed effects simultaneously. While this saturated fixed effect structure may appear attractive, in my setting it removes important variation necessary to estimate the return effect. In particular, the day fixed effect removes variation resulting from the fact that when a cueing firm announces earnings, this typically represents a cueing event for many firms.[15] Thus, if such a cue affects the return of many firms on day t, the day fixed effect will capture this, resulting in a small and insignificant coefficient on the cue dummy, which is precisely what I find in the sixth column. This is an important shortcoming of including both day fixed effects and firm-year-quarter fixed effects simultaneously. Therefore, in the remaining tests in the paper, I focus on the specification with the full set of controls and firm-year-quarter fixed effects.

In the seventh through the twelfth column, I present the same regressions, except that the dependent variable is the return over the window $[t, t+10]$. The magnitude of the coefficient on the cue dummy is markedly larger, suggesting that the effect of a cue takes some time to be fully impounded into prices. Taking the estimate from the eleventh column (with all controls and firm-year-quarter fixed effects), the effect of a cue is about 8.4 bps of abnormal return over a window of about 2 weeks.[16]

In Figure 2, I document the dynamics of this effect in more detail by plotting the coefficient on the cue dummy from 50 separate regressions. The dependent variable in each regression is the cumulative return over a different return window, ranging from $[t, t+1]$ to $[t, t+50]$. All regressions include firm-year-quarter fixed effects and the full set of controls from Table 2, and I cluster

[14] In Section 3.3, I show that cues from large firms are indeed driving the documented return effect. This result is plausible, since investors likely pay more attention to the earnings announcements of large firms.

[15] In my tests, I assume that a cueing firm's announcement of earnings triggers the recall of all firms that announced on the same day as the cueing firm in the previous quarter. Under this assumption, an announcement of a cueing firm triggers the recall of 48 firms on average (median: 39). Presumably, the announcement is a stronger cue for some firms than others, depending on the strength of the underlying memory association. In Section 3.4, I present targeted tests designed to examine how the strength of the underlying association affects the strength of the documented return effect.

[16] Some readers may worry that the stronger effects over this longer return window are merely picking up the joint effect of multiple cues that occur in the window. As Table 1 shows, a cue occurs on about 0.6% of firm-days in the sample. For about half of these days (0.3% of all firm-days in the sample), another cue occurs in the next 10 days. In Appendix Table B.1, I show that the results are similar when I control for all future cues that occur in the window $[t, t+10]$.

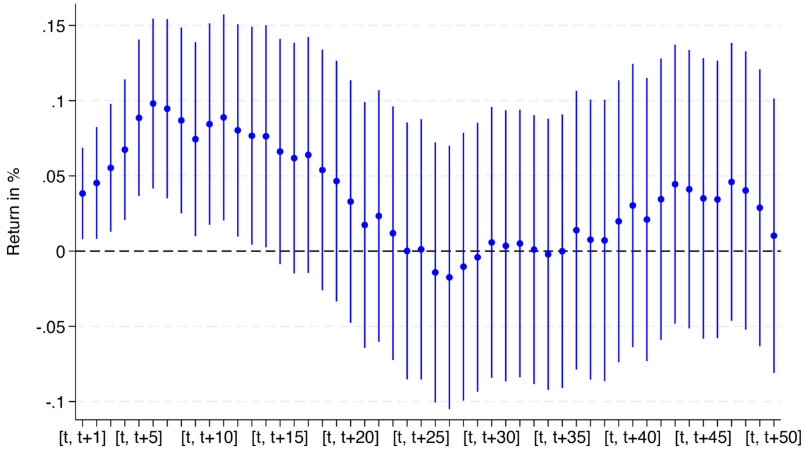

Figure 2
Return dynamics and long-run reversal
This figure traces out the return dynamics in response to a memory cue by showing results from 50 separate regressions. The dependent variable in each regression is the cumulative return over a different return window, ranging from $[t,t+1]$ to $[t,t+50]$. The dots in the figure represent the coefficients on the cue dummy, along with 95% confidence intervals, from each regression. All regressions include the full set of controls from Table 2 and firm-year-quarter fixed effects. Standard errors are clustered by firm and trading day.

standard errors by firm and day. The dots in the figure represent the coefficients on the cue dummy, along with 95% confidence intervals, from each regression. The first and the tenth dot in the figure correspond to the coefficients on the cue dummy from columns 5 and 11 in panel A of Table 2.

Figure 2 shows that the full effect of a memory cue realizes over the course of about 1-2 weeks, and maxes out at 10 bps around $t+6$. The effect then fully reverses and remains close to zero after about $t+25$. This reversal makes it unlikely that my results are driven by fundamental relationships. If they were, there should be no systematic reversal. Since pure memory-induced buying pressure carries no new information, prices should eventually revert to their fundamental values, which is precisely what I find. Overall, these results support the hypothesis that memory-induced attention can lead to temporary buying pressure in cued stocks.

To put the magnitude and horizon of this effect into perspective, it is helpful to compare it to the effects of externally generated attention and categorization documented in previous work. Da, Engelberg, and Gao (2011) find that an increase in retail investor attention leads to an abnormal return of about 34 bps over 2 weeks, with an (almost) full reversal after about 1 year. Thus, the effect of memory-induced attention materializes over a similar time period as the effect in Da, Engelberg, and Gao (2011), but is only about one-third as strong. Further, the 1-year reversal in Da, Engelberg, and Gao (2011) is considerably longer than the 5-week reversal that I document in this paper. On the other hand, Hartzmark and Shue (2018) find that the reversal of contrast effects takes about 50 trading days, which is more similar to my 5-week timeline. Hartzmark and Shue (2018) is a useful benchmark for my results because the contrast effects in that paper are also estimated using earnings announcements of other firms. In a study on categorization, Barberis, Shleifer, and Wurgler (2005) find that category and habitat effects have a horizon of at least 1 month. Since the memory effect that I document is almost fully reversed after 1 month, its horizon is at the lower end of the category and habitat effects documented in Barberis, Shleifer, and Wurgler (2005).[17]

As a final step in this section, I show the importance of an appropriate identification strategy to cleanly estimate memory effects. In my baseline design in panel A, the scheduling of earnings announcements is orthogonal to the news that firms disclose. In panel B of Table 2, I replicate the results from panel A, but instead of using only earnings announcements that are exogenously shifted by calendar rotations to estimate memory associations, I use all earnings announcements. This sample therefore also includes earnings announcements that are strategically rescheduled in response to information, as documented by a large literature (Penman 1987; Bagnoli, Kross, and Watts 2002; Johnson and So 2018).

Since I do not require data on Pattern firms for these tests, I am able to extend the sample period to years 1995 - 2020. I choose 1995 as the starting point, since DellaVigna and Pollet (2009) show that the accuracy of the earnings date is near perfect after December 1994. In the first column of panel B, I find that a cueing event leads to a daily abnormal return of 5 bps in this sample. This effect size is similar to the effect documented in panel A. However, as I increasingly control for firm news events across specifications, the magnitude of this effect drops substantially. Indeed, once I include firm-year-quarter fixed effects, the effect is zero and insignificant because these fixed effects capture persistent fundamental news about the firm. The results are similar for the window $[t, t+10]$ in columns 7 through 12. In sum, the results in panel B highlight the importance of an identification strategy that can cleanly separate the effects of a pure memory cue from potential information effects. Therefore, in all tests going forward, I will focus on the much cleaner sample used in panel A.

3.2 Retail order imbalance results

My tests in the previous section build on the strong existing evidence showing that individual investors are net buyers of attention-grabbing stocks, to the point where they can create positive price pressure (Barber and Odean 2008; Da, Engelberg, and Gao 2011). In this section, I provide direct evidence for the

[17] One possible reason the memory effect I document has a shorter lifespan than the category and habitat effects in Barberis, Shleifer, and Wurgler (2005) is that the memory effect documented in this paper is estimated based on contextual similarity alone, and ignores association due to similarity in firm characteristics. It is possible that memory associations based on similarity in characteristics are significantly more durable.

assumption that there is net retail buying in firm i's stock in response to a memory cue.

My tests are simple. I regress retail order imbalance on the cue dummy, controlling for the same firm-level controls as in Table 2 and firm-year-quarter fixed effects. Retail order imbalance is positive if there is net buying, and negative if there is net selling. Thus, I hypothesize that the coefficient on the cue dummy is positive.

