NBER WORKING PAPER SERIES

CORRELATION NEGLECT IN ASSET PRICES

Hongye Guo
Jessica Wachter

Working Paper 35753
http://www.nber.org/papers/w35753

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

The authors do not have conflicts to disclose. We thank Zahi Ben-David (discussant), Stefano Cassella (discussant), Zhi Da, Anastassia Fedyk (discussant), Sam Hartzmark (discussant), Sebastian Hillenbrand (discussant), Kristy Jansen, Peter Kelly, Wenhao Li, Zeming Li, Dong Lou, Jinfei Sheng (discussant), Ross Valkanov (discussant), Renxuan Wang (discussant), Florian Zimmermann, and seminar participants at Adam Smith Workshop, AFA, CEIBS-APS CAC Academic Conference, CICF, Cognitive Foundations of Finance Conference, EFA, LSE, NYU Stern, SFS Cavalcade, UBC Winter Conference, the University of Bristol, and University of Southern California for helpful comments. We are grateful to the Q-group for awarding this paper the 2025 Jack Treynor Prize. We gratefully acknowledge research support from National Institutes of Health Grant MH55687. Hongye Guo gratefully acknowledges research support from Hong Kong RGC GRF grant (#17507525). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Hongye Guo and Jessica Wachter. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Correlation Neglect in Asset Prices
Hongye Guo and Jessica Wachter
NBER Working Paper No. 35753
September 2026
JEL No. G12

ABSTRACT

We document a strong negative serial correlation between stock market returns in i) the second month of a quarter and ii) the first month of the subsequent quarter. This pattern arises as investors fail to fully recognize the predictably repetitive aggregate earnings "news" conveyed by the late earnings announcements in the second month of a quarter. The neglected correlation leads to overreaction, which is in turn corrected when earnings of a new quarter are announced. Consistent with this hypothesis, returns in the second month of a quarter positively correlate with those of the preceding month, and the return pattern amplifies when earnings announced in these two months are more similar. Survey data and industry-level returns lend further support to this mechanism. These results provide evidence of correlation neglect even among sophisticated, financially incentivized decision-makers, underscoring its importance as a behavioral phenomenon.

Hongye Guo
The University of Hong Kong (HKU)
hoguo@hku.hk

Jessica Wachter
University of Pennsylvania
The Wharton School
Department of Finance
and NBER
jwachter@wharton.upenn.edu

1 Introduction

Predicting stock market returns based on past returns has posed a decades-long challenge. Early research by Kendall and Hill (1953) and Fama (1965) revealed minimal serial correlation in stock returns. Fama (1970) introduced the efficient market hypothesis, which, in its weak form, postulates that prices already incorporate all available information from past prices. Poterba and Summers (1988) applied this hypothesis to the U.S. aggregate stock market, finding that monthly returns exhibit a small, statistically insignificant positive autocorrelation over a 12-month period.[1] The lack of correlation between past and future market returns is not only a well-recognized statistical phenomenon but also an emblem of market efficiency.

We start by documenting that the U.S. market return during the second month of a quarter negatively predicts the first month's return in the next quarter. These first months of a quarter are important because, owing to the U.S. corporate earnings cycle, they contain fresh earnings news about the aggregate economy. For example, at the end of December, firms close their books for Q4. They announce Q4 earnings in January, February, and March. January, therefore, contains the early earnings announcements and is the first time that investors learn about the economy's performance in Q4. The first months of the quarters are known as earnings seasons among practitioners and receive heightened attention. Throughout this paper, we refer to them as "newsy" months, as they produce fresh news about aggregate earnings.

The second month of each quarter contains roughly the same number of earnings announcements as does a newsy month. However, the aggregate earnings announced in the second month are to some degree repetitive of those announced in the preceding newsy month because the announced earnings are for the same fiscal quarter (Q4 in our example above). We call these second months "repetitive" months, as they contain earnings

[1] Furthermore, Poterba and Summers (1988) noted a negative autocorrelation over longer time horizons, while Lo and MacKinlay (1988) identified a positive autocorrelation in weekly returns for small stocks. Despite these variations, the consistently low month-to-month autocorrelation in the aggregate market returns remains an important piece of evidence supporting market efficiency.

announcements that produce repetitive aggregate earnings "news." Such repetition is not verbatim—different firms announce in the first and the second months of a quarter—but subtle in that the announcements pertain to earnings that share an aggregate component. In contrast, the third month of a quarter contains substantially fewer earnings announcements, and thus any repetition would likely be muted.

Embedded in our decision to look at this pattern of predictability is a hypothesis: imperfectly rational investors use announced earnings to forecast future earnings but fail to fully account for the repetitive nature of the earnings in the second month of a quarter. This hypothesis generates the prediction above: investors overreact in the repetitive months, and, when more about the true state is revealed in the next newsy month, there is a correction and thus the predictable reversal. Furthermore, because the "news" is the same in the second month, investors should react to it in the same direction as in the first month. This mechanism additionally predicts a companion continuation arm of return autocorrelation from the first month of a quarter to the second month. We show that both arms of return autocorrelation are economically and statistically significant in the data. Put compactly, the return pattern surrounding our mechanism is that a newsy month's return positively predicts the next repetitive month's return, which then negatively predicts the next newsy month's return.

