Memory Moves Markets*

Constantin Charles

November 2, 2022

Please click here for the latest version

Abstract

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 they were experienced in the same context by many investors. 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. Overall, my results suggest that economic models of human memory can explain behavior outside the laboratory and at the market level.

*I especially thank my advisor Cary Frydman for his invaluable advice and guidance. I also thank Kenneth Ahern, Colin Camerer, Tom Chang, Joey Engelberg, Sam Hartzmark, David Hirshleifer, Chad Kendall, Mete Kilic, Chris Parsons, Vesa Pursiainen, Kelly Shue, and Felix von Meyerinck for helpful comments and suggestions. The author is from Marshall School of Business, University of Southern California. Contact: ccharles@marshall.usc.edu.

1 Introduction

There is a large literature in finance analyzing the role of attention in financial markets (see Barber et al. (2019) for an overview). 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.

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 et al. (2020); Wachter and Kahana (2021)). 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 et al. (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

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

orthogonal to firm fundamentals. In addressing these challenges, I am guided by the model of Bordalo et al. (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 Coca-Cola 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 IBM, because IBM announced earnings on the same day as Coca-Cola 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 Coca-Cola and IBM in the previous quarter, the investor may recall one of these other firms instead of IBM. Thus, the memories of these other firms *interfere* with the recall of IBM on day t.

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

In my empirical tests, I follow the logic of this example and estimate which firms are associated in investors' memory using overlaps of earnings announcements in previous quarters. In my main specification, I regress the return of firm j on day t (IBM) on a dummy variable that is equal to one if at least one firm that announced earnings on the same day as firm j in the previous quarter (Coca-Cola) announces earnings on day t. In all my tests, I ensure that firm j does not have an own earnings announcement within $[t-3,\ t+3]$, to avoid confounding effects from a firm's own earnings announcement. Further, I only retain firm pairs that operate in very different industries, using the text-based industries of Hoberg and Phillips (2010, 2016).

[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 IBM.

I find that a memory cue leads to a daily abnormal return of 3 – 5 basis points for the cued firm. In annual terms, this corresponds to an abnormal return of 7.5 – 12.5%. Risk-based return movements are unlikely to explain these results since I use characteristic-adjusted returns as in Daniel et al. (1997). In terms of magnitude, this effect size is on par with the earnings announcement premium (Frazzini and Lamont (2007); Barber et al. (2013)). It is also possible to construct a trading strategy that takes advantage of predictable memory cues and generates a daily alpha of 8.7 basis points.

My empirical strategy is designed to capture memory associations that are orthogonal to firm fundamentals, by exploiting the fact that most firms follow rigid schedules for their earnings announcements (Hartzmark and Shue (2018); Noh et al. (2021)). For instance, Coca-Cola typically announces earnings on the third Tuesday, Wednesday, or Thursday of the first month of a calendar quarter, while IBM typically announces on the 18th or 19th of the first month of a calendar quarter. These rigid schedules result in overlaps for some quarterly earnings announcements but not for others, depending on how the calendar shakes out in each quarter.

Nevertheless, one may be concerned that not all firms follow rigid schedules or that some firms strategically advance or delay their earnings announcement depending on the earnings they plan to report (Penman (1987); Bagnoli et al. (2002); Johnson and So (2018)). To directly address these concerns, I replicate all my results using only earnings announcements that were exogenously shifted by calendar rotations (Noh et al. (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. All my results are very similar for the subsample of Pattern firms.

[3] Such patterns are sometimes explicitly required by the firms' bylaws (Noh et al. (2021)).

As a deeper test of the mechanism, I next examine how the psychological properties of memory modulate the strength of the effect. First, I test for contiguity in the memory effect. The intuition of contiguity can be best illustrated with an example. Recall that in the opening example Coca-Cola and IBM announced earnings on the same day last quarter, but this quarter they do not. Assume further that last quarter Apple announced earnings one day after Coca-Cola announced. Because Coca-Cola and Apple announced with a gap of one day in the previous quarter, they share a weaker memory association than Coca-Cola and IBM (who announced on the same day). Applying this intuition to my setting, I test whether an earnings announcement of Coca-Cola on day t results in a weaker return response for Apple than for IBM, since the memory association between Coca-Cola and Apple is weaker. I find strong evidence in favor of this prediction. The memory effect drops sharply if memories were encoded with a timing gap of just one day, which is consistent with the differing contexts during the encoding of the memory association.

These tests also help alleviate concerns that my results are driven by similar firms announcing at similar times during an earnings season. The identifying assumption in these tests is that amongst a set of firms announcing close in time to each other (e.g., late in the quarter), firms that announce *on the same day* are not systematically different from firms that announce with *a gap of one day*. This assumption is likely satisfied, especially for the set of Pattern firms whose earnings announcement dates are shifted by exogenous calendar rotations.

Second, I test whether recent experiences have a stronger effect than distant experiences, a classic prediction of memory models. Indeed, I find a stronger effect if memory associations were encoded in more recent quarters, and I find that the effect fades away with increasing quarterly lags. In addition to validating a key prediction of memory theory, this finding is helpful in further attenuating concerns that my results are driven by attention to fundamentally related firms. If my results were purely driven by the fact that fundamentally related firms are more likely to announce quarterly earnings on the same day, the coefficients on each quarterly lag should be approximately similar in magnitude. However, I observe a systematic decay in the effect, consistent with a memory channel.

Finally, I test whether interference weakens the documented effect. As discussed in the introductory example, if many firms announced earnings together with Coca-Cola and IBM, 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 et al. (2009) find that attention to one firm reduces attention to another firm on earnings announcement days. This finding adds plausibility to the idea that during the encoding of memory traces, earnings announcements by other firms distract from the encoding of an associative memory trace between Coca-Cola and IBM. When I test for interference directly, I find 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. These findings provide evidence of interference effects, a signature prediction of associative memory theory.

In additional tests, I find that the documented effect quickly reverses, a finding that is helpful in further addressing the concern that my results might be affected by fundamental relationships between firms. I also explore whether the earnings surprise of the cueing firm (Coca-Cola) predicts the return response of the cued firm (IBM). If my results were driven by information spillover from the cueing firm's earnings announcement, this might manifest itself in a systematic relationship between the earnings surprise and the return response. 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. Consistent with the latter, I do not find any relationship between the surprise of the cueing firm and the return response of the cued firm. I discuss further alternative explanations in Section 5.6.

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 et al. (2009); Da et al. (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 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 et al. (2020); Wachter and Kahana (2021); Bordalo et al. (2022); Da Silveira et al. (2020); Bodoh-Creed (2020); Nagel and Xu (2022)).

My paper differs from previous empirical tests of memory models, as these tests largely focus on individual decision-making (e.g., Enke et al. (2022); Gödker et al. (2022); Colonnelli et al. (2022); Goetzmann et al. (2022)). 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.

2 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 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 et al. (2009)).

I consider two firms as associated in memory if they announced earnings on the same day in any of the previous four fiscal quarters. This approach has the benefit of being simple while capturing the main forces of associative memory theory (Kahana (2012); Bordalo et al. (2020)). For instance, by focusing on associations that were encoded in the previous quarter vs. a more distant quarter, I can test for the effect of recency in recall (Murdock Jr (1962); Chang et al. (2017)). Similarly, 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 et al. (2020); Bordalo et al. (2022)).

My empirical approach builds on the fact that many firms follow rigid rules for their earnings announcement timing (Hartzmark and Shue (2018); Noh et al. (2021)). For instance, some firms follow the rule to announce on the kth day-of-week of a calendar month (e.g., third Thursday), while other firms announce on the kth day in a calendar month (e.g., the 18th day). These rules result in overlaps for some earnings announcements but not for others, depending on how the calendar shakes out in each month. In my baseline tests, I estimate regressions of the following type:

$$return_{j,t} = \alpha + \beta \cdot cue_{j,t} + \gamma_t + u_{j,t} \tag{1}$$

where $return_{j,t}$ is firm j's characteristic-adjusted return on day t, $cue_{j,t}$ is a dummy variable that is equal to one if at least one firm that announced earnings on the same day as firm j in any of the previous four fiscal quarters announces earnings on day t, and γ_t is a day fixed effect.

The coefficient β captures the effect of a memory cue for firm j on day t. My hypothesis is that such a cue creates memory-induced attention and leads to buying pressure. Therefore, I hypothesize that β is positive. When estimating this specification, I ensure that firm j does not have an own earnings announcement within $[t-3,\ t+3]$ to avoid confounding effects from a firm's own earnings announcement. I cluster standard errors in all regressions by firm and day.