Table 3 presents the results. In the first column, the dependent variable is average retail order imbalance constructed using the BHJOS algorithm over the window $[t, t+1]$. While the coefficient is not significantly different from zero, when I expand the window to $[t, t+10]$ in the second column, I do find a significantly positive effect. In the third and fourth columns, I replace the

Table 3
Retail order imbalance results

Dependent variable:	BHJOS retail OI in %		BJZZ retail OI in % (all lots)		BJZZ retail OI in % (odd lots)	
Window:	$[t,t+1]$ (1)	$[t,t+10]$ (2)	$[t,t+1]$ (3)	$[t,t+10]$ (4)	$[t,t+1]$ (5)	$[t,t+10]$ (6)
Cue $[t]$	0.179	0.220***	0.031	0.118*	0.280	0.329***
	(0.140)	(0.070)	(0.116)	(0.064)	(0.185)	(0.096)
Own EA $[t-10]$	0.173	0.096*	0.091	−0.039	−0.219	−0.175**
	(0.110)	(0.057)	(0.099)	(0.051)	(0.141)	(0.071)
Own EA $[t-9]$	0.274**	0.094	0.123	−0.036	0.013	−0.105
	(0.112)	(0.058)	(0.097)	(0.052)	(0.144)	(0.073)
Own EA $[t-8]$	0.428***	0.110*	0.143	−0.036	0.018	−0.114
	(0.111)	(0.059)	(0.097)	(0.053)	(0.143)	(0.074)
Own EA $[t-7]$	0.330***	0.134**	0.116	−0.018	−0.002	−0.115
	(0.111)	(0.060)	(0.096)	(0.053)	(0.135)	(0.077)
Own EA $[t-6]$	0.353***	0.136**	0.129	−0.017	−0.100	−0.112
	(0.109)	(0.060)	(0.095)	(0.054)	(0.141)	(0.079)
Own EA $[t-5]$	0.458***	0.156***	0.127	−0.031	−0.006	−0.121
	(0.113)	(0.060)	(0.100)	(0.054)	(0.144)	(0.079)
Own EA $[t-4]$	0.593***	0.200***	0.194**	−0.014	0.193	−0.070
	(0.112)	(0.061)	(0.097)	(0.054)	(0.136)	(0.080)
Own EA $[t-3]$	0.671***	0.223***	0.193**	−0.022	−0.045	−0.081
	(0.111)	(0.061)	(0.097)	(0.055)	(0.138)	(0.081)
Own EA $[t-2]$	0.474***	0.228***	−0.021	−0.034	−0.277*	−0.077
	(0.110)	(0.062)	(0.097)	(0.056)	(0.149)	(0.082)
Own EA $[t-1]$	0.491***	0.282***	−0.231**	−0.041	−0.752***	−0.161*
	(0.111)	(0.063)	(0.100)	(0.056)	(0.151)	(0.083)
Own EA $[t]$	0.312***	0.250***	−0.487***	−0.065	−1.202***	−0.279***
	(0.116)	(0.064)	(0.100)	(0.057)	(0.171)	(0.084)
Own EA $[t+1]$	3.665***	0.894***	2.677***	0.488***	2.661***	0.334***
	(0.114)	(0.061)	(0.102)	(0.054)	(0.166)	(0.083)
Own EA $[t+2]$	5.024***	1.199***	4.352***	0.741***	4.290***	0.582***
	(0.128)	(0.060)	(0.117)	(0.053)	(0.176)	(0.081)
Own EA $[t+3]$	2.352***	1.325***	2.203***	0.894***	1.744***	0.718***
	(0.119)	(0.059)	(0.106)	(0.053)	(0.162)	(0.081)
Own EA $[t+4]$	1.575***	1.421***	1.435***	0.991***	1.257***	0.825***
	(0.114)	(0.059)	(0.102)	(0.053)	(0.163)	(0.083)
Own EA $[t+5]$	1.074***	1.460***	1.045***	1.050***	0.899***	0.875***
	(0.112)	(0.059)	(0.102)	(0.053)	(0.151)	(0.083)
Own EA $[t+6]$	0.753***	1.491***	0.828***	1.122***	0.608***	0.888***
	(0.114)	(0.058)	(0.098)	(0.053)	(0.152)	(0.083)
Own EA $[t+7]$	0.685***	1.501***	0.814***	1.179***	0.413***	0.901***
	(0.111)	(0.059)	(0.101)	(0.053)	(0.156)	(0.083)

(Continued)

Table 3
Continued

Dependent variable:	BHJOS retail OI in %		BJZZ retail OI in % (all lots)		BJZZ retail OI in % (odd lots)	
Window:	[$t,t+1$] (1)	[$t,t+10$] (2)	[$t,t+1$] (3)	[$t,t+10$] (4)	[$t,t+1$] (5)	[$t,t+10$] (6)
Own EA [$t+8$]	0.707***	1.512***	0.769***	1.232***	0.391***	0.945***
	(0.112)	(0.059)	(0.103)	(0.054)	(0.150)	(0.083)
Own EA [$t+9$]	0.679***	1.566***	0.745***	1.329***	0.476***	1.023***
	(0.110)	(0.058)	(0.097)	(0.053)	(0.160)	(0.084)
Own EA [$t+10$]	0.783***	1.576***	0.643***	1.402***	0.451***	1.145***
	(0.106)	(0.058)	(0.094)	(0.053)	(0.153)	(0.082)
Own 8-K [t]	0.389***	0.044**	0.228***	−0.006	0.376***	0.030
	(0.051)	(0.019)	(0.045)	(0.019)	(0.075)	(0.030)
Own key development [t]	0.465***	0.133***	0.315***	0.091***	0.196***	0.080***
	(0.034)	(0.016)	(0.032)	(0.015)	(0.047)	(0.023)
Positive news coverage [t]	0.340***	0.108***	0.241***	0.091***	−0.073	0.043
	(0.049)	(0.024)	(0.046)	(0.022)	(0.075)	(0.035)
Negative news coverage [t]	−0.075	−0.061**	−0.098**	−0.057**	0.001	0.011
	(0.049)	(0.024)	(0.046)	(0.023)	(0.080)	(0.038)
log(1 + Increase in news volume)	0.398***	0.012	0.276***	−0.015	0.556***	0.028
	(0.044)	(0.019)	(0.036)	(0.017)	(0.059)	(0.028)
log(1 + Decrease in news volume)	0.160***	0.034	0.125***	0.031	0.333***	0.041
	(0.042)	(0.021)	(0.038)	(0.020)	(0.080)	(0.038)
Firm x Year-quarter FE	yes	yes	yes	yes	yes	yes
Observations	9,159,341	8,430,871	8,038,145	6,975,409	3,243,285	2,701,475
R-squared	.119	.398	.103	.376	.143	.494

This table shows how retail order imbalance varies in response to a memory cue. In columns 1 and 2, retail order imbalance from the Barber et al. (2024) (BHJOS) algorithm is regressed on the cue dummy and a battery of control variables. In columns 3 and 4, the dependent variable is retail order imbalance from the Boehmer et al. (2021) (BJZZ) algorithm. Columns 5 and 6 mirror columns 3 and 4, except that retail order imbalance is calculated using only odd lot trades (trades of fewer than 100 shares). All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *$p<.1$; **$p<.05$; ***$p<.01$.

dependent variable with retail order imbalance constructed using the BJZZ algorithm. For this measure of retail order imbalance, the coefficient on the cue dummy is smaller and statistical significance is weaker. One possible reason for this weaker result is that the BJZZ algorithm does not sign retail trades as accurately as the BHJOS algorithm (Barber et al. 2024). Finally, in the fifth and sixth columns, I use retail order imbalance from the BJZZ algorithm as the dependent variable, but construct it only from odd lot trades. Odd lot trades are trades of fewer than 100 shares. Given the smaller trade size, these trades are arguably more likely to be retail trades. Using this measure of retail order imbalance, I find that the effect in column 6 is approximately three times as strong as the effect in column 4 and about 50% stronger than the effect in column 2.[18]

In Figure 3, I show the dynamics of retail order imbalance. Each dot in the figure represents the coefficient on the cue dummy from a separate regression,

[18] In Appendix Table C.1, I show that the results are very similar when I calculate retail order imbalance using the number of retail buys and sells instead of retail buying and selling volume.

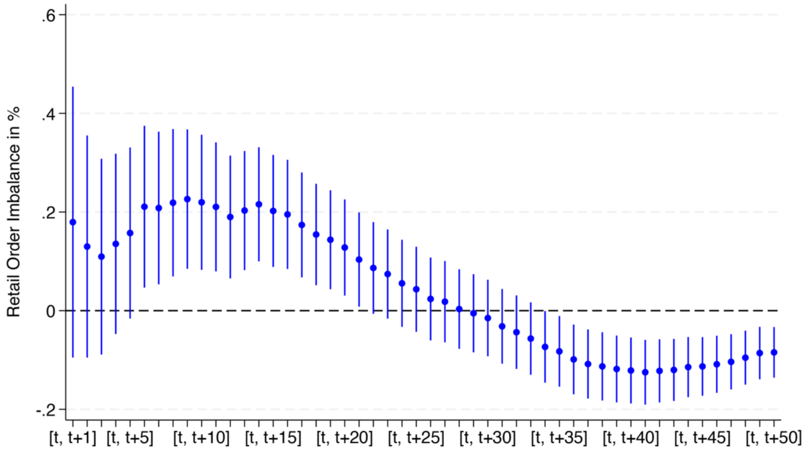

Figure 3
Retail order imbalance dynamics
This figure traces out the dynamics of retail order imbalance, constructed using the Barber et al. (2024) algorithm, in response to a memory cue by showing results from 50 separate regressions. The dependent variable in each regression is the average retail order imbalance over a different time window, ranging from $[t, t+1]$ to $[t, t+50]$. The dots in the figure represent the coefficients on the cue dummy, along with 95% confidence intervals, from each regression. All regressions include the full set of controls from Table 3 and firm-year-quarter fixed effects. Standard errors are clustered by firm and trading day.

where the dependent variable in each regression is the average retail order imbalance over a different time window, ranging from $[t, t+1]$ to $[t, t+50]$. All regressions include the full set of control variables and fixed effects from Table 3, and I cluster standard errors by firm and day.

Figure 3 shows the results for retail order imbalance constructed using the BHJOS algorithm. In the days after a cue, retail order imbalance is consistently positive and increasing. Since the dependent variable is average retail order imbalance for expanding time windows, a rise in the coefficient from one time period to the next means that there was more retail buying in that time period compared to the historical average. Conversely, a fall in the coefficient means that there was more selling in that time period compared to the historical average. The dynamics of retail order imbalance match the dynamics of the cumulative returns displayed in Figure 2. In particular, when there is more retail buying, there tends to be an increase in the return. Conversely, the return reversal occurs precisely when there is more selling. In Appendix Figures A.1 and A.2, I replicate Figure 3 for retail order imbalance constructed from the BJZZ algorithm, using either all lots or only odd lot trades. The look of these figures is similar to Figure 3.