Enke and Zimmermann (2019) previously identified this mechanism, correlation neglect, in experimental data. We formalize this mechanism using a model and show that it has the implication stated above. We also show how this mechanism arises naturally from basic memory processes, as in Wachter and Kahana (2024).[2] Correlation neglect can also be viewed as an instance of behavioral inattention (Gabaix, 2019), in which investors are inattentive to precisely the feature of the data-generating process that should lead them to pay less attention.[3] Our theory of correlation neglect generates additional predictions. The first

[2]A sophisticated theoretical literature develops models of correlation neglect in settings where agents learn from one another and from various contemporaneous signals (DeMarzo et al., 2003; Eyster and Rabin, 2010; Golub and Jackson, 2010; Bohren, 2016). Unlike in previous settings, our agents learn over time and the information becomes embedded in asset prices, requiring, as we show, a somewhat different approach in the specification of the signals and the fundamentals. Nonetheless, a basic similarity is that agents keep a "running tally" (DeMarzo et al., 2003), rather than unraveling the correlations with each new observation.

[3]In our sense, as in Enke and Zimmermann (2019), correlation neglect is the agent thinking they have

prediction gets to the core of our mechanism—a second, repetitive signal to which investors overreact. If this repetitive signal drives our return pattern, the pattern should amplify when it is more repetitive, that is, more similar to the first signal. We find that, indeed, our return pattern is substantially stronger when earnings announcements in the repetitive month more closely resemble those in the newsy month.

Second, because the mechanism is based on expectations, we should see it in expectations data. Controlling for average latency in revisions, earnings revisions in repetitive months are positively correlated with the previous month's revisions, whereas earnings revisions in the next quarter's newsy months are negatively correlated with the previous repetitive month's revisions.

Third, according to the theory, the eventual reversal occurs because of the initial overreaction. We show that these two arms of the return pattern are strongly correlated and fluctuate together. We further show that our return pattern weakens when the predictor month's return is driven by large, non-earnings-related shocks, and amplifies when it is not. This identifies a driving force behind the fluctuations in the strength of our return pattern over time.

Lastly, our model invokes fundamental aspects of behavior rather than specific aspects of the aggregate market. Thus, the pattern should also exist for industry returns relative to the market, which we also verify: excess returns in a repetitive month are positively correlated with the previous month's excess returns and negatively predict the following newsy month's excess returns. Furthermore, the cross-sectional results are stronger in larger industries, where investors are more likely to receive a repetitive signal in the second month. Each of these findings is predicted by the correlation neglect model.

These findings have broader implications. Correlation neglect represents a fundamental departure from rational expectations. Prior work links correlation neglect to failures of democratic institutions (Ortoleva and Snowberg, 2015) and places it among the factors

more information than they really have: the same evidence twice is treated as independent pieces of evidence. Wu (2023) documents a different kind of correlation neglect, in which forecasters underappreciate the correlation between the variables they forecast.

leading to the 2008 financial crisis (Akerlof and Shiller, 2009).[4] At the same time, an empirical literature provides direct evidence of correlation neglect in an experimental setting (Enke and Zimmermann, 2019), in the pricing of individual securities (Fedyk and Hodson, 2023; Tetlock, 2011), and in announcement day returns (Gilbert et al., 2012). Our study shows that correlation neglect is large enough to affect aggregate stock returns at a monthly horizon. For this to be the case, correlation neglect needs to be present among a broad set of sophisticated and financially incentivized agents. Our findings therefore support the view that correlation neglect is both pervasive and significant.

Our paper relates to several strands of literature, beyond that of correlation neglect described above. First, we build on the work of Guo (2025) in tracing return predictability to expectation errors rooted in the earnings cycle. Guo shows that a sum of the previous 4 newsy month returns predicts the returns in the intermediate two months with a positive sign and predicts the subsequent newsy month with a negative sign. In return predictability regressions that follow, we control for this effect, and find both effects are significant in the data. Guo's effect traces back to the mechanism of earnings persistence and parameter compression. Earnings have an autoregressive (AR) structure, in which past earnings observations, as far back as a year, predict future earnings. Layering the earnings cycle on top of this basic earnings process implies that calendar time and economic distance are not equivalent, and thus one set of AR parameters should be used to predict earnings in newsy months and a different set in non-newsy months.[5] Guo shows that if investors mistakenly compress both sets of parameters into one, the sum of the last 4 newsy month returns will predict non-newsy months with a positive sign, and newsy months with a negative sign. We show that Guo's parameter compression mechanism accounts for roughly half of the ability of the newsy month to predict the next repetitive month. It implies only a trace of the reversal between the repetitive month and the subsequent newsy month. The observed predictability patterns require two mechanisms; both parameter compression and correlation

[4]The repetition of narratives can create the illusion of independent signals, linking correlation neglect to broader processes of belief formation; see also Goetzmann et al. (2024). Levy and Razin (2015) develop a model in which correlation neglect can improve voting outcomes. Rees-Jones et al. (2024) demonstrate that correlation neglect exacerbates school matching inefficiencies and examine the associated policy implications.