I also frequently focus on associations that were encoded in the previous quarter, since these associations are likely strongest and therefore easiest to recall. In these tests, I estimate regressions of the following type:

$$return_{j,t} = \alpha + \beta \cdot cue_{j,t,q-1} + \gamma_t + u_{j,t} \tag{2}$$

In this specification, $cue_{j,t,q-1}$ is a dummy variable that is equal to one if at least one firm that announced earnings on the same day as firm j in the previous quarter announces earnings on day t. The subscript "$q-1$" refers to the fact that these associations were encoded in the previous quarter.

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, $cue_{j,t}$ (and $cue_{j,t,q-1}$) 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)). However, while helpful, this does not address cross-industry relationships or 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 et al. (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.

Given these challenges, in further tests I strengthen the identification by exploiting only variation that results from firms announcing on the same day vs. with a gap of one day. The identifying assumption in these tests is that amongst a set of firms announcing close in time to each other (e.g., late in the quarter), firms that announce *on the same day* are not systematically different from firms that announce with *a gap of one day*. For instance, consider a firm with bad earnings news. In order to look relatively better, this firm might prefer to announce among firms with even worse news. As a result, the firm might strategically reschedule its earnings announcement to a time later in the quarter, in order to announce amongst the set of firms that generally announce bad news. But it would be hard for the firm to reschedule its announcement to the precise day on which firms with even worse news are announcing. The reason is that even for firms with bad news, it is hard to predict how bad the news of other firms is (Hartzmark and Shue (2018)). Therefore, it is plausible that amongst the set of firms that announce late in the quarter, firms that announce on the same day are not systematically different than firms that announce with a gap of one day. This is the variation that I am exploiting in these tests, which I implement by estimating the following regression:

$$\begin{aligned} return_{j,t} = \alpha + \beta_1 \cdot cue^{\Delta(-2)}_{j,t,q-1} + \beta_2 \cdot cue^{\Delta(-1)}_{j,t,q-1} + \beta_3 \cdot cue^{\Delta(0)}_{j,t,q-1} \\ + \beta_4 \cdot cue^{\Delta(+1)}_{j,t,q-1} + \beta_5 \cdot cue^{\Delta(+2)}_{j,t,q-1} + \gamma_t + u_{j,t} \end{aligned} \tag{3}$$

In this regression, $cue^{\Delta(k)}_{j,t,q-1}$ is a dummy variable that is equal to one if at least one firm that announced earnings k days after firm j in the previous quarter announces earnings on day t. For instance, $cue^{\Delta(+1)}_{j,t,q-1}$ is a dummy variable that is equal to one if at least one firm that announced earnings one day after firm j in the previous quarter announces earnings on day t. Similarly, $cue^{\Delta(-1)}_{j,t,q-1}$ is a dummy variable that is equal to one if at least one firm that announced earnings one day before firm j in the previous quarter announces earnings on day t. Finally, $cue^{\Delta(0)}_{j,t,q-1}$ is a dummy variable that is equal to one if at least one firm that announced earnings on the same day as firm j in the previous quarter announces earnings on day t (i.e., this dummy variable corresponds to the main dummy variable in Equation (2)). By comparing β_3 to β_2 and β_4, I can gauge whether a cue has a stronger effect if the memory association was encoded when both firms announced on the same day vs. with a gap of one day.

However, given the large set of factors that could potentially affect earnings announcement timing, there may be endogenous factors that affect even the decision to announce on the same day vs. with a gap of one day. To directly address these concerns, I repeat my analysis for the subsample of earnings announcement dates that were exogenously shifted by calendar rotations (Noh et al. (2021)). As an illustration of how this works, consider two firms, A and B, whose fiscal quarters end on June 30th.[4] Further, assume that both firms are "Pattern firms", i.e., they both follow a strict quarter-specific pattern in their earnings announcement timing and have not deviated from this pattern by even one day for at least three years. Specifically, Firm A always announces on the first Thursday in August and Firm

[4]This example is adapted from Noh et al. (2021).

Figure 1. Calendar Rotations: An Example

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

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. All my results hold in the subsample of Pattern firms.

3 Data and Summary Statistics

3.1 Full Sample

The full sample spans years 1995 to 2020 and consists of all firm-days without an own earnings announcement within +/- 3 trading days. Throughout the paper, the term "days" always refers to trading days. I use data from I/B/E/S to identify quarterly earnings announcement dates. Since my empirical strategy relies on identifying overlaps of earnings announcements, it is crucial that the earnings announcement dates are measured without error. I therefore follow DellaVigna and Pollet (2009) and use only data from 1995 onwards, since the accuracy of the earnings date is near perfect after December 1994. I match I/B/E/S data with CRSP for information on stock returns and with Compustat for firm-specific accounting information.

I calculate characteristic-adjusted returns as in Daniel et al. (1997) and Hartzmark and Shue (2018). 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 due to 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. In all my tests, the term "returns" refers to these characteristic-adjusted returns. I winsorize all variables at the 1st and the 99th percentile.

On some firm-days, cueing events occur, which I capture with a simple dummy variable. A cueing event for firm j on day t occurs if at least one firm that announced earnings on the same day as firm j in any of the previous four fiscal quarters announces earnings on day t. I require cueing firms to be in a different industry than firm j 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 two firms operate in dissimilar industries.[5] Further,

[5] 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, which are calibrated to be as granular as two-digit SIC codes. Specifically, 4.5% of randomly drawn firms are deemed to be peers according to this classification.

similar to Hartzmark and Shue (2018), I require cueing firms to have a market capitalization (measured on 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, but my results also hold for different cutoffs, which I show later in the paper.

Of the 16,474 distinct firms in my sample, 10,023 firms (60.84%) have at least one cueing event. On firm-days with a cueing event, I also calculate the earnings surprise of the cueing firm(s). In calculating the surprise, I follow Hartzmark and Shue (2018) and identify each analyst's most recent forecast in I/B/E/S, and then take the median of all analyst forecasts made between 2 and 15 days prior to the earnings announcement. Then, I calculate the difference between the actual earnings announced by the cueing firm and the median earnings forecast and scale this difference by the share price of the firm from three trading days prior to the announcement. If there are multiple cueing firms that announce earnings on day t, I calculate the equally-weighted as well as the value-weighted average surprise of the cueing firms, using each cueing firm's market capitalization three days prior to the announcement as value weights.

3.2 Pattern Firm Sample

I also construct a second sample, which I label "Pattern firm sample". This sample is constructed near-identically to the full sample and consists of all firm-days without an own earnings announcement within +/- 3 trading days. I also apply all the other filters described above, such as the exclusion of firm-pairs that operate in the same TNIC-2 industry. The key difference of the Pattern firm sample vis-à-vis the full sample is that cueing events are estimated using only the earnings announcement dates of Pattern firms.

I calculate cueing events for the Pattern firm sample using data provided by Noh et al. (2021).[6] Since this data is available for years 2005 – 2019, I restrict the Pattern firm sample to these years. I focus on the *threshold3* dataset 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 three years or more. I cross-check the earnings announcement dates in the *threshold3* dataset with those in my dataset and drop any discrepancies (less than

[6] Available here.

1% of observations). In the Pattern firm sample, a cueing event for Pattern firm j on day t occurs if at least one Pattern firm that announced earnings on the same day as Pattern firm j in any of the previous four fiscal quarters announces earnings on day t. That is, in this sample, only Pattern firms can be cueing or cued firms. Note, however, that this sample also contains non-Pattern firms, since the sample contains all firm-days without an earnings announcement within +/- 3 days. It is just that non-Pattern firms cannot be cueing or cued firms in this sample. As a result, of the 9,244 distinct firms in the Pattern firm sample, only 3,963 firms (42.87%) have at least one cueing event.

3.3 Summary Statistics

Table 1 presents summary statistics for the full sample in Panel A and for the Pattern firm sample in Panel B. In both samples, returns in excess of the characteristic-matched portfolio are slightly negative. This is because both samples systematically exclude days with an own earnings announcement, and therefore systematically exclude the earnings announcement premium (Frazzini and Lamont (2007); Barber et al. (2013)).

In the full sample (Panel A), there is a cueing event on 5.22% of days. This captures cues from memory associations that were encoded in any of the previous four fiscal quarters. However, in many tests, I focus on memory associations that were encoded in the previous quarter, since these are likely strongest and easiest to recall. The dummy variable Cue_{q-1} captures cueing events from such associations and shows that on 1.79% of days, there is a cueing event from a memory association that was encoded in the previous quarter. The subscript "$q-1$" refers to the fact that these associations were encoded in the previous quarter. While there is only one cue on the median day with a cueing event, there are also days with multiple cues. On such days, I calculate the earnings surprise of the cue as either the equally-weighted or value-weighted average surprise of the cueing firms. Earnings surprises are typically close to zero.

The Pattern firm sample (Panel B) looks very similar. The main difference is that cueing events are much less frequent, since they are calculated using only overlaps of earnings announcements of Pattern firms. In the Pattern firm sample, there is a cueing event on 0.66% of days and 0.21% of cueing events are from associations that were encoded in the previous quarter.