Overall, these results illustrate the chain of events more clearly, by connecting the dots between the cueing event, the trading behavior of retail investors, and the ultimate return effect documented in the previous section. Further, the finding that there is aggregate retail trading in response to memory

Table 4
Cross-sectional cuts along firm size

Dependent variable:	Return $[t, t+10]$ (%)		
Size of cued firm:	All (1)	Small (2)	Large (3)
Cue (size of cueing firm > 90th pctile)	0.093***	0.168***	0.021
	(0.032)	(0.062)	(0.032)
Cue (size of cueing firm in 51st - 90th pctile)	0.009	0.052	−0.036
	(0.024)	(0.048)	(0.024)
Cue (size of cueing firm in 10th - 50th pctile)	−0.037	−0.057	−0.038*
	(0.023)	(0.045)	(0.023)
Cue (size of cueing firm < 10th pctile)	0.017	−0.054	0.068**
	(0.035)	(0.066)	(0.031)
Controls	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes
Observations	16,502,158	8,238,578	8,262,465
R-squared	.156	.161	.178

This table shows how the return effect of a memory cue varies with the size of the cueing and cued firm. In all columns, the dependent variable is the return over $[t, t+10]$. The main independent variables are four dummy variables that capture cues from firms at different points of the size distribution. Column 1 shows the results for the full sample, while columns 2 and 3 split the sample along the median market capitalization, and show the results for small and large cued firms, respectively. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. These fixed effects can result in singleton observations that are dropped during the estimation, which explains why the sum of the number of observations in columns 2 and 3 is slightly lower than the number of observations in column 1. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. $*p < .1$; $**p < .05$; $***p < .01$.

cues is consistent with Charles (2022), who also finds that memory associations determine which stocks investors choose to trade. Finally, by showing the underlying retail trading behavior, the results in this section help rule out the alternative explanation that overlapping earnings announcements may create return comovement among associated firms, which might be picked up by institutional algorithm-based trading in the next quarter.

3.3 Cross-sectional cuts along firm size

In this section, I test whether the documented return effect varies with firm size. In my tests so far, I focus on large and salient cues by requiring cueing firms to have a market capitalization above the 90th percentile. However, it is possible that earnings announcements from smaller firms may also attract enough attention to serve as cues. To test for this possibility, I augment the baseline regression from Table 2 with additional cue dummies that capture the incremental effects of cueing events from smaller firms.

The first column of Table 4 presents the results. I find that cues from the largest firms are driving the effect, while cues from firms with market capitalization in the 1st to the 90th percentile do not lead to significant effects. These results are plausible, as investors are more likely to hear about an earnings announcement if the firm is large.

I also explore whether the documented return effect varies with the size of the cued firm. On the one hand, conditional on a cue, large firms might come easier to mind. On the other hand, large firms have more liquid stocks, making buying pressure less likely to occur. To test for these possibilities, I split the sample along median firm size, using the market capitalization from $t-3$, and show the effects separately for large and small cued firms in columns 2 and 3 of Table 4. I find that the effect is driven by small cued firms.

3.4 Testing the psychological mechanism

The goal of this section is to dig deeper into the psychological mechanism that drives the documented results. I begin by testing whether attention during the encoding of a memory association modulates the strength of the association. Then, I turn to testing the two key forces of associative memory theory: similarity and interference. I close by exploring whether the timing (or contiguity) of the earnings announcements matters for the strength of the underlying association.

3.4.1 Attention during encoding First, I evaluate the role of attention during encoding for the strength of the underlying memory association. The intuition of my tests can be illustrated best with the introductory example. Recall that last quarter, firms A and B announced earnings on the same day. I hypothesize that the association between the two firms is stronger if investors paid more attention to them. While I do not have a perfect measure of investor attention, Barber and Odean (2008) suggest three proxies for whether investors are paying attention to a firm: high abnormal trading volume, high media coverage, and extreme returns. I test all three proxies in turn, beginning with abnormal trading volume.

In my tests, I classify an association as “strong” if abnormal trading volume of both firms was above the median on the day of the encoding. I calculate abnormal trading volume as in Barber and Odean (2008). When calculating the median cutoff, I use only observations of Pattern firms on earnings announcement dates. I focus on earnings announcement dates because abnormal trading volume is systematically higher on earnings announcement days than on other days, and I want to identify cases where abnormal volume is high even for an earnings announcement day. I choose to focus only on Pattern firms when calculating the cutoff to ensure that a sufficient number of Pattern firms falls above and below the cutoff. Finally, I calculate the median cutoff for each year separately to account for time-varying trends in the amount of trading.

Having identified strong associations in this way, I regress the return of firm i over the window $[t, t+10]$ on four dummies. The first dummy is equal to one if the cueing firm’s size is above the 90th percentile and if the underlying association is strong (i.e., both the cueing and cued firm had above-median trading volume on the day of the encoding). The second dummy captures all

Table 5
Testing the psychological mechanism

Panel A

Dependent variable:	Return $[t, t+10]$ in %				
Association stronger due to joint:	High trading volume (1)	High BHJOS retail volume (2)	High BJZZ retail volume (3)	High media coverage (4)	Top mover (5)
Size of cueing firm > 90th pctile					
Cue (strong association)	0.176***	0.248***	0.149*	0.074*	0.257
	(0.053)	(0.078)	(0.088)	(0.039)	(0.566)
Cue (all other associations)	0.039	0.058*	0.080**	0.097**	0.085***
	(0.035)	(0.033)	(0.033)	(0.046)	(0.031)
Size of cueing firm in 51st - 90th pctile					
Cue (strong association)	0.031	0.047	−0.008	−0.005	−0.028
	(0.040)	(0.053)	(0.066)	(0.035)	(0.210)
Cue (all other associations)	−0.020	−0.011	−0.001	0.001	−0.001
	(0.030)	(0.027)	(0.026)	(0.032)	(0.025)
Controls	yes	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.156	.156	.156	.156	.156
p-value (strong = other, large cues)	.0232	.0217	.4551	.6953	.7620
p-value (strong = other, small cues)	.2664	.2937	.9103	.8863	.8989

(Continued)

other cues from firms with size above the 90th percentile. These other cues target associations that are not classified as strong based on this classification. I hypothesize that the coefficient on the "strong cue" dummy is larger in magnitude than the coefficient on the dummy for all other cues. The third and fourth dummies mirror the first and second, except that they capture the incremental effects of (strong) cues from smaller firms, namely those in the 51st to the 90th percentile of size. While Table 4 shows that cues from smaller firms do not lead to significant return effects on average, it is possible that cues from smaller firms that target strong associations do lead to significant effects. The third and fourth dummies allow me to evaluate this possibility.

I present the results from this regression in the first column in panel A of Table 5. Focusing first on cues from large firms, I can test whether cues that target strong underlying associations have stronger effects than cues that target other associations by comparing the coefficients on the first two dummies. Indeed, the coefficient on the strong cue dummy is more than four times as large. I show the p-value of a test for the equality of the two coefficients at the bottom of the table and find that the difference between the two coefficients is statistically significant ($p < .025$). Looking at cues from smaller firms, I do not find significant effects, even for cues that target a strong underlying association.

In the second and third columns, I define an association as strong if both firms had above-median abnormal retail trading volume on the day of the encoding. In the second column, I use abnormal retail volume from the Barber et al. (2024) algorithm. I again find that the effect is about four times larger if the underlying association is strong, a difference that is statistically significant

Table 5
Continued
Panel B

Dependent variable:	Return $[t, t+10]$ in %						
Association stronger due to joint:	Pos. surprise (6)	Neg. surprise (7)	Extreme pos. surprise (8)	Extreme neg. surprise (9)	Low interference (10)	EA market hours (11)	EA after hours (12)
Size of cueing firm > 90th pctile							
Cue (strong association)	0.082	0.220	0.860	−5.055***	0.092**	0.079*	0.090
	(0.057)	(0.143)	(1.994)	(0.645)	(0.041)	(0.045)	(0.061)
Cue (all other associations)	0.087**	0.070**	0.085***	0.086***	0.083*	0.086**	0.084**
	(0.035)	(0.031)	(0.031)	(0.031)	(0.048)	(0.040)	(0.034)
Size of cueing firm in 51st - 90th pctile							
Cue (strong association)	−0.012	0.385***	−0.268	−1.280	0.022	0.018	−0.003
	(0.045)	(0.093)	(0.806)	(0.874)	(0.032)	(0.037)	(0.042)
Cue (all other associations)	0.001	−0.016	−0.001	−0.001	−0.028	−0.012	−0.001
	(0.027)	(0.026)	(0.025)	(0.025)	(0.040)	(0.030)	(0.028)
Controls	yes	yes	yes	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes	yes	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.156	.156	.156	.156	.156	.156	.156
p-value (strong = other, large cues)	.9355	.2939	.6973	.0000	.8925	.9113	.9277
p-value (strong = other, small cues)	.7877	.0000	.7406	.1432	.3268	.4638	.9637

This table shows how the return effect of a memory cue varies with the strength of the underlying memory association. In all columns, the dependent variable is the return over $[t, t+10]$. The main independent variables are four dummies. The first dummy is equal to one if the cue targets a strong underlying memory association and if the cueing firm's size is above the 90th percentile. The column header indicates which features during the encoding determine whether the underlying association is strong. The second dummy captures all other cues from firms with size above the 90th percentile. The third and fourth dummies mirror the first and second, except that they capture the incremental effects of (strong) cues from firms in the 51st to the 90th percentile of size. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, two dummy variables that capture positive and negative news coverage, respectively, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Panel B is a continuation of panel A. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. p-values for tests on the equality of the coefficients on first and second dummy, and the third and fourth dummy, are displayed at the bottom of the table. *$p < .1$; **$p < .05$; ***$p < .01$.