[5]See his Figures 4 and 5 for details.

neglect operate to produce the observed results.[6]

Second, our results bear on the broader literature on predicting the market, with implications for the efficient markets hypothesis. Goyal and Welch (2008) and Goyal et al. (2024) provide comprehensive examinations of market return predictors in the literature. They find that most of the predictors are useless in real time and generate negative out-of-sample R^2s. Even when these out-of-sample R^2s are positive, they tend to be economically small, on the order of tens of basis points,[7] a finding that we replicate on the 46 predictors that Goyal et al. (2024) examine. The reversal arm in our return pattern alone generates an out-of-sample R^2 of more than 2%, and combining both arms of the return patterns improves the out-of-sample R^2 by more than 3 percentage points beyond what Guo's predictor achieves. Even the out-of-sample R^2 of 2% is one order of magnitude larger than that of the median successful predictor and more than twice that of the most successful one. Evidence against the weak form of the Efficient Market Hypothesis only increases when our evidence is combined with what Guo (2025) documents. As we show, there can be subtle relations between past prices that imply abnormal returns, even in the highly liquid aggregate market.

Finally, our paper contributes to the behavioral literature on forecasting. A long-running debate centers on underreaction versus overreaction and which dominates; much of the evidence for underreaction comes from consensus survey forecasts (Coibion and Gorodnichenko, 2012; Bouchaud et al., 2019). The fact that both appear makes it challenging to find a cognitive foundation of behavioral anomalies. From a perspective rooted in diagnostic expectations (Bordalo et al., 2019), in which beliefs overweight outcomes whose likelihood has risen the most in light of recent news, Bordalo et al. (2020) argue that it is overreaction that dominates: individual forecasters overreact to their own information, while consensus forecasts appear to underreact because they aggregate information dispersed across forecasters. In this paper, we show that what might appear to be underreaction followed by overreaction (because there is continuation and then reversal), is in fact overreaction only. The continu-

[6]Lou et al. (2019, 2026) also find predictability that follows calendar-based cycles – in their case the cycles are in hours of the day rather than months of the year. They attribute their findings to clientele effects, an alternative explanation which we discuss in what follows.

[7]See also Campbell and Thompson (2008), Rapach et al. (2010), and Rapach and Zhou (2022).

ation is reversed, and the reason behind the continuation is attributing information content to a signal that is already priced in.

Though our model is one of over-reaction, its mechanism is distinct from diagnostic expectations. Viewed as a process for belief formation, diagnostic expectations involves a distortion in updating, expressed as an extra term mediating between prior, likelihood, and posterior. Correlation neglect is Bayesian, but with an impoverished model that fails to account for features of the data generating process. A memory model provides a foundation for both, allowing them to be disentangled; diagnostic expectations relies on the mechanism of interference, whereas correlation neglect relies on the more basic idea that items that are repeated tend to be easier to remember.

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 describes our return predictability regressions, both in and out-of-sample. Section 4 describes the model and its additional predictions, which Section 5 tests. Section 6 evaluates several alternative explanations. Section 7 concludes.

2 Data

Our return data come from CRSP and Ken French's website, accounting data come from Compustat North America, and LTG forecast data come from IBES Detail History (adjusted). These are common data sources used by a large number of studies. There are two main sources of earnings announcement dates in the United States: Compustat and IBES. Both data sources have been used individually in major studies. These announcement days often disagree, though the disagreements are usually small and mostly concentrated before December 1994 (Dellavigna and Pollet 2009). We implement the algorithm in Dellavigna and Pollet (2009), which combines the Compustat and IBES announcement dates to form the best estimate of the actual announcement dates.[8] If the adjusted announcement date is the same as the IBES announcement date, we also shift the date to the next trading date if the IBES announcement time is after the market closure, following Johnson and So (2018a).

[8]The algorithm is effectively as follows: 1) when the two sources differ, use the earlier date, and 2) when they agree and the date is before January 1, 1990, shift the day to the previous trading date.

This gives us earnings announcement dates since 1971, although NYSE listing requirements have mandated quarterly earnings announcements since the 1910s (Kraft et al. 2017).[9]

In terms of specific variables, we use the market factor from Fama and French (1993) to represent aggregate market excess returns and the SIC code to represent industries. We use "Report Date of Quarterly Earnings" (`rdq` from Compustat) and "Announcement Date, Actual" (`anndats_act` from IBES) for earnings announcement dates, "Common/Ordinary Equity–Total" (`ceqq`) for book value of equity, and "Income Before Extraordinary Items" (`ibq`) for earnings.

3 Predicting aggregate stock market returns

3.1 The earnings cycle

We first characterize the US earnings cycle, which originates from i) aligned fiscal quarters and ii) heterogeneous announcement lags. Panel A of Table 1 shows that the vast majority of US firms close their books for a given fiscal quarter on the last day of a calendar quarter. Therefore, the fiscal quarters over which the earnings are recorded are effectively the calendar quarters. Although fiscal quarters are aligned, firms do not announce their quarterly earnings together. Almost all firms announce within three months of the fiscal quarter end date, but there is much variation in their announcement lags. Panel B shows that among the on-cycle announcements, about 45% occur in the first month of a quarter, about 48% in the second month, and about 8% in the last month.