Table 1: Summary Statistics

Panel A: Full Sample	**Mean**	**p25**	**p50**	**p75**	**Std. Dev.**	**Min**	**Max**	**N**
Return (%)	-0.0168	-1.3331	-0.0801	1.1828	3.2857	-10.9855	12.8454	31,392,090
Cue (dummy)	0.0522	0.0000	0.0000	0.0000	0.2225	0.0000	1.0000	31,392,090
Cue$_{q-1}$ (dummy)	0.0179	0.0000	0.0000	0.0000	0.1325	0.0000	1.0000	31,392,090
Number of cues	2.2170	1	1	2	2.4032	1	35	1,639,021
Surprise of cue (EW)	0.0008	0.0001	0.0005	0.0011	0.0024	-0.0087	0.0156	1,639,021
Surprise of cue (VW)	0.0007	0.0001	0.0004	0.0011	0.0022	-0.0081	0.0142	1,639,021
Panel B: Pattern Firm Sample	**Mean**	**p25**	**p50**	**p75**	**Std. Dev.**	**Min**	**Max**	**N**
Return (%)	-0.0145	-1.1239	-0.0522	1.0218	2.7287	-10.9855	12.8454	15,773,961
Cue (dummy)	0.0066	0.0000	0.0000	0.0000	0.0809	0.0000	1.0000	15,773,961
Cue$_{q-1}$ (dummy)	0.0021	0.0000	0.0000	0.0000	0.0456	0.0000	1.0000	15,773,961
Number of cues	1.8256	1	1	2	1.7427	1	19	103,809
Surprise of cue (EW)	0.0007	0.0000	0.0005	0.0012	0.0017	-0.0052	0.0082	103,809
Surprise of cue (VW)	0.0007	0.0000	0.0005	0.0013	0.0016	-0.0050	0.0082	103,809

Notes: This table contains summary statistics of the two samples used in the empirical analysis. Panel A contains the full sample, which covers years 1995-2020. Panel B contains the Pattern firm sample, which covers years 2005-2019. Return is the raw return of a stock on day t minus the value-weighted return of a portfolio of stocks matched on size, book-to-market, and momentum. Cue is a dummy equal to one if at least one firm that announced earnings on the same day as firm j in any of the previous four fiscal quarters announces earnings on day t. Cue$_{q-1}$ is a similar dummy but focuses on firms that announced on the same day as firm j in the previous quarter. Number of cues is the number of cueing firms announcing on days with Cue equal to one. Surprise is the earnings surprise of the cueing firm(s) announcing on days with Cue equal to one. On days with multiple cues, I calculate both the equally-weighted (EW) and value-weighted (VW) earnings surprise.

4 Main Results

4.1 Baseline Results

In my first test, I estimate Equation (1) for the full sample and regress the return of firm j on day t on a dummy variable that is equal to one if there is a cueing event for firm j on day t. This dummy variable is equal to one if at least one firm that announced earnings on the same day as firm j in any of the previous four fiscal quarters 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 a daily abnormal return of 2.3 basis points, which corresponds to 5.8% in annual terms. In the second 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 et al. (2022)), but the effect is virtually identical when controlling for day fixed effects.

In the third column, I estimate Equation (2), which is similar but focuses on cueing events from associations that were encoded in the previous quarter. The idea behind this test is that recent associations might be easier to recall, potentially leading to stronger effects (although such associations are also captured by the catchall dummy in the first two columns).[7] Indeed, the coefficient is slightly larger and implies that such a cue leads to a daily abnormal return of 3 basis points, which annualizes to 7.5%. Controlling for day fixed effects in the fourth column also does not affect the size or significance of this coefficient. In the fifth and sixth columns, I break out the effect separately for days with one cue, two cues, and three or more cues. The effect tends to become stronger as the number of cues increases, which is plausible, since a larger number of cues is more likely to direct memory-induced attention to the cued firm.[8]

In Panel B, I replicate these results for the Pattern firm sample. In the first column of Panel B, I find that the effect of a cue is similar, and indeed slightly stronger, in this sample. A cueing event leads to a daily abnormal return of 3.8 basis points, which annualizes to 9.5%. When I test for the equality of coefficients across samples, I find that the difference is not significant ($p = 0.117$). In the second column, in which I account for day fixed effects, the coefficient slightly diminishes in magnitude and is even more similar to the corresponding coefficient in the full sample. In the third and fourth columns I again focus on cues from associations that were encoded in the previous quarter. The effect size is slightly larger for these cues, again suggesting that these associations might be stronger and easier to recall. Finally, in the fifth and sixth columns, I find that days with two cues tend to lead to stronger effects than days with a single cue. However, the coefficient on the dummy indicating three or more cues is small and insignificant in this sample, which might be driven by the fact that

[7]I explore the role of recency in recall in more detail in Table 4.

[8]In Appendix A, I show that I find similar results when I focus only on cueing events from the largest firm that had an overlapping earnings announcement with firm j in any of the previous four quarters.

Table 2: Baseline Results

Panel A: Full Sample

Dependent variable:	Return on day t (%)					
	(1)	(2)	(3)	(4)	(5)	(6)
Cue (dummy)	0.023***	0.021***				
	(0.005)	(0.006)				
Cue_{q-1} (dummy)			0.030***	0.029***		
			(0.007)	(0.008)		
Single Cue (dummy)					0.022***	0.021***
					(0.005)	(0.006)
Two Cues (dummy)					0.018**	0.014
					(0.007)	(0.009)
Three or More Cues (dummy)					0.028***	0.026**
					(0.009)	(0.011)
Day FE	no	yes	no	yes	no	yes
Observations	31,392,090	31,392,090	31,392,090	31,392,090	31,392,090	31,392,090
R-squared	0.000	0.003	0.000	0.003	0.000	0.003

Panel B: Pattern Firm Sample

Dependent variable:	Return on day t (%)					
	(1)	(2)	(3)	(4)	(5)	(6)
Cue (dummy)	0.038***	0.027**				
	(0.011)	(0.013)				
Cue_{q-1} (dummy)			0.062***	0.053***		
			(0.015)	(0.018)		
Single Cue (dummy)					0.035***	0.023*
					(0.011)	(0.013)
Two Cues (dummy)					0.077***	0.056**
					(0.023)	(0.026)
Three or More Cues (dummy)					0.014	0.018
					(0.024)	(0.028)
Day FE	no	yes	no	yes	no	yes
Observations	15,773,961	15,773,961	15,773,961	15,773,961	15,773,961	15,773,961
R-squared	0.000	0.004	0.000	0.004	0.000	0.004

Notes: This table presents results from regressions of the return on day t on different measures of memory cues. Return on day t is the raw return of a stock minus the value-weighted return of a portfolio of stocks matched on size, book-to-market, and momentum. In the first two columns, the independent variable is the dummy variable Cue, which is equal to one if at least one firm that announced earnings on the same day as firm j in any of the previous four fiscal quarters announces earnings on day t. In the next two columns, the independent variable is Cue_{q-1}, which is a similar dummy but focuses on firms that announced on the same day as firm j in the previous quarter. In the last two columns, the independent variables are dummies that indicate the number of cueing firms on days with Cue equal to one. Columns (2), (4), and (6) also include day fixed effects. Panel A shows results for the full sample and Panel B for the Pattern firm sample. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

it is very rare for this many cues to occur in the Pattern firm sample.

The Pattern firm sample provides a powerful setting in which to test robustness of my results, since the cues in this sample are driven by plausibly exogenous calendar rotations. However, one downside of this setting is that it favors identification at the expense of generalizability (Noh et al. (2021)). Given this, it is reassuring that the effect sizes are similar in both samples, suggesting that the documented effects hold for a broad set of firms. Overall, the results Table 2 support the hypothesis that internally-generated attention can lead to buying pressure in memory-associated stocks.

4.2 Contiguity

In this next set of tests, I estimate Equation (3) from Section (2). These tests serve a dual purpose: first, they tighten the identification; second, they allow me to explore an important property of memory, called the "law of contiguity". This law states that two items share a stronger association in memory if they were experienced closer in time together (Kahana (2012)). The intuition of contiguity, and how I test for it in my setting, can be illustrated best with the introductory example. Recall that last quarter, Coca-Cola and IBM announced earnings on the same day. Now, also assume that last quarter Apple announced one day after Coca-Cola. According to the law of contiguity, Coca-Cola and Apple should share a weaker association than Coca-Cola and IBM.

Therefore, in the tests that follow, I compare whether an earnings announcement of Coca-Cola on day t results in a weaker return response for Apple than for IBM, since the association between Coca-Cola and Apple is weaker than the association between Coca-Cola and IBM.