($p < .025$). In the third column, I instead use abnormal retail volume from the Boehmer et al. (2021) algorithm. Here, too, I find a larger effect when the underlying association is strong, but in this case, the difference is smaller and not statistically significant. Further, I do not find a significant effect from strong cues if the cueing firm is small for either measure of abnormal retail volume.

As a second proxy for investor attention, I test whether the underlying association is stronger if both firms were covered heavily in the media. I use RavenPack's Aggregate Event Volume to capture the amount of media coverage that a firm received. I do not find that high joint media coverage during encoding leads to a stronger effect. One reason for this result might be that investors are sufficiently aware of the earnings announcements of large firms, even if these announcements are not covered heavily in the media.

As a third proxy for investor attention, I focus on extreme returns. Many news outlets prominently display lists of stocks with extreme returns. One particularly salient list is the Top-Mover list of the fintech brokerage Robinhood (Barber et al. 2022).[19] I recreate the Top-Mover list for each earnings announcement day by ranking all stocks in CRSP based on their daily absolute return, and then identify the top 20 stocks on this list. When constructing the ranking, I follow Barber et al. (2022) and only consider stocks with a market capitalization above \$300 million on that day. Using the resulting Top-Mover list, I define an association as strong if both firms were on the Top-Mover list on the day of the encoding.

In the fifth column, I present the results. While the coefficient on the strong cue dummy is positive and sizable, it is not significant. One possible reason for this lack of significance is lack of statistical power: it is rare for two firms to both be on the Top-Mover list, and it is especially rare for large firms to be on the Top-Mover list because their returns are simply not volatile enough.

Overall, the results in columns 1 through 5 are very suggestive that attention during the encoding plays an important role in modulating the strength of the underlying association. However, an important caveat to these results is that the effects of strong cues and other cues are not always statistically different from each other.

3.4.2 Similarity during encoding I now evaluate the role of similarity, one of the two key forces of associative memory theory. In associative memory theory, two experiences share a stronger association if their features are more similar. To test this prediction in my setting, I test whether firms that have similar earnings surprises during the encoding also share a stronger association. Specifically, I classify an association as strong either if both firms had positive earnings surprises during the encoding, or if both firms had negative earnings surprises during the encoding. This approach also allows me to test whether

[19] While the trading app of Robinhood only became available in 2015, similar lists are available from many other news outlets in the earlier years of my sample (e.g., Yahoo! Finance, *Wall Street Journal*, and CNBC).

the nature of the experience (positive or negative) affects the nature of the underlying association.

Columns 6 and 7 in Panel B present the results of these tests. I find that only cues targeting negative associations lead to stronger return effects, and that this effect is particularly strong for small cueing firms. Notably, the return effects are positive, even though both firms had negative earnings surprises during the encoding. While not conclusive, these results provide support for the prediction that similarity strengthens the underlying association, even if the nature of the experience is negative.

These results appear to suggest that the nature of the experience does not affect the nature of the underlying association. However, in the test presented in column 7, many of the earnings surprises are only mildly negative. It may be that when two firms have *extremely* negative earnings surprises during encoding, the resulting association leads to a negative return response. To test for this possibility, I identify extremely positive and extremely negative associations based on the extremeness of the earnings surprises. I classify associations as extremely positive if the absolute earnings surprise of both firms was in the top decile and positive, and I classify associations as extremely negative if the absolute surprise of both firms was in the top decile and negative.

In columns 8 and 9, I find that cues targeting extremely positive associations lead to strong positive return effects in terms of magnitude, but these effects are not statistically significant. Conversely, cues targeting extremely negative associations lead to strong, significant, and very negative return responses. This result suggests that the return response is only negative when the underlying memory association is extremely negative and encoded based on rare and salient tail events.

3.4.3 Interference during encoding Next, I test the second key force of associative memory theory, interference, and explore whether it dampens the documented return effect. The intuition is straightforward. If a memory association between firms A and B was encoded on a day that many other firms also announced earnings, the strength of this association should be weaker. The reason is that on such days investors encode associations not only between firms A and B but also between firm A and these other firms. As a result, when cued with an earnings announcement by firm A on day t, investors might not recall firm B, but one of these other firms instead. Put differently, the memories of these other firms interfere with the recall of firm B. Thus, I hypothesize that the return response of firm B is stronger if interference is lower.

Applying this intuition to my setting, I classify associations as “strong” if they were encoded on days that the number of other firms announcing was below the median. In column 10, I show that the effect is slightly stronger for cues from associations with low interference. However, this difference is not statistically significant and hence this result is only suggestive.

3.4.4 Contiguity during encoding In this next set of tests, I explore whether the timing of the earnings announcements matters for the strength of the underlying association. Two firms that both announce during market hours might share a stronger association than a pair of firms where one announces during market hours and the other one after market hours. In the terms of associative memory theory, the two firms announcing during market hours share a stronger association because the context of their earnings announcements is more similar. The idea that two items share a stronger association in memory if they were experienced closer in time together is typically referred to as "contiguity" (Kahana 2012). In columns 11 and 12, I do not find that associations are stronger if both announcements occurred either during or after market hours.

In sum, when evaluating the results in all of the columns jointly, Table 5 provides suggestive, but not conclusive, evidence of associative memory as the underlying psychological mechanism. Similarity, interference, and contiguity are deep laws of memory that memory theorists have discussed at length (Kahana 2012). Tests of these laws have typically been done at the individual level in an experimental setting, and even recent work that applies memory theory to economics and finance focuses on experimental tests (Bordalo et al. 2023, Enke, Schwerter, and Zimmermann 2024, Gödker, Jiao, and Smeets Forthcoming). In contrast, in Table 5, I construct tests of these laws using market data. Overall, the results are largely consistent with the predictions of the theory, but the differences between coefficients are often not statistically significant. Compared to experimental tests at the individual level, implementing tests at the market level is challenging because the data have been aggregated and there is a substantial amount of noise. When evaluating the results jointly, however, the signature patterns of associative memory theory do stand out.

3.5 Alternative cues

In my tests so far, I focus on earnings announcements as cueing events. This decision is based on associative memory theory, which predicts that when two firms announce earnings on the same day, the shared context of this joint announcement generates the memory association. Thus, cues that bring to mind the context "earnings announcement" are effective at targeting the association. Naturally, an earnings announcement itself is a strong cue for this context and hence triggers the recall of the associated firm through the mechanism of contextual similarity.

Here, I test whether other firm events can also serve as cues for the recall of these associations. As a first set of firm events, I use Capital IQ Key Developments. I focus on the set of events used in Kwon and Tang (2023), since these are events that occur frequently and to which investors are likely paying attention. In Appendix Table D.1, I classify each event as being either an "accounting-based event" or a "non-accounting-based event". Since accounting-based events are similar in context to earnings announcements, they might be strong cues for the memory associations in my sample, which are estimated based on overlapping earnings announcements. Examples of accounting-based events include announcements of operating results, changes in corporate guidance, and impairments/write-offs. Similarly, examples of non-accounting-based events include business expansions, executive changes, and product and client announcements. Beyond Capital IQ Key Developments, I also use filing dates of Form 8-Ks as a second set of firm events.

To implement my tests, I augment Equation (1) with several dummy variables that are equal to one if these alternative events occur at a memory-associated firm on day t. Importantly, this regression also includes the baseline cue dummy that captures the cueing effect of an earnings announcement. Including this control is important because it rules out that other firm events that occur close in time to an earnings announcement are erroneously identified as effective cues. Further, to compare the effect sizes of alternative cues and earnings announcement cues, I consider all alternative cueing events that occur in the same calendar quarter as the earnings announcement of a cueing firm. Finally, I include all cues from firms with size above the 90th percentile of market capitalization, since Table 4 shows that these are by far the strongest cues.

Table 6 presents the results. I begin by regressing firm i's return over the window $[t, t+10]$ on four dummies, where each dummy is equal to one if one of the aforementioned events occurs at a memory-associated (cueing) firm on day t. Column 1 shows that, unconditionally, none of the alternative cues leads to significant effects. However, it is possible that these alternative cues are only effective if the underlying memory association is particularly strong. Therefore, in the remaining columns of the table, I break out the effect in a similar way as in Table 5.

The results can be summarized in the following way. Alternative cues can lead to significant effects, but only if the underlying memory association is particularly strong. Non-accounting-based Key Developments can act as significant cues if both firms had high abnormal trading volume on the day of the memory encoding. In terms of magnitude, the effect of such an alternative cue is about 20% as strong as the effect of an earnings announcement cue. Similarly, accounting-based Key Developments can act as strong and significant cues if both firms were Top-Movers on the day of the memory encoding. While both accounting-based and non-accounting-based Key Developments can sometimes act as cues, overall, earnings announcements are by far the strongest cues. These results are consistent with associative memory theory and highlight the importance of contextual similarity between the cue and the underlying association during recall.