The quarterly cycles of earnings announcements imply that each of the three months of the quarter has a different status. The first month is the first time that investors learn about the *aggregate* economy's performance in the previous quarter. This economy-level learning is possible because corporate earnings have a common time component: the aggregate economic condition in fiscal quarter 4 (FQ4) of 2023 influences all firms,[10] and consequently, if

[9] When the Securities Exchange Act of 1934 was passed by the US congress, most companies did not find the mandated quarterly financial reporting new, owing to the earlier requirement from NYSE.

[10] Appendix D empirically analyzes this common time component and demonstrates its economic significance.

early announcers' FQ4 earnings are good, then FQ4 is more likely than not a good quarter for everyone. Earnings announced in the first month of a quarter therefore provide new information about aggregate earnings earned in the previous quarter. We refer to these first months of quarters as newsy earnings months, or newsy months for convenience.

The second months, in contrast, see earnings that are predictably repetitive of those announced in the first months. This is again owing to the shared aggregate component of corporate earnings: earnings announced in February, like those announced in January, are earned in Q4 of the previous year and therefore inherently similar. We refer to the second months of the quarters as the repetitive earnings months, or simply repetitive months. The third months of the quarters also contain announcements of predictably repetitive earnings. However, since they contain substantially fewer on-cycle earnings announcements, the repetition is muted.

3.2 Forecasting regressions: a first look at the return pattern

Two arms of serial correlation As noted in the introduction, the mechanism of correlation neglect gives rise to a reversal arm and a continuation arm of return serial correlation. We first plainly illustrate them with the following time-series regressions:

$$mkt_t = \alpha + \beta mkt_{nr(t)} + \epsilon_t. \tag{1}$$

Here, mkt_t is the return of the US stock market in excess of the risk-free rate in month t, and $nr(t)$ is the most recent month before t that is newsy or repetitive. Also, we do not include cases in which t is the third month of a quarter. For example, if t is February, then $nr(t)$ is January, and if t is April, $nr(t)$ is February.

Panel A, Column 1 of Table 2 performs regression 1 on the sub-sample in which the dependent variable month t is newsy. The negative coefficient of -0.277 implies that a repetitive month's return negatively predicts the return in the next newsy month. This reversal arm of return autocorrelation is our main result. Column 2 shows the companion continuation arm. The positive coefficient of 0.271 shows that a newsy month's return positively predicts

Figure 1: The earnings cycle and return predictability

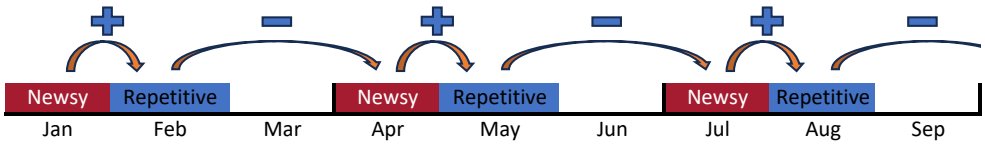

the return in the subsequent repetitive month. These values are economically sizable, given they are bounded by -1 and 1 in a stationary time series. Figure 1 qualitatively illustrates this pattern. Without the third months of quarters that contain muted earnings news, the first-order autocorrelation of the monthly stock market returns alternates between strongly positive (newsy predicting repetitive) and strongly negative (repetitive predicting newsy).

Investors making full use of both arms of the correlation should refer to Column 3, which extracts the difference between Columns 1 and 2. Economically, it demonstrates that the monthly autocorrelation of stock market returns varies strongly and predictably. The associated t-statistic of -4.29 corresponds to an in-sample Sharpe ratio of around 0.4 that investors would have earned historically in a *beta-neutral* market-timing strategy that takes advantage of both the continuation and reversal arms. This is economically sizable, given that the strategy takes a zero equity position on average. Panels B, C, and D of Table 2 report the same set of numbers in the first half, second half, and post-WWII portion of the sample, respectively. The pattern operates in each of these episodes and is not the sole result of a small number of observations.

Relation to Guo (2025) Guo (2025) documents that, owing to a separate mechanism of "parameter compression," another pattern of return predictability revolves around the earnings cycle. His flagship predictor, introduced in his Table 3, is the sum of the *past four newsy month* returns; it positively forecasts repetitive months and negatively forecasts newsy months. His pattern relates to ours. The most recent newsy month, the independent variable of our continuation arm, is one of the four months in his predictor. And, as the introduction notes, his mechanism implies a trace of our reversal arm, because the repetitive month's return echoes the newsy-month returns that precede it.

Figure 2 shows how the two mechanisms divide the continuation arm. The figure plots the coefficients in a regression of repetitive month returns on the past four newsy month returns. The estimated loadings are 0.27 on the first lag and about 0.05 on each of the others. Parameter compression implies loadings that decline geometrically with the lag, and his flagship predictor, which weights the four lags equally, implies the equal loading of 0.13 shown by the dashed line. A single rate of geometric decline cannot produce the estimated pattern: matching the drop from the first lag to the second implies near-zero loadings on the third and fourth, whereas matching the flat loadings on the later lags implies little decline at the first. The pattern is instead consistent with the two mechanisms operating together: a slowly declining component from parameter compression, and a component from correlation neglect concentrated on the most recent newsy month, the aggregate component of whose earnings is the one re-announced in the repetitive month.