To implement these tests, I estimate Equation (3), which includes several dummy variables that capture different types of cues. For instance, $cue_{j,t,q-1}^{\Delta(-1)}$ is equal to one if at least one firm that announced earnings one day before firm j in the previous quarter announces earnings on day t. Thus, this dummy variable captures the return response of Apple if Coca-Cola announces earnings on day t. Similarly, $cue_{j,t,q-1}^{\Delta(0)}$ is equal to one if at least one firm that announced earnings on the same day as firm j in the previous quarter announces earnings on day t (i.e., this is the standard dummy that I use in most specifications). This

Figure 2. Contiguity
Return Response (%) on the Day of a Memory Cue

(a) Panel A: Full Sample

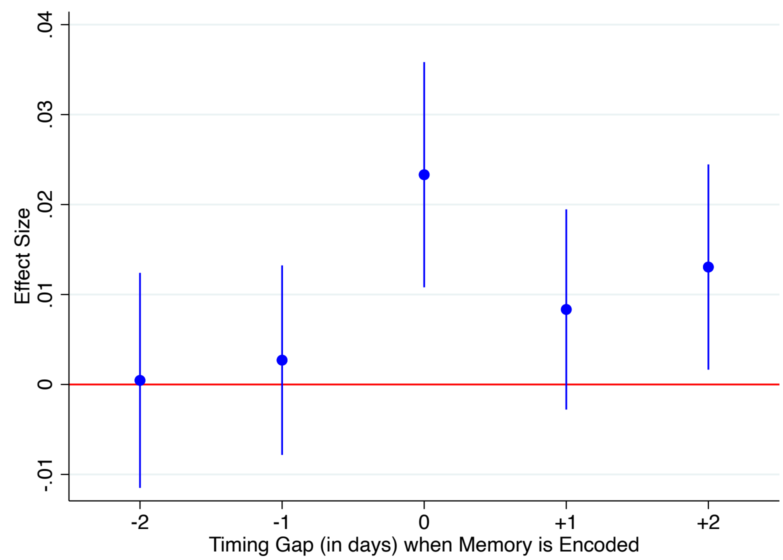

(b) Panel B: Pattern Firm Sample

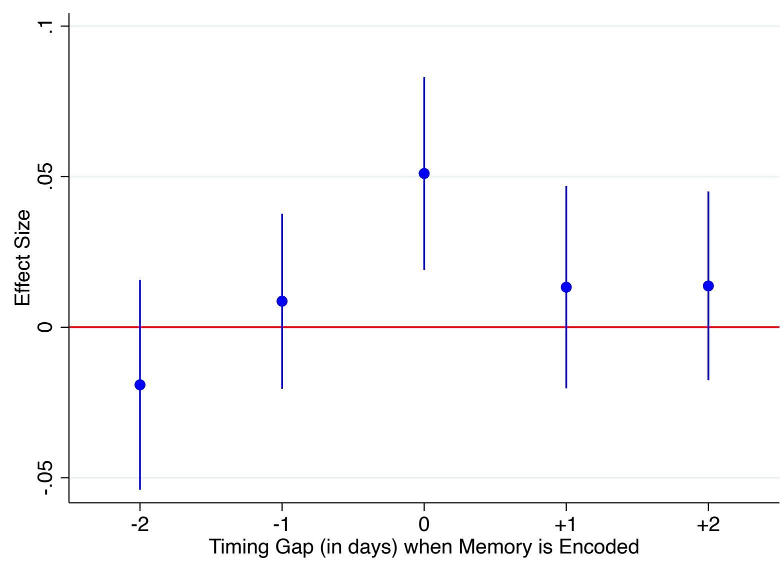

Notes: This figure plots the coefficient estimates from Table 3 along with 95% confidence intervals. Panel A shows results for the full sample and Panel B for the Pattern firm sample.

Table 3: Contiguity

Dependent variable:	Return on day t (%)	
Sample:	Full	Pattern Firm
	(1)	(2)
$\text{Cue}_{q-1}^{\Delta(-2)}$ (dummy)	0.000	-0.019
	(0.006)	(0.018)
$\text{Cue}_{q-1}^{\Delta(-1)}$ (dummy)	0.003	0.009
	(0.005)	(0.015)
$\text{Cue}_{q-1}^{\Delta(0)}$ (dummy)	0.023***	0.051***
	(0.006)	(0.016)
$\text{Cue}_{q-1}^{\Delta(+1)}$ (dummy)	0.008	0.013
	(0.006)	(0.017)
$\text{Cue}_{q-1}^{\Delta(+2)}$ (dummy)	0.013**	0.014
	(0.006)	(0.016)
Day FE	yes	yes
Observations	31,392,090	15,773,961
R-squared	0.003	0.004

Notes: This table presents results from regressions of the return on day t on dummy variables that capture different types of cues. The dummies are all of the type $cue_{j,t,q-1}^{\Delta(k)}$ and are equal to one if at least one firm that announced earnings k days after firm j in the previous quarter announces earnings on day t. Column (1) shows results for the full sample and column (2) for the Pattern firm sample. Both columns also include day fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

dummy captures the return response of IBM if Coca-Cola announces earnings on day t.

Table 3 presents the results from these regressions and Figure 2 plots the coefficients together with their corresponding 95% confidence intervals. The effect is strongest and most significant for associations that were encoded on the same day, and it drops of sharply for associations that were encoded with a gap of just one day. These results are fully consistent with the law of contiguity and show that memory associations are stronger if they were encoded in a more similar context.

In addition to providing support for an important property of memory, these tests also help alleviate concerns that my results are driven by similar firms announcing at similar times during an earnings season. As discussed in Section (2), these tests exploit variation from firms announcing on the same day vs. with a gap of one day in the previous quarter. Thus, the identifying assumption in these tests is that amongst a set of firms announcing close in time to each other (e.g., late in the quarter), firms that announce *on the same day* are not systematically different from firms that announce with *a gap of one day*. This assumption is likely satisfied, especially for the set of Pattern firms, whose earnings announcement dates are shifted by exogenous calendar rotations.

4.3 Recency

Perhaps the best-known property of human memory is that recent experiences are easier to recall than distant experiences, a property which I explore in the tests that follow. To do so, I augment Equation (2) with dummies for each of the previous four quarters, indicating cueing events from memory associations that were encoded in the respective quarter. These dummies allow me to back out the recency slope non-parametrically.

The first column in Panel A of Table 4 shows that when I estimate this specification for the full sample, the effect is strongest for the most recent quarter and fades away with further lags. This finding provides direct evidence that recently encoded memory associations have stronger effects. In Figure 3, I also show this effect graphically. Each dot in the figure is a coefficient estimate from the regression and the bars show corresponding 95% confidence intervals. The figure shows a clear systematic decay in the effect. Not only is this result consistent with a long psychology literature on memory, but also with the findings in Chang

Figure 3. Recency
Return Response (%) on the Day of a Memory Cue

(a) Panel A: Full Sample

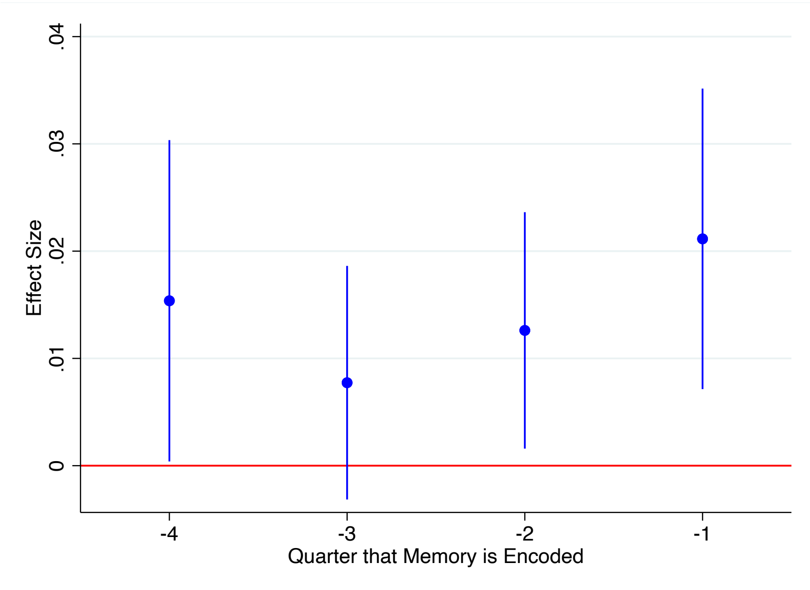

(b) Panel B: Pattern Firm Sample

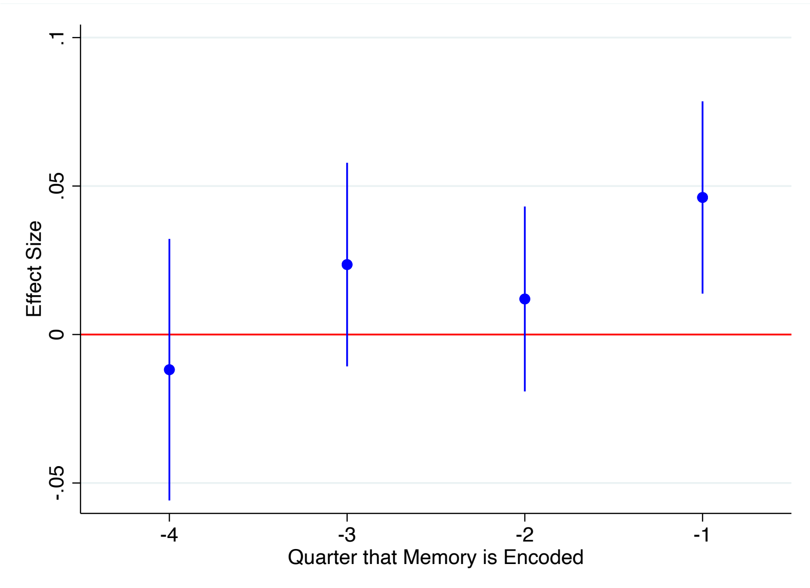

Notes: This figure plots the coefficient estimates from the first column of Table 4 along with 95% confidence intervals. Panel A shows results for the full sample and Panel B for the Pattern firm sample.