Table 6
Alternative cues
Panel A

Dependent variable:	Return $[t, t+10]$ in %					
Association stronger due to joint:	n/a (1)	High trading volume (2)	High BHJOS retail volume (3)	High BJZZ retail volume (4)	High media coverage (5)	Top mover (6)
Earnings announcement						
Cue (baseline)	0.086**					
	(0.034)					
Cue (strong association)		0.198***	0.282***	0.149	0.074*	0.294
		(0.056)	(0.083)	(0.093)	(0.043)	(0.579)
Cue (all other associations)		0.035	0.057	0.081**	0.099**	0.085**
		(0.039)	(0.036)	(0.036)	(0.050)	(0.034)
Key Development (accounting)						
Cue (Baseline)	0.019					
	(0.015)					
Cue (strong association)		0.029	0.045	−0.017	0.023	0.689**
		(0.027)	(0.043)	(0.050)	(0.021)	(0.345)
Cue (all other associations)		0.015	0.016	0.021	0.014	0.018
		(0.018)	(0.016)	(0.016)	(0.023)	(0.015)
Key Development (nonaccounting)						
Cue (baseline)	0.014					
	(0.012)					
Cue (strong association)		0.038**	0.052**	0.026	0.020	0.095
		(0.018)	(0.026)	(0.030)	(0.014)	(0.184)
Cue (all other associations)		0.004	0.009	0.013	0.008	0.014
		(0.013)	(0.012)	(0.012)	(0.017)	(0.012)
Form 8-K						
Cue (baseline)	−0.000					
	(0.013)					
Cue (strong association)		0.003	0.006	0.000	−0.030*	0.329
		(0.023)	(0.035)	(0.040)	(0.017)	(0.320)
Cue (all other associations)		−0.002	−0.001	0.000	0.031	−0.001
		(0.014)	(0.013)	(0.013)	(0.019)	(0.013)
Controls	yes	yes	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.156	.156	.156	.156	.156	.156
p-value (strong = other, EA)	-	.0095	.0094	.4873	.6931	.7186
p-value (strong = other, KeyDev acc.)	-	.6368	.5110	.4426	.7680	.0517
p-value (strong = other, KeyDev nonacc.)	-	.0742	.0996	.6695	.5653	.6592
p-value (strong = other, 8-K)	-	.8636	.8450	.9930	.0135	.3024

(Continued)

Table 6
Continued
Panel B

Dependent variable:	Return $[t,t+10]$ in %						
Association stronger due to joint:	Pos. surprise (7)	Neg. surprise (8)	Extreme pos. surprise (9)	Extreme neg. surprise (10)	Low interference (11)	EA market hours (12)	EA after hours (13)
Earnings announcement							
Cue (strong association)	0.077	0.411***	0.779	−5.257***	0.105**	0.088*	0.095
	(0.061)	(0.150)	(1.983)	(1.472)	(0.043)	(0.048)	(0.066)
Cue (all other associations)	0.088**	0.074**	0.086**	0.086**	0.065	0.084*	0.084**
	(0.039)	(0.034)	(0.034)	(0.034)	(0.054)	(0.044)	(0.037)
Key Development (accounting)							
Cue (strong association)	0.011	0.012	−0.167	−1.164	0.030	0.014	0.054
	(0.027)	(0.081)	(2.058)	(2.765)	(0.020)	(0.022)	(0.036)
Cue (all other associations)	0.021	0.019	0.019	0.019	0.003	0.021	0.013
	(0.017)	(0.016)	(0.015)	(0.015)	(0.025)	(0.020)	(0.016)
Key Development (nonaccounting)							
Cue (strong association)	0.015	0.010	0.290	0.233	0.023	0.006	0.042*
	(0.018)	(0.046)	(0.631)	(0.871)	(0.014)	(0.015)	(0.023)
Cue (all other associations)	0.014	0.014	0.014	0.014	−0.000	0.020	0.008
	(0.013)	(0.012)	(0.012)	(0.012)	(0.020)	(0.014)	(0.012)
Form 8-K							
Cue (strong association)	0.026	0.031	−2.054**	−0.210	0.014	−0.015	−0.008
	(0.022)	(0.057)	(0.975)	(1.124)	(0.017)	(0.019)	(0.027)
Cue (all other associations)	−0.007	−0.001	0.000	−0.000	−0.016	0.009	0.001
	(0.015)	(0.013)	(0.013)	(0.013)	(0.019)	(0.016)	(0.014)
Controls	yes	yes	yes	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes	yes	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.156	.156	.156	.156	.156	.156	.156
p-value (strong = other, EA)	.8766	.0233	.7264	.0003	.5580	.9472	.8697
p-value (strong = other, KeyDev acc.)	.7446	.9344	.9280	.6687	.3991	.7915	.2816
p-value (strong = other, KeyDev nonacc.)	.9770	.9252	.6620	.8015	.3364	.4604	.1388
p-value (strong = other, 8-K)	.1931	.5783	.0352	.8521	.2253	.2981	.7589

This table shows how the return effect of a memory cue varies with the type of cueing event as well as with the strength of the underlying memory association. In all columns, the dependent variable is the return over $[t,t+10]$. The independent variables in all columns capture cueing events from firms above the 90th percentile of size. In column 1, the main independent variables are four dummies that capture the effects of different types of cueing events: Earnings Announcements, Key Developments (accounting-based), Key Developments (nonaccounting-based), and filings of Form 8-Ks. In all other columns, the main independent variables are eight dummies, two for each type of cueing event, that break out whether the cueing event targets a strong underlying memory association. The column header indicates which features during the encoding determine whether the underlying association is strong. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10,t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, two dummy variables that capture positive and negative news coverage, respectively, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Panel B is a continuation of panel A. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. For each type of cueing event, the p-value from a test on the equality of the coefficients on the two dummies is displayed at the bottom of the table. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

Table 7
Anticipation of earnings announcements

Return window:	$[t-3]$ (1)	$[t-2]$ (2)	$[t-1]$ (3)	$[t-1]$ (4)
Cue $[t]$	−0.012	−0.004	−0.003	
	(0.009)	(0.009)	(0.010)	
Cue after hours $[t-1]$				−0.005
				(0.015)
Cue market hours $[t]$				−0.008
				(0.011)
Controls	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes
Observations	16,501,071	16,501,152	16,501,645	16,501,645
R-squared	.015	.014	.014	.014

This table tests whether anticipation of a cueing event can lead to return effects. In columns 1 - 3, the return on days $t-3$, $t-2$, and $t-1$, respectively, is regressed on a dummy variable that captures a cue on day t. In column 4, the return on day $t-1$ is regressed on a dummy that captures a cue that occurred after the market closes on day $t-1$ as well as a dummy that captures a cue that occurred during market hours on day t. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

4. Further Results and Robustness

4.1 Anticipation of earnings announcements

In this section, my goal is to evaluate whether the anticipation of an earnings announcement is sufficient to trigger a market move. Prior evidence indicates that attention to firms increases immediately before they announce earnings (Chapman 2018). Further, since the earnings announcements of Pattern firms are on a perfectly predictable schedule, investors might anticipate these earnings announcements ahead of time.[20] On the other hand, it might be that the announcement itself is required to act as a salient cue.

To distinguish these two possibilities, I replace the dependent variable in Equation (1) with returns from days prior to the cue. In the first three columns of Table 7, I show the results for the return on days $t-3$, $t-2$, and $t-1$ as the dependent variable, respectively. The coefficient on the cue dummy is a relatively precisely estimated zero in all three columns. This strongly suggests that the anticipation of an announcement is not strong enough to act as a cue.

As an alternative falsification test, I leverage the fact that about one-third of earnings announcements occur after the market closes. I test whether an announcement that occurred on day $t-1$ after market hours affects the return on day $t-1$. If anticipation of that event suffices to act as a cue, the effect might already manifest itself in the return on day $t-1$. Conversely, if it is the announcement itself that triggers the effect, there should be no return effect on

[20] More generally, earnings announcements are typically scheduled at least a week ahead of time (Boulland and Dessaint 2017).

Table 8
Surprise of the cue

Dependent variable:	Return $[t, t+10]$ in %			
Surprise:	EW		VW	
	(1)	(2)	(3)	(4)
Cue $[t]$	0.089**		0.090**	
	(0.035)		(0.035)	
Cue x Surprise $[t]$	−7.678		−10.606	
	(18.797)		(18.246)	
Cue x Surprise quintile 1 $[t]$		0.029		0.022
		(0.047)		(0.047)
Cue x Surprise quintile 2 $[t]$		−0.030		0.124
		(0.142)		(0.141)
Cue x Surprise quintile 3 $[t]$		0.175***		0.169***
		(0.062)		(0.061)
Cue x Surprise quintile 4 $[t]$		0.172***		0.151**
		(0.060)		(0.060)
Cue x Surprise quintile 5 $[t]$		0.028		0.046
		(0.067)		(0.066)
Controls	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes
Observations	16,502,158	16,502,158	16,502,158	16,502,158
R-squared	.156	.156	.156	.156

This table tests whether the earnings surprise of the cueing firm(s) can predict the return response of the cued firm. The earnings surprise is the difference between the actual earnings announced by the cueing firm and the median analyst earnings forecast, scaled by the share price of the firm from three trading days prior to the announcement. If there are multiple cues for firm i on day t, columns 1 and 2 use the equally weighted average surprise, and columns 3 and 4 use the value-weighted average surprise. In columns 1 and 3, the cue dummy is interacted with the surprise, while in columns 2 and 4, the cue dummy is interacted with dummy variables for each quintile of the surprise distribution. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

day $t-1$. In the fourth column of Table 7, I show that the coefficient on the after hours cue is zero and insignificant, providing further evidence that the anticipation of an announcement is not sufficient to trigger an effect.