Table 3 separates the mechanisms directly by including both predictors in one regression. Our coefficients change only modestly relative to Table 2: the reversal arm is -0.247 ($t = -2.73$), the continuation arm is 0.207 ($t = 2.15$), and the difference is -0.454 ($t = -3.44$). His predictor remains significant as well, with a difference coefficient of -0.140 ($t = -3.00$): both mechanisms operate in the data. Consistent with the trace implied by his mechanism, his predictor's own reversal-arm coefficient is small (-0.062, $t = -1.87$). In later tests of return predictability, we use the difference coefficient estimated controlling for his predictor as the key coefficient that captures correlation neglect.

3.3 Formal performance evaluation via out of sample metrics

In an influential critique of the return predictability literature, Goyal and Welch (2008) note that when data available in real time are used to forecast returns, predictors often generate negative out-of-sample (OOS) R^2. This means that equity premium estimates generated using these predictors often fail to outperform the historical average despite being (sometimes highly) significant in-sample. As Lewellen and Shanken (2002) emphasize, predictability that is visible ex post need not be exploitable in real time. Out-of-sample R^2 has since been established as the formal evaluation metric for aggregate stock market return predictors,

whereas in-sample metrics such as the t-statistic are only suggestive. We now evaluate the forecasting performance of our predictor, with an emphasis on the reversal arm.

Table 4 presents the results. Following Goyal and Welch (2008), we compute OOS R^2s as $1 - \frac{\sum_{t=1}^{n}(mkt_t - f_{t-1}(z_{t-1}))^2}{\sum_{t=1}^{n}(mkt_t - \overline{mkt}_{t-1})^2}$, where $\overline{mkt}_{t-1}$ is the expanding window mean of past stock market excess returns up to month $t-1$, and $f_{t-1}(z_{t-1})$ is the forecast being evaluated, formed based on a predictor z_{t-1} also using information up to month $t-1$. The value $\overline{mkt}_{t-1}$ makes no use of the predictor z_{t-1} and serves as a benchmark of evaluation. The R^2s reflect whether the evaluated forecasts—which attempt to use the predictor z in real time—generate higher or lower forecast errors than the benchmark.

In our case, we compute the R^2 where the target month t is either repetitive or newsy, and the underlying predictor is $z_{t-1} \equiv mkt_{nr(t)} - \overline{mkt}_{t-1}$ for a repetitive t, $z_{t-1} \equiv -(mkt_{nr(t)} - \overline{mkt}_{t-1})$ for a newsy t. Following Campbell and Thompson (2008), we extend the CRSP aggregate market return series to 1872 using data from Global Financial Database, and evaluate the R^2 from 1926. For each target month t, 2 coefficients $\hat{\beta}_{t-1}$ and $\hat{\alpha}_{t-1}$—the former for the predictor and latter for the constant term—are extracted from a simple expanding-window OLS regression of historical market excess returns on historical predictor values and a constant using data up to $t-1$. The forecast for month t that uses the predictor z_{t-1} is then $\hat{\beta}_{t-1}z_{t-1} + \overline{mkt}_{t-1}$.

This approach allows us to separately discern the forecasting success of the continuation and the reversal arm of the return pattern. In the first column, the forecast takes the neutral value of $\overline{mkt}_{t-1}$ when t is repetitive, and $\hat{\beta}_{t-1}z_{t-1} + \overline{mkt}_{t-1}$ when t is newsy. In other words, it uses only the reversal arm of our return pattern. In the second column, the forecast is $\hat{\beta}_{t-1}z_{t-1} + \overline{mkt}_{t-1}$ when t is repetitive, and $\overline{mkt}_{t-1}$ when t is newsy, using only the continuation arm. In the third column, the forecast is $\hat{\beta}_{t-1}z_{t-1} + \overline{mkt}_{t-1}$ in all cases, using both arms of our return pattern. The fourth column uses the predictor that Guo (2025) proposes and computes its R^2 as a benchmark. The fifth column combines the forecast based on our predictor and the forecast based on his by simply taking their average, following Rapach et al. (2010). The difference of about 3.27% between Columns 5 and 4 is therefore a point estimate of the marginal contribution of our predictor beyond his, whereas

the first column, which uses only the reversal arm, serves as a lower bound.

Column 1 of Table 4 shows that the reversal arm alone generates an OOS R^2 of 2.75%. Figure 3 places this value among the 19 positive monthly OOS R^2s of the 46 predictors that Goyal et al. (2024) survey, following their computation; their median is 0.26% and their maximum is 1.09%.[11] Even with no loading on the continuation arm, the reversal arm generates an OOS R^2 an order of magnitude above the median successful predictor and more than twice the maximum. The continuation arm alone generates 3.20% (Column 2), and the two arms together generate 5.95% (Column 3). The predictor that Guo (2025) proposes generates 3.68% (Column 4), in line with the values that he reports.[12] Combining the two predictors yields the largest value in the table, 6.95% (Column 5). All of these values are computed in real time and use past returns only.

This section shows that our statistical findings would have helped to forecast the return on the market portfolio, to an economically significant degree. This is a basic plausibility check that, nonetheless, many predictors fail, and one that complements the in-sample findings in Section 3.2. Taken together with the findings that Guo (2025) reports, these results further suggest revising one's view of the weak form of the Efficient Market Hypothesis. Having established the significance of this predictability, we now propose and test a mechanism.