Table 4: Recency

Panel A: Full Sample

Dependent variable:	Return on day t (%)				
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} (dummy)	0.021***	0.029***			
	(0.007)	(0.008)			
Cue_{q-2} (dummy)	0.013**		0.022***		
	(0.006)		(0.007)		
Cue_{q-3} (dummy)	0.008			0.018***	
	(0.006)			(0.007)	
Cue_{q-4} (dummy)	0.015**				0.022***
	(0.008)				(0.008)
Day FE	yes	yes	yes	yes	yes
Observations	31,392,090	31,392,090	31,392,090	31,392,090	31,392,090
R-squared	0.003	0.003	0.003	0.003	0.003

Panel B: Pattern Firm Sample

Dependent variable:	Return on day t (%)				
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} (dummy)	0.046***	0.053***			
	(0.017)	(0.018)			
Cue_{q-2} (dummy)	0.012		0.025		
	(0.016)		(0.018)		
Cue_{q-3} (dummy)	0.024			0.032*	
	(0.017)			(0.019)	
Cue_{q-4} (dummy)	-0.012				-0.001
	(0.022)				(0.023)
Day FE	yes	yes	yes	yes	yes
Observations	15,773,961	15,773,961	15,773,961	15,773,961	15,773,961
R-squared	0.004	0.004	0.004	0.004	0.004

Notes: This table presents results from regressions of the return on day t on memory cues that were encoded in different quarters. The independent variables are dummy variables for each of the previous four fiscal quarters, and are equal to one if at least one firm that announced earnings on the same day as firm j in that quarter announces earnings on day t. Column (1) includes all dummies simultaneously and the remaining columns each include only one dummy at a time. All columns include day fixed effects. Panel A shows results for the full sample and Panel B for the Pattern firm sample. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

et al. (2017), which shows that investors tend to overweight the most recent quarter.

In the second through fifth column of Table 4, I also estimate separate regressions, each with a dummy variable for a different quarterly lag. These separate regressions further address the worry that my results are picking up fundamental relationships. A possible interpretation of the full regression presented in the first column is that each quarterly lag proxies for fundamentally related firms, but the prior quarter is the best proxy. Under this interpretation, when included separately, each dummy should be a strong predictor on its own, getting a bit noisier with further lags. In contrast, under a memory channel, the coefficients from the separate regressions should look similar to those from the full regression and also show a systematic decay. This is precisely what I find.

In Panel B, I replicate these results in the Pattern firm sample. Here, too, I find that the coefficient on the previous quarter is strongest and that the effect becomes weaker as the lag increases. The drop off from the first to the second lag is even sharper in this sample than in the full sample. In sum, the results in Table 4 not only validate a key prediction of memory theory, but also help attenuate concerns that my results are driven by attention to fundamentally related firms.

4.4 Interference

I next turn to tests that explore whether interference dampens the documented effect. The intuition is straightforward. If a memory association between Coca-Cola and IBM 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 not only encode associations between Coca-Cola and IBM, but also associations between Coca-Cola and these other firms. As a result, when cued with an earnings announcement by Coca-Cola on day t, investors might not recall IBM, but one of these other firms instead. Put differently, the memories of these other firms interfere with the recall of IBM. Thus, I hypothesize that the return response of IBM is weaker if interference is higher.

Applying this intuition to my setting, I classify associations as having "low interference" if they were encoded on days that the number of other firms announcing was below the median. Conversely, I classify associations as having "high interference" if the number of

Figure 4. Interference
Return Response (%) on the Day of a Memory Cue

(a) Panel A: Full Sample

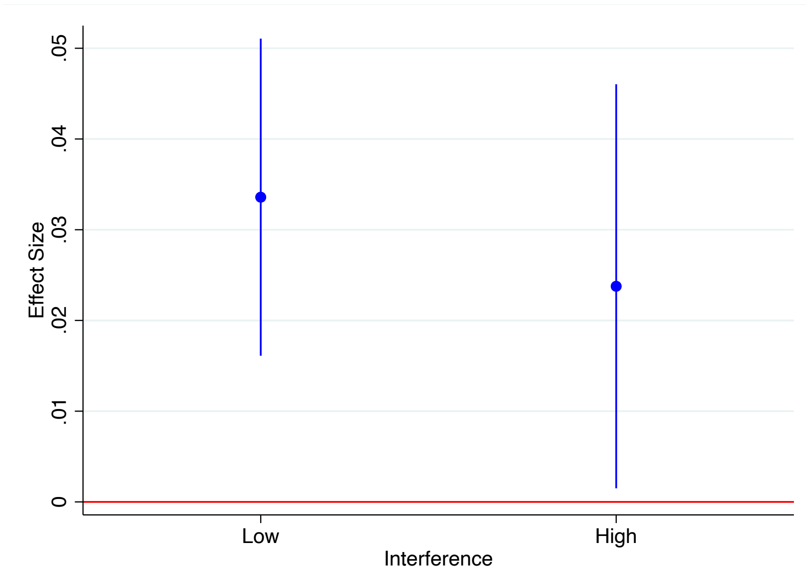

(b) Panel B: Pattern Firm Sample

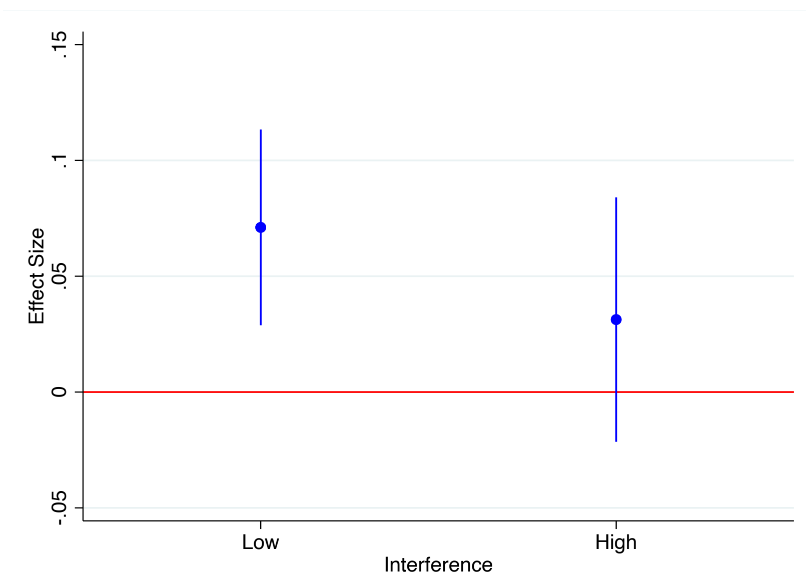

Notes: This figure plots the coefficient estimates from Table 5 along with 95% confidence intervals. Panel A shows results for the full sample and Panel B for the Pattern firm sample.

Table 5: Interference

Dependent variable:	Return on day t (%)	
Sample:	Full	Pattern Firm
	(1)	(2)
Cue_{q-1} x Low Interference	0.034***	0.071***
	(0.009)	(0.022)
Cue_{q-1} x High Interference	0.024**	0.031
	(0.011)	(0.027)
Day FE	yes	yes
Observations	31,392,090	15,773,961
R-squared	0.003	0.004

Notes: This table presents results from regressions of the return on day t on memory cues that were encoded with low and high interference. Cue_{q-1} x Low Interference is equal to one if at least one firm that announced earnings on the same day as firm j in the previous quarter announces earnings on day t, and if the number of firms announcing on that day in the previous quarter was below the median. Cue_{q-1} x High Interference is defined equivalently, except that the number of firms announcing on that day in the previous quarter was above the median. Column (1) shows results for the full sample and column (2) for the Pattern firm sample. Both columns also include day fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

firms announcing on the same day was above the median. I calculate the median cutoff for each year separately to account for time-varying trends in the number of firms announcing.