4.2 Surprise of the cue

In this next set of tests, I explore whether the earnings surprise of the cueing firm predicts the return response of the cued firm. These tests are designed to address the potential concern that my results might be driven by information spillover from the cueing firm's earnings announcement. Such spillovers might manifest themselves in a systematic relationship between the earnings surprise of the cueing firm and the return response of the cued firm. For instance, more positive surprises might lead to higher returns and more negative surprises might lead to lower returns (Thomas and Zhang 2008). In contrast, if the earnings announcement purely acts as a memory cue that directs attention, the earnings surprise is unlikely to play an important role for the strength of the effect.

In Table 8, I regress the return of firm i over the window $[t, t+10]$ on the cue dummy as well as the interaction of the cue dummy with the earnings surprise of the cueing firm(s). I also test whether the effect varies along the distribution of the cueing firms' earnings surprise, to account for potential nonlinear relationships. On days with multiple cues, I calculate either the equally weighted average (first and second columns) or the value-weighted average (third and fourth columns) of the cueing firms' earnings surprise.

I find that the surprise of the cueing firm does not have predictive power for the return response of the cued firm. While at first blush the coefficient on the interaction term in the first and third columns might appear economically large, this magnitude is driven by the fact that the earnings surprise variable has a tiny standard deviation (see Table 1). For instance, the coefficient in the first column implies that a one-standard-deviation increase in the earnings surprise of the cueing firm would decrease the return response of the cued firm by 1 bp. The nonparametric estimations in the second and fourth columns also show that there is no monotonic relationship between the earnings surprise of the cueing firm and the return response of the cued firm. In fact, the effect is strongest for cues with earnings surprises close to zero (quintile 3). Taken together, these findings help address the concern that my results might be picking up information spillovers from cueing to cued firms. In contrast, these results are wholly consistent with a memory-based explanation, in which a cueing firm's earnings announcement simply directs attention to the memory-associated firm.

4.3 Trading strategy

The earnings announcements of Pattern firms are perfectly predictable. Therefore, it is possible to construct a trading strategy that takes advantage of the buying pressure caused by memory-induced trading. This intentionally simple trading strategy is a long-short strategy that goes long stocks that were cued on day t and short stocks that were not cued. The long and short leg of the strategy are value-weighted portfolios using the market capitalization of each stock on day $t-3$. I use small firms for which the market capitalization from $t-3$ is below the median to form these portfolios, since these firms drive the effect (see Table 4).

To account for the potential role of risk factors, I regress the time series of returns generated by this trading strategy on the market, size, value, momentum, and short-term reversal factors, which are sourced from the Kenneth French Data Library. Table 9 shows the return of the trading strategy for three sets of five years each: 2005 - 2009, 2010 - 2014, and 2015 - 2019. For the most recent set of years, the strategy yields an alpha of 69.3 bps over the window $[t, t+10]$, which is significant at the 1% level. This corresponds to an annualized abnormal return of about 16%. It is worth noting that the trading strategy can only be implemented if there is a cueing event (i.e., an earnings announcement) and if there are stocks that fall into the long and short leg of

Table 9
Trading strategy over time

Return window:	$[t, t+10]$		
Time period:	2005 - 2009 (1)	2010 - 2014 (2)	2015 - 2019 (3)
Alpha [%]	−0.265	0.224	0.693***
	(0.414)	(0.219)	(0.228)
Mkt	0.292	−0.027	0.124
	(0.329)	(0.264)	(0.287)
SMB	0.087	−0.163	0.637
	(0.569)	(0.477)	(0.456)
HML	−0.539	−0.141	0.573
	(0.599)	(0.529)	(0.486)
Momentum	−0.144	0.233	0.162
	(0.443)	(0.363)	(0.341)
ST reversal	0.249	0.211	0.487
	(0.340)	(0.410)	(0.412)
Observations	374	516	507
R-squared	.007	.002	.010

This table shows how the trading strategy performs in three different sets of 5 years each. The trading strategy return in each time period is regressed on the market, size, value, momentum, and short-term reversal factors, which are sourced from the Kenneth French Data Library. The trading strategy return is the return of a long-short strategy. The long leg of the strategy consists of a value-weighted portfolio of stocks with a cueing event on day t, and the short leg of the strategy consists of a value-weighted portfolio of stocks without a cueing event on day t. The weights in these portfolios are the market capitalization of each stock on day $t-3$. These portfolios are formed using firms whose market capitalization from $t-3$ is below the median. Standard errors are displayed in parentheses below the coefficients. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

the strategy. In my sample, the strategy can be implemented on an average of 100 trading days per year.

As Table 9 shows, the strategy is only profitable from 2015 onwards. While I can only speculate on what might be driving these time trends, these results might reflect the fact that the rise of trading platforms and apps has dramatically changed the way retail investors trade. One example is Robinhood's trading app, which launched in 2015 and which allows retail investors to trade more easily. The resulting increase in retail trading might lead to more memory-induced trades and could be driving some of these time trends. Such trading platforms also make certain types of curated information much more salient and available to retail investors, possibly affecting how information is processed, stored, and retrieved. For instance, Robinhood prominently displays the day's biggest winners and losers to investors via the Top-Mover list (Barber et al. 2022).

From a practical perspective, it may not be possible to profitably implement the proposed trading strategy. The strategy requires shorting a large number of small stocks, which might be practically infeasible. Further, the costs associated with shorting these stocks might be substantial, and could wipe out the profitability of the strategy. However, I want to emphasize that the purpose of illustrating this trading strategy is to show that the main result holds in a different specification, with a different risk adjustment. To this end, the results in Table 9 show that the loadings on the factors are insignificant

Table 10
Trading strategy: Long and short leg

Return window:	$[t,t+1]$	$[t,t+10]$	$[t,t+1]$	$[t,t+10]$
Strategy:	Long leg		Short leg	
	(1)	(2)	(3)	(4)
Alpha [%]	0.467**	0.853***	−0.028	0.160
	(0.184)	(0.271)	(0.046)	(0.147)
Mkt	0.776***	0.609*	0.908***	0.485***
	(0.232)	(0.342)	(0.058)	(0.186)
SMB	1.361***	1.814***	0.719***	1.177***
	(0.369)	(0.543)	(0.092)	(0.295)
HML	0.302	0.406	0.000	−0.167
	(0.393)	(0.578)	(0.098)	(0.315)
Momentum	−0.192	−0.463	−0.252***	−0.626***
	(0.276)	(0.406)	(0.069)	(0.221)
ST reversal	0.540	0.742	0.174**	0.254
	(0.333)	(0.490)	(0.083)	(0.267)
Observations	507	507	507	507
R-squared	.075	.053	.472	.089

This table shows how the long and short leg of the trading strategy perform in 2015−2019. Returns of the long and short leg are regressed on the market, size, value, momentum, and short-term reversal factors, which are sourced from the Kenneth French Data Library. In columns 1 and 2, the dependent variable is the return of the long leg of the strategy, which consists of a value-weighted portfolio of stocks with a cueing event on day t. In columns 3 and 4, the dependent variable is the return of the short leg of the strategy, which consists of a value-weighted portfolio of stocks without a cueing event on day t. The weights in these portfolios are the market capitalization of each stock on day $t-3$. These portfolios are formed using firms whose market capitalization from $t-3$ is below the median. Standard errors are displayed in parentheses below the coefficients. *$p<.1$; **$p<.05$; ***$p<.01$.

and economically small, and do not wash out the positive and significant alpha in recent years. Thus, a risk-based explanation is unlikely to explain my results.

In Table 10, I use the trading strategy setup to provide further evidence in favor of the hypothesis of memory-induced buying pressure. Under this hypothesis, the abnormal return of the trading strategy should be driven by the long leg of the strategy. When I break out the return of the strategy separately for the long and the short leg, I find that the return of the strategy is indeed driven entirely by the long leg. Overall, the results in this section highlight the robustness of my findings using calendar-time asset pricing methods.

4.4 Alternative explanations

My results support the hypothesis that memory-induced attention leads to buying pressure. In several tests aimed at the mechanism, I find support for key predictions of associative memory theory. Specifically, I find that higher attention during the encoding increases the strength of the underlying association. Further, I find that the similarity of the experiences during encoding (e.g., if both firms have negative earnings surprises) modulates the strength as well as the nature of the association. I also find suggestive evidence that the effect weakens if there were more distracting earnings announcements by other firms during the encoding of a memory association between two firms. These distracting events lead to interference in recall on the day of the cue.

Here, I discuss two alternative explanations for these results. The first possibility is that the results could be driven by fundamental relationships between cueing and cued firms, and/or information spillover from the cueing firm's earnings announcement. My tests are designed to rule out these possibilities. First, my tests leverage the fact that earnings announcement dates of Pattern firms are exogenously shifted by calendar rotations. Thus, the resulting overlaps cannot be driven by fundamental information. Second, I find that the return effect I document is temporary and reverses to zero after about 25 trading days. If my results were driven by fundamental relationships, there should be no systematic reversal. And third, I find that the earnings surprise of the cueing firm has no predictive power for the return response of the cued firm, making information spillover an unlikely explanation.

The second possibility is that the results are not driven by investors' memories, but instead by some form of external information archive that mimics the properties of memory. To organize the discussion, recall the introductory example where firms A and B announced earnings on the same day last quarter, but this quarter they do not. This alternative explanation posits that when firm A announces earnings this quarter, investors rediscover firm B, for example, by reading an archived newspaper article from last quarter, in which both firms A and B are covered. While this explanation might plausibly explain the baseline results, it must also explain the results from the mechanism tests in Tables 5 and 6. Specifically, it must explain why media coverage during encoding does *not* modulate the strength of the effect. If the effect was driven by investors rediscovering firms when accessing historically archived media coverage, the effect should be significantly stronger for firms whose announcements were covered more heavily.