4 Model

A parsimonious mechanism that jointly explains the continuation and reversal arms of the return pattern is correlation neglect: investors do not fully discount the inherently repetitive earnings announced in the second month of a quarter and thus overreact to them, creating the return continuation that gets corrected in the next newsy month. Below, we build a model of correlation neglect, which we show, in Appendix B, emerges from basic principles of human memory. While our model shares a basic mechanism with that of Enke and Zimmermann (2019), it differs from theirs in several respects. In Appendix C, we explain

[11] The thresholds for lower-frequency predictors predicting quarterly or annual returns are naturally higher. All OOS R^2s we report are monthly to make the comparison apples-to-apples.

[12] His published values range from 3.6% to 3.9%.

why it is necessary to depart from the Enke and Zimmermann model to apply their intuition to asset prices.

Assume that the economy is either in a high or low-productivity state. If the economy is in the high-productivity state, the aggregate market is worth 1, whereas it is worth 0 if it is in a low-productivity state. The high-productivity state occurs with probability p, which is unknown to investors.

This model has three periods. Period 1 corresponds to the first month of a quarter, namely a newsy month. Period 2 is the repetitive month that follows. Period 3 is the first month of the following quarter. In period 1, investors receive a signal $\hat{s}_1 = s_1$, with $s_1 \sim \text{Bernoulli}(p)$. This signal represents the information pertaining to the aggregate state that investors extract from firms that announce in that month. In period 2, this signal is potentially repeated (as opposed to an independent signal being drawn from Bernoulli(p)). Whether or not the signal is repeated is independent of the original signal. We let p_x denote the probability that a signal is repeated, and let $x \sim \text{Bernoulli}(p_x)$ determine if it is repeated or not. If $x = 1$, then s_1 is brought again to investors' attention, namely $\hat{s}_2 = s_1 = \hat{s}_1$. If $x = 0$, investors see an independent draw from Bernoulli(p), $\hat{s}_2 = s_2$. This second signal, like the first, represents the information pertaining to the aggregate state that investors extract from firms announcing in that month. This information is potentially repetitive because there is no a priori reason why the aggregate market state as of the previous quarter would have affected the earnings of firms announcing early or late differently.[13] It is potentially, but not necessarily, repetitive because the information in the first month of the quarter may have been mainly idiosyncratic. There is random variation in the degree to which the information in the first month reflects the second month; this random variation is what the binary variable x is meant to capture.

To focus on the mechanism of correlation neglect, we assume investors are risk-neutral and the riskfree rate is equal to zero. We assume that the agent is using signals extracted from individual firms to estimate the true state of the market, which will be revealed in the

[13]Separately, the aggregate state may evolve over time. Investors' consideration of this evolution, and potential biases stemming from the earnings announcement cycle, is the subject of Guo (2025). We compare the mechanisms in Section 3.2.

subsequent quarter. The real world can therefore be viewed as a sequence of independent but overlapping three-period models, in which we abstract from autocorrelation in the true state of the economy. Investors start with a uniform prior distribution over p.[14] We also assume that investors are, on average, rational, in that the true value of p is drawn from the prior distribution.

Under these conditions, the value of the market is equal to the posterior mean of p. If, out of T observations, N indicate that the economy is in a high-productivity state, this posterior mean equals

$$E[p|N \text{ out of } T] = \frac{N+1}{T+2}. \tag{2}$$

Equation 2 is Laplace's rule of succession. The investors' prior implies that the price prior to observing the signal in the first period is $P_0 = E[p] = 1/2$. After observing signal s_1 in the first period, the mean shifts so that the price equals

$$P_1 = E[p|s_1] = \frac{s_1+1}{3}.$$

For simplicity, we define returns as differences in prices rather than percent changes. The return $R_1 = P_1 - P_0$. Note that $E[R_1] = 0$.

Equilibrium with rational investors

The rational investor computes the period-2 price P_2 given observed $\hat{s}_2$. The investors do not see the outcome of x, but incorporate the possibility that $x = 1$, and therefore, that $\hat{s}_2$ is a repeat of an earlier signal. In expectation, given s_1,

$$E[P_2|s_1] = E[P_2|x=1, s_1]P(x=1|s_1) + E[P_2|x=0, s_1]P(x=0|s_1) \tag{3}$$

Because x is independent of s_1, (3) reduces to

$$E[P_2|s_1] = E[P_2|x=1, s_1]p_x + E[P_2|x=0, s_1](1-p_x).$$

[14]The results that follow hold under a general Beta prior; the uniform prior eliminates unnecessary notation.