In Table 5, I show that the effect is stronger for cues from associations with low interference. In the full sample, the effect is almost 50% larger if interference is low. In the Pattern firm sample, the effect is more than double if interference is low. Figure 4 also visualizes this result as a coefficient plot with 95% confidence intervals.

5 Further Results and Robustness

5.1 Reversals

Since memory-induced buying pressure carries no new information, prices should eventually revert to their fundamental values. In contrast, if my results were driven by fundamental relationships, there should be no systematic reversal. To test for these possibilities, I examine the effect for return windows that extend beyond day t in Table 6. When testing for reversal, it is important to ensure that no new cues occur in the return window that is being analyzed, since these new cues would confound the effect of the reversal. Therefore, in these tests I exclude all observations for which there is an additional cue on day $t+1$ or $t+2$. Further, since the return windows extend beyond day t, it is important to ensure that the returns in these windows are not affected by (the anticipation of) an own earnings announcement. To err on the conservative side in these tests, I extend my baseline window and exclude all days within $[t-10, t+10]$ of an own earnings announcement.

Panel A presents the results for the full sample. In the first column of Table 6, I reproduce the main effect for this restricted sample. The coefficient on Cue_{q-1} is very similar to the baseline effect in Table 2. The second column shows that there is full reversal of the effect on day $t+1$, and the remaining columns show that there is no significant effect for return windows that extend beyond day t. Panel B presents similar results for the Pattern firm sample. In this sample, there is partial reversal on day $t+1$, but the effects are not as clear cut. However, even in this sample, I cannot detect a significant effect for return windows that extend beyond day t. Taken together, these results suggest that the mispricing due to

Table 6: Reversals

Panel A: Full Sample

Return window:	$[t]$	$[t+1]$	$[t+2]$	$[t,\, t+1]$	$[t,\, t+2]$
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} (dummy)	0.035**	-0.037***	-0.001	-0.001	-0.006
	(0.015)	(0.014)	(0.015)	(0.021)	(0.025)
Day FE	yes	yes	yes	yes	yes
Observations	28,180,614	28,163,550	28,146,559	28,163,550	28,147,444
R-squared	0.004	0.004	0.004	0.004	0.004

Panel B: Pattern Firm Sample

Return window:	$[t]$	$[t+1]$	$[t+2]$	$[t,\, t+1]$	$[t,\, t+2]$
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} (dummy)	0.083***	-0.046	0.044	0.031	0.075
	(0.032)	(0.030)	(0.028)	(0.045)	(0.053)
Day FE	yes	yes	yes	yes	yes
Observations	14,215,217	14,209,807	14,204,430	14,209,807	14,204,847
R-squared	0.005	0.005	0.005	0.005	0.004

Notes: This table presents results from regressions of the return for various return windows on Cue_{q-1}. The return windows are indicated in the column headers. The sample excludes firm-days with an additional cue on day $t+1$ or $t+2$. All firm-days within $[t-10,\, t+10]$ of an own earnings announcement of firm j are excluded from the sample. All columns include day fixed effects. Panel A shows results for the full sample and Panel B for the Pattern firm sample. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

memory-induced buying pressure is quickly corrected by the market.

5.2 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 7, I regress the return of firm j on day t on the earnings surprise of the cueing firm. I also test whether the effect varies along the distribution of the cueing firms' earnings surprise, to account for potential non-linear relationships. The sample in these tests is restricted to days with a cueing event, since these are the only days on which an earnings surprise of cueing firms can be calculated. On days with multiple cues, I calculate either the equally-weighted average (first and second column) or the value-weighted average (third and fourth column) of the cueing firms' earnings surprise. Panel A presents the results of these tests for the full sample, and Panel B for the Pattern firm sample.

Table 7: Surprise of the Cue

Panel A: Full Sample

Dependent variable:	Return on day t (%)			
Surprise:	EW		VW	
	(1)	(2)	(3)	(4)
Surprise	-2.304		-2.715	
	(2.101)		(2.105)	
Surprise Quintile 2 (dummy)		0.002		0.015
		(0.013)		(0.014)
Surprise Quintile 3 (dummy)		0.030**		0.030**
		(0.014)		(0.014)
Surprise Quintile 4 (dummy)		0.023		0.015
		(0.014)		(0.014)
Surprise Quintile 5 (dummy)		0.024		0.015
		(0.016)		(0.015)
Day FE	yes	yes	yes	yes
Observations	1,639,021	1,639,021	1,639,021	1,639,021
R-squared	0.006	0.006	0.006	0.006

Panel B: Pattern Firm Sample

Dependent variable:	Return on day t (%)			
Surprise:	EW		VW	
	(1)	(2)	(3)	(4)
Surprise	-10.786		-10.006	
	(10.520)		(9.539)	
Surprise Quintile 2 (dummy)		-0.010		-0.001
		(0.039)		(0.037)
Surprise Quintile 3 (dummy)		0.022		0.032
		(0.041)		(0.037)
Surprise Quintile 4 (dummy)		-0.009		-0.006
		(0.042)		(0.037)
Surprise Quintile 5 (dummy)		-0.021		-0.061
		(0.047)		(0.047)
Day FE	yes	yes	yes	yes
Observations	103,784	103,784	103,784	103,784
R-squared	0.021	0.021	0.021	0.021

Notes: This table presents results from regressions of the return on day t on the earnings surprise of the cueing firm(s). The sample is restricted to days with a cue. 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 on day t, columns (1) and (2) use the equally-weighted average surprise, and columns (3) and (4) use the value-weighted average surprise. Columns (1) and (3) include the surprise directly as an independent variable, while columns (2) and (4) include dummy variables that indicate the second through fifth quintile of the surprise distribution. Panel A presents results for the full sample and Panel B for the Pattern firm sample. All columns include day fixed effects. These fixed effects can result in singleton observations, which are dropped during the estimation (specifically, 25 observations in Panel B). Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

I find that the surprise of cueing firms does not have predictive power for the return response of the cued firm, neither in the full sample nor in the Pattern firm sample. While at first blush the coefficients in the first and third column 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 of Panel A 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 half a basis point. The small economic magnitudes of the earnings surprise coefficients are also apparent in the non-parametric estimations in the second and fourth column. 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, regardless of the sign or magnitude of the earnings surprise.

5.3 Large vs. Small Firms

I also explore whether the documented effect is stronger for large or small firms. 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 firms, both in the full sample and in the Pattern firm sample. Table 8 presents the results. I find that the effect is driven mostly by small firms. However, particularly in the Pattern firm sample, the effect can also be documented for large firms.

Table 8: Large vs. Small Firms

Dependent variable:	Return on day t (%)			
Sample:	Full		Pattern firm	
Firm size:	Small (1)	Large (2)	Small (3)	Large (4)
Cue_{q-1} (dummy)	0.054*** (0.014)	0.010* (0.006)	0.106*** (0.039)	0.028** (0.014)
Day FE	yes	yes	yes	yes
Observations	15,696,043	15,696,047	7,886,980	7,886,981
R-squared	0.008	0.002	0.010	0.001

Notes: This table splits the sample along the median market capitalization of firm j from $t-3$ and presents results from regressions of the return on day t on Cue_{q-1}. Columns (1) and (2) show results for the full sample and columns (3) and (4) for the Pattern firm sample. All columns include day fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

5.4 Trading Strategy

Since earnings announcements are usually scheduled at least a week ahead of time (Boulland and Dessaint (2017)), it is possible to construct a trading strategy that takes advantage of the buying pressure caused by memory-induced attention. This intentionally simple trading strategy is a daily long-short strategy that is held on day t. The strategy goes long stocks for which Cue_{q-1} is equal to one and it goes short stocks for which Cue_{q-1} is equal to zero. 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 (market capitalization from $t-3$ is below the median) to form these portfolios, since these firms drive most of the effect (see Table 8).

To account for the potential role of risk factors, I regress the time series of daily 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 that the strategy yields a daily alpha of 8.7 basis points, which is significant at the 5% level. This daily alpha would correspond to an annual abnormal return of about 22% if the trading strategy could be implemented every trading day. However, 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 the strategy. In my sample, the strategy can be implemented on an average of 110 trading days per year. Thus, the strategy generates an annual abnormal return of about 9.5%.

Whether this strategy continues to yield positive abnormal returns after accounting for transaction costs depends on the size and nature of these costs, which likely differ across investors. However, the purpose of illustrating this trading strategy is not to identify the highest possible alpha, but rather to show that the main result holds in a different specification, with a different risk adjustment. To this end, Table 9 shows that while there are loadings on the factors, these loadings are economically small and do not wash out the positive and significant alpha. Thus, a risk-based explanation is unlikely to explain the results. Overall, these results highlight the robustness of my findings using calendar-time asset pricing methods.