Thus, for this explanation to work, information must be archived and accessed in very particular ways. Furthermore, to have aggregate effects, many investors must be using the same (or very similarly organized) archives. Associative memory provides one such archive, one with clear predictions from decades of experimental work. While it is difficult to fully rule out the alternative explanation of some external archive, associative memory provides a very parsimonious explanation.

5. Conclusion

In this paper, I provide evidence of memory effects in financial markets. I show that memory-induced attention creates buying pressure in the cued firm's stock. In tests aimed at the mechanism, I show that the documented effect varies with the strength of the underlying memory association. In particular, I show that attention during the encoding is an important determinant of the strength of the underlying association. Further, I find suggestive evidence that the two key forces of associative memory theory – similarity and interference – modulate the strength of the effect. I also find that the nature of the experience (positive or negative) drives the direction of the return effect, but that only extremely negative associations lead to negative return effects. Finally, I show that various firm events can serve as memory cues, but that the contextual similarity between the cue and the underlying association is a key determinant of the strength of the cue during recall.

Most existing tests of human memory are conducted at the individual level. Several recent studies focus on how memory constraints affect beliefs and individual decision-making (e.g., Charles 2022; Enke, Schwerter, and Zimmermann 2024; Jiang et al. 2024; Gödker, Jiao, and Smeets Forthcoming). In contrast, my setting allows me to show that the constraints of human memory can aggregate and affect asset prices. Overall, my results suggest that economic models of human memory can explain behavior outside the laboratory and at the market level.

My results also provide evidence consistent with the idea that internally generated attention can have effects on financial markets. This suggests that there is a whole class of internal attention sources, a class distinct from the external sources that have previously been investigated. In other words, the set of attention sources that is relevant for financial decision-making is potentially much larger than previously thought. Fleshing out these internal sources in more detail and testing their effects on financial markets could be a promising direction for future work.

Finally, models of human memory can potentially provide a microfoundation for the strong degree of categorization documented in the prior literature (Barberis and Shleifer 2003; Barberis, Shleifer, and Wurgler 2005; Peng and Xiong 2006; Huang 2019). In associative memory theory, the strength of an association between two items is determined both by the similarity of their features as well as the similarity of the context in which the two items were experienced. As such, associative memory theory naturally creates the categorization along characteristics (or features) documented in previous work, and suggests that other features – like contextual similarity – can also be the source of categorization. Models of human memory might thus be a useful organizing framework for categorical thinking and may help uncover other dimensions along which firms are categorized in investors' minds.

Code Availability: The replication code is available in the Harvard Dataverse at https://doi.org/10.7910/DVN/HEGXCC.

Appendix

A. Retail Order Imbalance Dynamics Using the BJZZ Algorithm

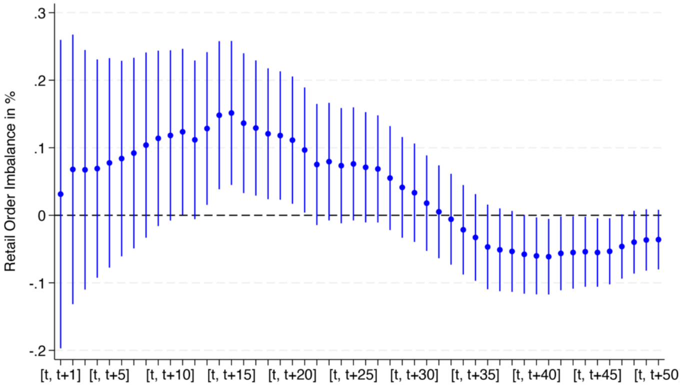

Figure A.1
Retail order imbalance dynamics (all lots)
This figure traces out the dynamics of retail order imbalance, constructed using the Boehmer et al. (2021) algorithm, in response to a memory cue by showing results from 50 separate regressions. The dependent variable in each regression is the average retail order imbalance over a different time window, ranging from $[t, t+1]$ to $[t, t+50]$. The dots in the figure represent the coefficients on the cue dummy, along with 95% confidence intervals, from each regression. All regressions include the full set of controls from Table 3 and firm-year-quarter fixed effects. Standard errors are clustered by firm and trading day.

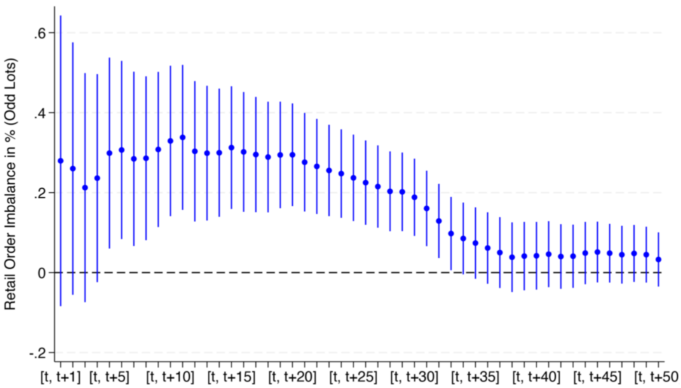

Figure A.2
Retail order imbalance dynamics (odd lots)
This figure traces out the dynamics of retail order imbalance, constructed using the Boehmer et al. (2021) algorithm, in response to a memory cue by showing results from 50 separate regressions. The dependent variable in each regression is the average retail order imbalance, constructed using only odd lot trades, over a different time window, ranging from $[t, t+1]$ to $[t, t+50]$. The dots in the figure represent the coefficients on the cue dummy, along with 95% confidence intervals, from each regression. All regressions include the full set of controls from Table 3 and firm-year-quarter fixed effects. Standard errors are clustered by firm and trading day.

B. Controlling for Future Cues

Table B.1
Controlling for future cues

Dependent variable:	Return $[t, t+10]$ in %	Retail order imbalance $[t, t+10]$ in %		
Retail OI measure:	n/a (1)	BHJOS (2)	BJZZ (all lots) (3)	BJZZ (odd lots) (4)
Cue $[t]$	0.065**	0.196***	0.106*	0.285***
	(0.029)	(0.062)	(0.057)	(0.085)
Cue $[t+1]$	0.047*	0.193***	0.114**	0.253***
	(0.028)	(0.057)	(0.052)	(0.079)
Cue $[t+2]$	0.066**	0.169***	0.100*	0.214***
	(0.028)	(0.057)	(0.052)	(0.078)
Cue $[t+3]$	0.054*	0.147***	0.097*	0.191**
	(0.028)	(0.055)	(0.050)	(0.079)
Cue $[t+4]$	0.055**	0.121**	0.080	0.156**
	(0.028)	(0.056)	(0.049)	(0.077)
Cue $[t+5]$	0.075***	0.019	0.030	0.073
	(0.027)	(0.055)	(0.048)	(0.075)
Cue $[t+6]$	0.046	0.182***	0.162***	0.127
	(0.028)	(0.056)	(0.049)	(0.079)
Cue $[t+7]$	0.025	0.194***	0.179***	0.166**
	(0.028)	(0.060)	(0.053)	(0.081)
Cue $[t+8]$	0.006	0.211***	0.202***	0.170**
	(0.026)	(0.061)	(0.056)	(0.082)
Cue $[t+9]$	0.003	0.277***	0.233***	0.205**
	(0.027)	(0.064)	(0.058)	(0.085)
Cue $[t+10]$	−0.036	0.414***	0.359***	0.258***
	(0.028)	(0.070)	(0.064)	(0.093)
Controls	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes
Observations	16,502,158	8,430,871	6,975,409	2,701,475
R-squared	.156	.398	.376	.494

This table shows the main results from Tables 2 and 3, when controlling for 10 dummies that capture future cues occurring in the window $[t, t+10]$. In column 1, the dependent variable is the return over $[t, t+10]$. In columns 2 - 4, the dependent variable is average retail order imbalance over $[t, t+10]$ from the Barber et al. (2024) (BHJOS) algorithm, the Boehmer et al. (2021) (BJZZ) algorithm, and the BJZZ algorithm using only odd lot trades. All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10, t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. $*p < .1$; $**p < .05$; $***p < .01$.

C. Retail Order Imbalance Using Number of Trades

Table C.1
Retail order imbalance using number of trades

Dependent variable:	BHJOS retail OI in %		BJZZ retail OI in % (all lots)		BJZZ retail OI in % (odd lots)	
Window:	$[t,t+1]$ (1)	$[t,t+10]$ (2)	$[t,t+1]$ (3)	$[t,t+10]$ (4)	$[t,t+1]$ (5)	$[t,t+10]$ (6)
Cue $[t]$	0.102	0.207***	−0.024	0.121*	0.247	0.339***
	(0.132)	(0.067)	(0.113)	(0.063)	(0.161)	(0.084)
Controls	yes	yes	yes	yes	yes	yes
Firm x Year-quarter FE	yes	yes	yes	yes	yes	yes
Observations	9,159,341	8,430,871	8,038,145	6,975,409	3,243,285	2,701,475
R-squared	.141	.424	.140	.427	.191	.539

This table replicates Table 3 with retail order imbalance constructed using the number of retail trades each day. In columns 1 and 2, retail order imbalance from the Barber et al. (2024) (BHJOS) algorithm is regressed on the cue dummy and a battery of control variables. In columns 3 and 4, the dependent variable is retail order imbalance from the Boehmer et al. (2021) (BJZZ) algorithm. Columns 5 and 6 mirror columns 3 and 4, except that retail order imbalance is calculated using only odd lot trades (trades of fewer than 100 shares). All columns include the following control variables: 21 dummy variables that capture a potential own earnings announcement in the window $[t-10,t+10]$, a dummy variable that captures the filing of an own 8-K, a dummy variable that captures an own Key Development, a dummy variable each for positive/negative news coverage, and the logarithm of 1 plus the absolute increase or decrease in the volume of news coverage. All columns also include firm x year-quarter fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *$p<.1$; **$p<.05$; ***$p<.01$.