If rational investors were to know $x = 1$, they would not update based on s_2. Therefore $E[P_2|x = 1, s_1] = P_1$. Moreover, if investors were to know $x = 0$, they would update based on the outcome of s_2 as an independent draw. Because (2) follows a martingale as observations accumulate, we again have $E[P_2|x = 0, s_1] = P_1$.[15] It follows that $E[R_2] = E[E[R_2|s_1]] = E[E[P_2 - P_1|s_1]] = 0$ and that

$$E[R_2R_1] = E[E[R_2R_1|P_1]] = E[E[(P_2 - P_1)R_1|P_1]] = E[E[P_2 - P_1|s_1]R_1] = 0.$$

The second-to-last equality follows because P_1 reveals s_1 and because R_1 is known given s_1. The last equality follows because $E[P_2 - P_1|s_1] = 0$ for all values of s_1. Therefore $\text{Cov}(R_1, R_2) = 0$, and there is no continuation in the case of a rational investor. Similar reasoning establishes $E[R_3R_2] = 0$ and thus that there is no reversal in the case of a rational investor.

Equilibrium with correlation neglect

The investor with correlation neglect calculates

$$P_2 = \frac{\hat{s}_1 + \hat{s}_2 + 1}{4}$$

[15]Note that

$$\begin{aligned} E[P_2|s_1, x = 0] &= \frac{s_1 + 1}{4}(1 - P_1) + \frac{s_1 + 2}{4}P_1 \\ &= \left(\frac{s_1 + 2}{4} - \frac{s_1 + 1}{4}\right)P_1 + \frac{s_1 + 1}{4} \\ &= \frac{1}{4}P_1 + \frac{3}{4}P_1 = P_1. \end{aligned}$$

Using an argument similar to (3),

$$\begin{aligned} E[P_2|s_1] &= E\left[p_x\frac{2s_1+1}{4}+(1-p_x)\frac{s_1+s_2+1}{4}\right] \\ &= E\left[p_x\left(\frac{2s_1+1}{4}-P_1\right)+p_xP_1+(1-p_x)E\left[\frac{s_1+s_2+1}{4}\middle| s_1\right]\right] \\ &= E\left[p_x\left(\frac{2s_1+1}{4}-P_1\right)+p_xP_1+(1-p_x)P_1\right], \end{aligned}$$

where the last line follows from the aforementioned martingale property. Therefore

$$E[P_2]=p_xE\left[\frac{2s_1+1}{4}-P_1\right]+E[P_1]=E[P_1],$$

establishing that $E[R_2]=0$. We have used the fact that

$$E\left[\frac{2s_1+1}{4}-P_1\right]=E\left[\frac{2s_1-1}{12}\right]=0,$$

because $E[s_1]=E[p]=1/2$.[16]

Similarly, if we let $P_2^{\mathcal{R}}$ denote the period-2 price for the rational investor,

$$E[R_3]=E[P_3-P_2^{\mathcal{R}}+P_2^{\mathcal{R}}-P_2]=E[P_2^{\mathcal{R}}-P_2]=E[P_1-P_2]=0,$$

where we twice use the martingale property. The last line follows from $E[R_2]=0$ as mentioned previously. The investor with correlation neglect is not directionally right or wrong.

However, unlike in the case with the rational agent, R_2 is positively correlated with R_1 and R_3 negatively correlated with R_2. Intuitively, the agent double-counts s_1 in the second period, not adjusting for the fact that it may be the same signal as before. Because R_2 has mean zero, it suffices to show that $E[R_2R_1]>0$. Again applying the argument from (3),

$$E[R_2R_1]=E[E[R_2R_1|\hat{s}_1]|x=1]p_x+E[E[R_2R_1|\hat{s}_1]|x=0](1-p_x).$$

[16] We can make a stronger statement: because s_1 is a binary variable, it must have a Bernoulli distribution. Because $E[s_1]=E[E[s_1|p]]=1/2$, the properties of the Bernoulli distribution indicate that the unconditional probability of $s_1=1$ is $1/2$.

The second term equals zero. Therefore

$$E[R_2R_1] = E[E[R_2R_1|\hat{s}_1]|x=1]p_x = E\left[\left(\frac{2s_1-1}{12}\right)\left(\frac{2s_1-1}{6}\right)\right]p_x > 0, \tag{4}$$

where the second equality follows because R_1 is a constant in $\hat{s}_1$.

Finally, we show reversal: $E[R_3R_2] < 0$. We follow the same line of argument as above for $E[R_2R_1] > 0$, breaking up the unconditional expectation into terms that condition on the outcome of x:

$$E[R_3R_2] = E[E[R_3R_2|\hat{s}_1,\hat{s}_2]|x=1]p_x + E[E[R_3R_2|\hat{s}_1,\hat{s}_2]|x=0](1-p_x).$$

The second term equals zero. To evaluate the first term, note that $E[P_3|x=1,\hat{s}_1,\hat{s}_2] = P_1$ because there is no information in period 2. By definition, $P_3 - P_1 = P_3 - P_2 + P_2 - P_1$. Therefore $E[R_3 + R_2|x=1,\hat{s}_1,\hat{s}_2] = 0$. It follows that any reaction to "news" in the second period will be an over-reaction that will, on average, be corrected in the third and final period:

$$\begin{aligned} E[R_3R_2] &= E[E[R_3R_2|x=1,\hat{s}_1,\hat{s}_2]]p_x \\ &= E[E[R_3|x=1,\hat{s}_1,\hat{s}_2]R_2]p_x, \\ &= -E[R_2^2]p_x < 0, \end{aligned} \tag{5}$$

where the second line follows because R_2 is a constant in $\hat{s}_1, \hat{s}_2$.