Table 9: Trading Strategy

Dependent variable:	Return on day t (%)
	(1)
Alpha [%]	0.087**
	(0.039)
Mkt	0.124***
	(0.034)
SMB	0.323***
	(0.060)
HML	0.061
	(0.060)
Momentum	-0.115***
	(0.043)
ST Reversal	-0.003
	(0.044)
Observations	2,753
R-squared	0.023

Notes: This table presents results from a regression of daily trading strategy returns on the daily market, size, value, momentum, and short-term reversal factors, which are sourced from the Kenneth French Data Library. The daily trading strategy returns are the daily returns of a long-short strategy. The long leg of the strategy consists of a value-weighted portfolio of stocks with Cue_{q-1} equal to one, and the short leg of the strategy consists of a value-weighted portfolio of stocks with Cue_{q-1} equal to zero. 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. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

5.5 Robustness

In my tests so far, I focus on large and salient cues by requiring each cueing firm to have a market capitalization that is above the NYSE's 90th percentile of market capitalization. Here, I show that these results are robust to using the 85th and 95th percentile as cutoffs. In Table 10, I rerun my main regressions using these different cutoffs. Panel A shows the results for the full sample while Panel B shows the results for the Pattern firm sample. Overall, my results are robust to using different size cutoffs.

In my main tests, I ensure that firm j does not have an own earnings announcement in the window $[t-3, t+3]$ to avoid confounding effects from a firm's own earnings announcement. However, this window is somewhat ad hoc. Therefore, in Table 10, I explore the robustness of my results for different windows. Specifically, I replicate my main results using samples in which I ensure that there is no own earnings announcement in $[t-1, t+1]$ and in $[t-5, t+5]$. Panel A presents the results for the full sample and Panel B presents results for the Pattern firm sample. For both alternative windows, I find very similar results as with the baseline window.

An open question is whether my results reflect investors' memory, or whether they reflect analysts' memory. For instance, assume that an analyst follows Coca-Cola and IBM, and has historically experienced both firms. If Coca-Cola announces earnings on day t, the analyst will likely recall IBM, and might decide to update his forecast for IBM. In this scenario, the stock price reaction of IBM on day t might be a side-effect of the forecast update of the analyst. In order to rule out this scenario, I exclude all firm-pairs that have any overlap in analyst following in the days $t-45$ to t. Using this restricted sample, I rerun my main regressions and present the results in the fifth column of Table 10. The coefficient estimates are virtually unchanged relative to the main results in Table 2 and support the hypothesis that it is investors' memory, rather than analysts' memory, that is driving my results.

5.6 Alternative Explanations

My results strongly 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

Table 10: Robustness

Panel A: Full Sample

Dep. var.:	Return on day t (%)				
Sample:	Cue above 85th pctile	Cue above 95th pctile	No own EA in $[t-1,\, t+1]$	No own EA in $[t-5,\, t+5]$	No Analyst Overlap
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1}	0.026***	0.027***	0.028***	0.029***	0.029***
	(0.008)	(0.009)	(0.008)	(0.009)	(0.008)
Day FE	yes	yes	yes	yes	yes
Observations	31,392,090	31,392,090	32,000,990	30,781,708	31,392,090
R-squared	0.003	0.003	0.003	0.003	0.003

Panel B: Pattern Firm Sample

Dep. var.:	Return on day t (%)				
Sample:	Cue above 85th pctile	Cue above 95th pctile	No own EA in $[t-1,\, t+1]$	No own EA in $[t-5,\, t+5]$	No Analyst Overlap
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1}	0.053***	0.046**	0.047***	0.054***	0.053***
	(0.016)	(0.022)	(0.017)	(0.019)	(0.018)
Day FE	yes	yes	yes	yes	yes
Observations	15,773,961	15,773,961	16,111,016	15,435,974	15,773,961
R-squared	0.004	0.004	0.004	0.004	0.004

Notes: This table presents results from regressions of the return on day t on Cue_{q-1}. Column (1) only uses cueing firms that have a market capitalization (measured on day $t-3$) above the NYSE's 85th percentile of market capitalization in that month, and column (2) uses only cueing firms that have a market capitalization (measured on day $t-3$) above the NYSE's 95th percentile of market capitalization in that month. Column (3) drops all firm-days within $[t-1,\, t+1]$ of an own earnings announcement of firm j, and column (4) drops all firm-days within $[t-5,\, t+5]$ of an own earnings announcement. Column (5) only includes cues from firms that have no overlap in analyst following with firm j in $[t-45,\, t]$. All columns include day fixed effects. Panel A shows results for the full sample and Panel B for the Pattern firm sample. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

associative memory theory. Specifically, I find that if two firms historically announced with a timing gap between their earnings announcements, they share a weaker memory association. Further, I find that associations that were encoded in recent quarters have a stronger effect than associations that were encoded in more distant quarters. And finally, 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. However, my tests are designed to rule out these possibilities. First, I find that there is a sharp discontinuity in the effect if two firms announced with a gap of just one day, making it unlikely that the effect is driven by similar firms announcing during similar times in an earnings season. Second, I find that the earnings surprise of the cueing firm has no predictive power for the return response of the cued firm. And third, I replicate all my results in the sample of Pattern firms, whose earnings announcement dates are exogenously shifted by calendar rotations.

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 Coca-Cola and IBM announced earnings on the same day last quarter, but this quarter they do not. This alternative explanation posits that when Coca-Cola announces earnings this quarter, investors rediscover IBM, for example by reading an archived newspaper article from last quarter, in which both Coca-Cola and IBM are covered. While this explanation might plausibly explain the baseline results, it must also explain the results from the mechanism tests. Specifically, it must explain:

1. Why there is such a sharp discontinuity in the effect if two firms announced with a gap of just one day. If investors access historically archived newspaper articles, they might plausibly discover firms that announced one day before or after the cueing firm, since these firms could be covered in the same articles as well.

2. Why interference dampens the effect. If anything, we might expect more newspaper articles to be written on days with many earnings announcements. Thus, investors should be more likely to discover a firm if it announced on a busy day.

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.

6 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. Memories that were encoded on the same day have a much stronger effect than memories that were encoded with a time gap. Further, I find that recently encoded memories have the strongest effect, and that the effect fades away with time. Finally, the documented effect is stronger if there are fewer interfering events during the encoding of a memory association.

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 et al. (2022); Gödker et al. (2022)). In contrast, my setting allows me to show that the constraints of human memory can aggregate and affect asset prices. 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 that is 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. Future research may flesh out these internal sources in more detail and test their effects on financial markets.

Appendix

A Robustness using only the Largest Cue

Table A.1: Baseline Results using only Largest Cue

Dependent variable:	Return on day t (%)			
Sample:	Full		Pattern firm	
	(1)	(2)	(3)	(4)
Cue_{q-1} of largest cue (dummy)	0.024***	0.021**	0.061***	0.053**
	(0.007)	(0.008)	(0.018)	(0.021)
Day FE	no	yes	no	yes
Observations	31,392,090	31,392,090	15,773,961	15,773,961
R-squared	0.000	0.003	0.000	0.004

Notes: This table presents results from regressions of the return on day t on Cue_{q-1} of the largest cue. This dummy is equal to one if the largest firm that announced earnings on the same day as firm j in the previous fiscal quarter announces earnings on day t. Columns (1) and (2) show results for the full sample and columns (3) and (4) for the Pattern firm sample. Columns (2) and (4) also include day fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.2: Recency using only Largest Cue

Panel A: Full Sample

Dependent variable:	Return on day t (%)				
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} of largest cue (dummy)	0.018**	0.021**			
	(0.008)	(0.008)			
Cue_{q-2} of largest cue (dummy)	0.014*		0.017**		
	(0.007)		(0.007)		
Cue_{q-3} of largest cue (dummy)	0.008			0.012*	
	(0.007)			(0.007)	
Cue_{q-4} of largest cue (dummy)	0.018*				0.020**
	(0.009)				(0.009)
Day FE	yes	yes	yes	yes	yes
Observations	31,392,090	31,392,090	31,392,090	31,392,090	31,392,090
R-squared	0.003	0.003	0.003	0.003	0.003

Panel B: Pattern Firm Sample

Dependent variable:	Return on day t (%)				
	(1)	(2)	(3)	(4)	(5)
Cue_{q-1} of largest cue (dummy)	0.048**	0.053**			
	(0.021)	(0.021)			
Cue_{q-2} of largest cue (dummy)	0.010		0.020		
	(0.017)		(0.018)		
Cue_{q-3} of largest cue (dummy)	0.035*			0.039**	
	(0.019)			(0.020)	
Cue_{q-4} of largest cue (dummy)	-0.019				-0.013
	(0.024)				(0.024)
Day FE	yes	yes	yes	yes	yes
Observations	15,773,961	15,773,961	15,773,961	15,773,961	15,773,961
R-squared	0.004	0.004	0.004	0.004	0.004