D. Classification of Key Developments

Table D.1
Classification of key developments

Event	Accounting based event	*keydeveventtypeid* in WRDS
Announcement of operating results	Yes	226
Annual general meeting	No	62
Business expansions	No	31
Changes in company bylaws/rules	No	77
Client announcements	No	23
Corporate guidance - Lowered	Yes	26
Corporate guidance - New/confirmed	Yes	29
Corporate guidance - Raised	Yes	27
Credit rating - S&P - Credit Watch/Outlook Action	No	72
Discontinued operations/Downsizings	No	21
Dividend affirmations	Yes	45
Earnings calls	Yes	48
Executive/Board changes - Other	No	16
Executive changes - CEO	No	101
Executive changes - CFO	No	102
Impairments/Write-offs	Yes	73
Lawsuits/Legal issues	No	25
M&A rumors and discussions	No	65
M&A transaction announcements	No	80
M&A transaction closings	No	81
Product-related announcements	No	41
Seeking acquisitions/investments	No	3
Strategic alliances	No	22

This table shows the Key Developments used in Kwon and Tang (2023) (except for earnings announcements). Each event is classified as an accounting-based or non-accounting-based event. Identifiers from the WRDS Capital IQ Key Developments database for each event type are also shown (*keydeveventtypeid*).

References

Azeredo da Silveira, R., Y. Sung, and M. Woodford. 2024. Optimally imprecise memory and biased forecasts. *American Economic Review* 114:3075–118.

Bagnoli, M., W. Kross, and S. G. Watts. 2002. The information in management's expected earnings report date: A day late, a penny short. *Journal of Accounting Research* 40:1275–96.

Barber, B. M., E. T. De George, R. Lehavy, and B. Trueman. 2013. The earnings announcement premium around the globe. *Journal of Financial Economics* 108:118–38.

Barber, B. M., X. Huang, P. Jorion, T. Odean, and C. Schwarz. 2024. A (sub)penny for your thoughts: Tracking retail investor activity in taq. *Journal of Finance* 79:2403–27.

Barber, B. M., X. Huang, T. Odean, and C. Schwarz. 2022. Attention-induced trading and returns: Evidence from robinhood users. *Journal of Finance* 77:3141–90.

Barber, B. M., S. Lin, and T. Odean. 2019. Mediating investor attention. Working Paper, University of California, Berkeley.

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

Barberis, N., and A. Shleifer. 2003. Style investing. *Journal of Financial Economics* 68:161–99.

Barberis, N., A. Shleifer, and J. Wurgler. 2005. Comovement. *Journal of Financial Economics* 75:283–317.

Bodoh-Creed, A. L. 2020. Mood, memory, and the evaluation of asset prices. *Review of Finance* 24:227–62.

Boehmer, E., C. M. Jones, X. Zhang, and X. Zhang. 2021. Tracking retail investor activity. *Journal of Finance* 76:2249–305.

Bordalo, P., J. J. Conlon, N. Gennaioli, S. Y. Kwon, and A. Shleifer. 2023. Memory and probability. *Quarterly Journal of Economics* 138:265–311.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2020. Memory, attention, and choice. *Quarterly Journal of Economics* 135:1399–442.

Boulland, R., and O. Dessaint. 2017. Announcing the announcement. *Journal of Banking & Finance* 82:59–79.

Chapman, K. 2018. Earnings notifications, investor attention, and the earnings announcement premium. *Journal of Accounting and Economics* 66:222–43.

Charles, C. 2022. Memory and trading. Working Paper, London School of Economics.

Chen, T., Z. Gao, J. He, W. Jiang, and W. Xiong. 2019. Daily price limits and destructive market behavior. *Journal of Econometrics* 208:249–64.

Chen, X., L. An, Z. Wang, and J. Yu. 2023. Attention spillover in asset pricing. *Journal of Finance* 78:3515–59.

Chen, Y., R. B. Cohen, and Z. K. Wang. 2022. Famous firms, earnings clusters, and the stock market. Working Paper, University of Rochester.

Colonnelli, E., N. J. Gormsen, and T. McQuade. 2024. Selfish corporations. *Review of Economic Studies* 91:1498–536.

Da, Z., J. Engelberg, and P. Gao. 2011. In search of attention. *Journal of Finance* 66:1461–99.

Daniel, K., M. Grinblatt, S. Titman, and R. Wermers. 1997. Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52:1035–58.

DellaVigna, S., and J. M. Pollet. 2009. Investor inattention and friday earnings announcements. *Journal of Finance* 64:709–49.

Enke, B., F. Schwerter, and F. Zimmermann. 2024. Associative memory, beliefs and market interactions. *Journal of Financial Economics* 157:1038–53.

Frazzini, A., and O. Lamont. 2007. The earnings announcement premium and trading volume. Working Paper, Yale School of Management.

Frésard, L., G. Hoberg, and G. M. Phillips. 2020. Innovation activities and integration through vertical acquisitions. *Review of Financial Studies* 33:2937–76.

Frydman, C., and B. Wang. 2020. The impact of salience on investor behavior: Evidence from a natural experiment. *Journal of Finance* 75:229–76.

Gilboa, I., and D. Schmeidler. 1995. Case-based decision theory. *Quarterly Journal of Economics* 110:605–39.

Gödker, K., P. Jiao, and P. Smeets. Forthcoming. Investor memory. *Review of Financial Studies* .

Goetzmann, W. N., A. Watanabe, and M. Watanabe. 2022. Evidence on retrieved context: How history matters. Working Paper, Yale University.

Hartzmark, S. M. 2015. The worst, the best, ignoring all the rest: The rank effect and trading behavior. *Review of Financial Studies* 28:1024–59.

Hartzmark, S. M., S. D. Hirshman, and A. Imas. 2021. Ownership, learning, and beliefs. *Quarterly Journal of Economics* 136:1665–717.

Hartzmark, S. M., and K. Shue. 2018. A tough act to follow: Contrast effects in financial markets. *Journal of Finance* 73:1567–613.

Hartzmark, S. M., and D. H. Solomon. 2018. Recurring firm events and predictable returns: The within-firm time series. *Annual Review of Financial Economics* 10:499–517.

Hirshleifer, D., S. S. Lim, and S. H. Teoh. 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance* 64:2289–325.

Hirshleifer, D., and I. Welch. 2002. An economic approach to the psychology of change: Amnesia, inertia, and impulsiveness. *Journal of Economics & Management Strategy* 11:379–421.

Hoberg, G., and G. Phillips. 2010. Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies* 23:3773–811.

———. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124:1423–65.

Huang, X. 2019. Mark twain's cat: Investment experience, categorical thinking, and stock selection. *Journal of Financial Economics* 131:404–32.

Jiang, L., J. Liu, L. Peng, and B. Wang. 2022. Investor attention and asset pricing anomalies. *Review of Finance* 26:563–93.

Jiang, Z., H. Liu, C. Peng, and H. Yan. 2024. Investor memory and biased beliefs: Evidence from the field. Working Paper, Kellogg School of Management.

Johnson, T. L., and E. C. So. 2018. Time will tell: Information in the timing of scheduled earnings news. *Journal of Financial and Quantitative Analysis* 53:2431–64.

Kahana, M. J. 2012. *Foundations of human memory*. Oxford, UK: Oxford University Press.

Kwon, S. Y., and J. Tang. 2023. Extreme events and overreaction to news. Working Paper, Harvard University.

Lou, D. 2014. Attracting investor attention through advertising. *Review of Financial Studies* 27:1797–829.

Mullainathan, S. 2002. A memory-based model of bounded rationality. *Quarterly Journal of Economics* 117:735–74.

Nagel, S., and Z. Xu. 2022. Asset pricing with fading memory. *Review of Financial Studies* 35:2190–245.

Noh, S., E. C. So, and R. S. Verdi. 2021. Calendar rotations: A new approach for studying the impact of timing using earnings announcements. *Journal of Financial Economics* 140:865–93.

Patton, A. J., and M. Verardo. 2012. Does beta move with news? firm-specific information flows and learning about profitability. *Review of Financial Studies* 25:2789–839.

Peng, L., and W. Xiong. 2006. Investor attention, overconfidence and category learning. *Journal of Financial Economics* 80:563–602.

Penman, S. H. 1987. The distribution of earnings news over time and seasonalities in aggregate stock returns. *Journal of Financial Economics* 18:199–228.

Schmidt, D. 2019. Distracted institutional investors. *Journal of Financial and Quantitative Analysis* 54:2453–91.

Seasholes, M. S., and G. Wu. 2007. Predictable behavior, profits, and attention. *Journal of Empirical Finance* 14:590–610.

Thomas, J., and F. Zhang. 2008. Overreaction to intra-industry information transfers? *Journal of Accounting Research* 46:909–40.

Wachter, J. A., and M. J. Kahana. 2024. A retrieved-context theory of financial decisions. *Quarterly Journal of Economics* 139:1095–147.