Notes: This table presents results from regressions of the return on day t on memory cues that were encoded in different quarters. The independent variables are dummy variables for each of the previous four fiscal quarters, and are equal to one if the largest firm that announced earnings on the same day as firm j in that quarter announces earnings on day t. Column (1) includes all dummies simultaneously and the remaining columns each include only one dummy at a time. All columns include day fixed effects. Panel A shows results for the full sample and Panel B for the Pattern firm sample. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.3: Interference using only Largest Cue

Dependent variable:	Return on day t (%)	
Sample:	Full	Pattern Firm
	(1)	(2)
Cue_{q-1} x Low Interference	0.030***	0.078***
	(0.010)	(0.024)
Cue_{q-1} x High Interference	0.011	0.025
	(0.013)	(0.034)
Day FE	yes	yes
Observations	31,392,090	15,773,961
R-squared	0.003	0.004

Notes: This table presents results from regressions of the return on day t on memory cues that were encoded with low and high interference. Cue_{q-1} x Low Interference is equal to one if the largest firm that announced earnings on the same day as firm j in the previous quarter announces earnings on day t, and if the number of firms announcing on that day in the previous quarter was below the median. Cue_{q-1} x High Interference is defined equivalently, except that the number of firms announcing on that day in the previous quarter was above the median. Column (1) shows results for the full sample and column (2) for the Pattern firm sample. Both columns also include day fixed effects. Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

Table A.4: Surprise of Largest Cue

Dependent variable:	Return on day t (%)			
Sample:	Full		Pattern Firm	
	(1)	(2)	(3)	(4)
Surprise	0.678		8.294	
	(3.157)		(7.677)	
Surprise Quintile 2 (dummy)		-0.004		-0.009
		(0.020)		(0.041)
Surprise Quintile 3 (dummy)		0.002		0.079*
		(0.019)		(0.042)
Surprise Quintile 4 (dummy)		-0.015		0.050
		(0.019)		(0.045)
Surprise Quintile 5 (dummy)		0.002		0.005
		(0.019)		(0.046)
Day FE	yes	yes	yes	yes
Observations	821,168	821,168	77,970	77,970
R-squared	0.008	0.008	0.024	0.024

Notes: This table presents results from regressions of the return on day t on the earnings surprise of the largest cueing firm. The sample is restricted to days with a cue from the largest firm. 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. Columns (1) and (3) include the surprise directly as an independent variable, while columns (2) and (4) include dummy variables that indicate the second through fifth quintile of the surprise distribution. Columns (1) and (2) present results for the full sample and columns (3) and (4) for the Pattern firm sample. All columns include day fixed effects. These fixed effects can result in singleton observations, which are dropped during the estimation (specifically, 1 observation in Panel A and 25 observations in Panel B). Standard errors are clustered by firm and trading day and are displayed in parentheses below the coefficients. *, **, and ***, indicate statistical significance at the 10%, 5%, and 1% level, respectively.

REFERENCES

An, Li, Xin Chen, Zhengwei Wang, and Jianfeng Yu, 2022, Attention spillover in asset pricing, Working Paper.

Bagnoli, Mark, William Kross, and Susan G Watts, 2002, The information in management's expected earnings report date: A day late, a penny short, *Journal of Accounting Research* 40, 1275–1296.

Barber, Brad M, Emmanuel T De George, Reuven Lehavy, and Brett Trueman, 2013, The earnings announcement premium around the globe, *Journal of Financial Economics* 108, 118–138.

Barber, Brad M, Xing Huang, Terrance Odean, and Christopher Schwarz, 2021, Attention induced trading and returns: Evidence from robinhood users, *Journal of Finance* forthcoming.

Barber, Brad M, Shengle Lin, and Terrance Odean, 2019, Mediating investor attention, Working paper.

Barber, Brad M, and Terrance 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.

Bodoh-Creed, Aaron L, 2020, Mood, memory, and the evaluation of asset prices, *Review of Finance* 24, 227–262.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer, 2022, Memory and probability, *Quarterly Journal of Economics* conditionally accepted.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *Quarterly Journal of Economics* 135.

Boulland, Romain, and Olivier Dessaint, 2017, Announcing the announcement, *Journal of Banking & Finance* 82, 59–79.

Chang, Tom Y, Samuel M Hartzmark, David H Solomon, and Eugene F Soltes, 2017, Being surprised by the unsurprising: Earnings seasonality and stock returns, *Review of Financial Studies* 30, 281–323.

Charles, Constantin, 2022, Memory and trading, Working paper.

Chen, Ting, Zhenyu Gao, Jibao He, Wenxi Jiang, and Wei Xiong, 2019, Daily price limits and destructive market behavior, *Journal of Econometrics* 208, 249–264.

Chen, Yixin, Randolph B Cohen, and Zixuan Kevin Wang, 2022, Famous firms, earnings clusters, and the stock market, Working Paper.

Colonnelli, Emanuele, Niels J Gormsen, and Timothy McQuade, 2022, Selfish corporations, Working paper.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, 2011, In search of attention, *Journal of Finance* 66, 1461–1499.

Da Silveira, Rava Azeredo, Yeji Sung, and Michael Woodford, 2020, Optimally imprecise memory and biased forecasts, Working paper.

Daniel, Kent, Mark Grinblatt, Sheridan Titman, and Russ Wermers, 1997, Measuring mutual fund performance with characteristic-based benchmarks, *Journal of Finance* 52, 1035–1058.

DellaVigna, Stefano, and Joshua M Pollet, 2009, Investor inattention and friday earnings announcements, *Journal of Finance* 64, 709–749.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, 2022, Associative memory and belief formation, Working paper.

Frazzini, Andrea, and Owen Lamont, 2007, The earnings announcement premium and trading volume, Working paper.

Frydman, Cary, and Baolian Wang, 2020, The impact of salience on investor behavior: Evidence from a natural experiment, *Journal of Finance* 75, 229–276.

Gilboa, Itzhak, and David Schmeidler, 1995, Case-based decision theory, *Quarterly Journal of Economics* 110, 605–639.

Gödker, Katrin, Peiran Jiao, and Paul Smeets, 2022, Investor memory, Working paper.

Goetzmann, William N, Akiko Watanabe, and Masahiro Watanabe, 2022, Evidence on retrieved context: How history matters, Working Paper.

Hartzmark, Samuel M, 2015, The worst, the best, ignoring all the rest: The rank effect and trading behavior, *Review of Financial Studies* 28, 1024–1059.

Hartzmark, Samuel M, Samuel D Hirshman, and Alex Imas, 2021, Ownership, learning, and beliefs, *Quarterly Journal of Economics* 136, 1665–1717.

Hartzmark, Samuel M, and Kelly Shue, 2018, A tough act to follow: Contrast effects in financial markets, *Journal of Finance* 73, 1567–1613.

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

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh, 2009, Driven to distraction: Extraneous events and underreaction to earnings news, *Journal of Finance* 64, 2289–2325.

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

Hoberg, Gerard, and Gordon Phillips, 2010, Product market synergies and competition in mergers and acquisitions: A text-based analysis, *Review of Financial Studies* 23, 3773–3811.

Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.

Jiang, Lei, Jinyu Liu, Lin Peng, and Baolian Wang, 2022, Investor attention and asset pricing anomalies, *Review of Finance* 26, 563–593.

Johnson, Travis L, and Eric C So, 2018, Time will tell: Information in the timing of scheduled earnings news, *Journal of Financial and Quantitative Analysis* 53, 2431–2464.

Kahana, Michael J, 2012, *Foundations of human memory* (OUP USA).

Lou, Dong, 2014, Attracting investor attention through advertising, *Review of Financial Studies* 27, 1797–1829.

Mullainathan, Sendhil, 2002, A memory-based model of bounded rationality, *Quarterly Journal of Economics* 117, 735–774.

Murdock Jr, Bennet B, 1962, The serial position effect of free recall, *Journal of Experimental Psychology* 64, 482.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *Review of Financial Studies* 35, 2190–2245.

Noh, Suzie, Eric C So, and Rodrigo S Verdi, 2021, Calendar rotations: A new approach for studying the impact of timing using earnings announcements, *Journal of Financial Economics* 140, 865–893.

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

Schmidt, Daniel, 2019, Distracted institutional investors, *Journal of Financial and Quantitative Analysis* 54, 2453–2491.

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

Thomas, Jacob, and Frank Zhang, 2008, Overreaction to intra-industry information transfers?, *Journal of Accounting Research* 46, 909–940.

Wachter, Jessica A, and Michael J Kahana, 2021, A retrieved-context theory of financial decisions, Working paper.