Clearly, both continuation and reversal are increasing in p_x, which leads to testable predictions that we explore in the next section. Moreover, if the signal in period 2 is similar to that of period 1, continuation and reversal will be greater than otherwise. In our stylized set-up, we compare continuation and reversal under $\hat{s}_1 = \hat{s}_2$ to $\hat{s}_1 \neq \hat{s}_2$. The probabilistic environment is designed to capture the main feature of the analysis in Section 5.1 below, in which samples are divided based on the similarity of the signals in period 2. This is an ex-post conditioning that can affect measurement of the population moments: the outcome $\hat{s}_2 \neq \hat{s}_1$ (or $\hat{s}_2 = \hat{s}_1$) conveys information about the unknown p that is not realized until

the second period. Under the uniform prior, the conditional means of R_1 and R_2 vanish by symmetry; under a general Beta prior they do not, and so we center the moments below. In either case, $E[R_3|\hat{s}_2 \neq \hat{s}_1] = E[R_3|\hat{s}_2 = \hat{s}_1] = 0$ because the information is already embedded in P_2 and because $E[R_3] = 0$.

Our model is simplistic in that signals are zero or 1, and thus they can either be equal or diametrically opposite. In this extreme case, when $\hat{s}_1 \neq \hat{s}_2$, we no longer see continuation:

$$E\left[(R_2 - E[R_2|\hat{s}_1 \neq \hat{s}_2])(R_1 - E[R_1|\hat{s}_1 \neq \hat{s}_2]) \mid \hat{s}_1 \neq \hat{s}_2\right] = -E\left[(P_1 - E[P_1|\hat{s}_1 \neq \hat{s}_2])^2\right] < 0, \tag{6}$$

noting that P_2 is a constant under the conditioning set. Moreover, because $P_2 = P_2^{\mathcal{R}}$, the rational forecast in period 2 under this condition, there is no reversal.[17]

We now show that there is continuation in the case of correlation neglect and $\hat{s}_1 = \hat{s}_2$.[18] When $\hat{s}_1 = \hat{s}_2$, in the economy with correlation neglect, $P_2 = (2s_1 + 1)/4$, so that $R_2 = (2s_1 - 1)/12$. Recall that $R_1 = (2s_1 - 1)/6$. Finally,

$$\begin{aligned} E\left[(R_2 - E[R_2|\hat{s}_1 = \hat{s}_2])(R_1 - E[R_1|\hat{s}_1 = \hat{s}_2])|\hat{s}_1 = \hat{s}_2\right] &= \frac{1}{18}E[(s_1 - E[s_1|\hat{s}_1 = \hat{s}_2])^2|\hat{s}_1 = \hat{s}_2] \\ &> 0 \end{aligned} \tag{7}$$

where the equality follows because all constant terms cancel out with their expectations. This establishes continuation in the case of $\hat{s}_1 = \hat{s}_2$. Finally, because $E[R_3R_2]$ is a weighted average of its value under the conditions $\hat{s}_1 = \hat{s}_2$ and $\hat{s}_1 \neq \hat{s}_2$, and because it is equal to zero under the latter condition (when $P_2 = P_2^{\mathcal{R}}$), it must be negative when $\hat{s}_1 = \hat{s}_2$. The argument suffices because $E[R_3|\hat{s}_2 \neq \hat{s}_1] = E[R_3|\hat{s}_2 = \hat{s}_1] = 0$ as mentioned earlier. Therefore, the continuation and reversal pattern is stronger when signals are more similar than otherwise, an implication of the model which we test below.

[17] P_3 is a draw from a Bernoulli distribution with mean $P_2^{\mathcal{R}}$, and $P_3 - P_2^{\mathcal{R}}$ is independent of anything in the information set as of Period 2.

[18] One might think that this step would be unnecessary, since we know that there is continuation in population and also that there will be no continuation in samples with $\hat{s}_1 \neq \hat{s}_2$. However, while $E[R_1R_2]$ is a weighted average of $E[R_1R_2|\hat{s}_1 = \hat{s}_2]$ and $E[R_1R_2|\hat{s}_1 \neq \hat{s}_2]$, the same cannot be said for $\mathrm{Cov}(R_1, R_2)$ under a general prior, because the conditional means need not agree across the two conditioning sets.

5 Evaluating the model's predictions

The model in Section 4 explains why returns in the first month of each quarter predict those of the second month with a positive sign and those of the second month predict those of the subsequent first month of the quarter with a negative sign. We now derive several additional predictions.

The first prediction lies at the heart of correlation neglect—similar signals. If they indeed drive our return pattern, the pattern should amplify when the signals are more similar. Equations 6–7 and the surrounding discussion show that samples in which the second signal is relatively similar to the first are those in which continuation and reversal will be enhanced:

Prediction 1 (Similarity of announcements)**.** *Return predictability will be stronger when the announcement in the repetitive month is more similar to the previous newsy month.*

Section 5.1 discusses tests of this hypothesis.

The second prediction is straightforward. Given that the driving force is investors' beliefs, we should see analogous results in survey data. Note that the expectation revision in a newsy month is $E^{\mathrm{CN}}[p|s] - 1/2$, where $E^{\mathrm{CN}}[\cdot]$ denotes the subjective expectation of the investors exhibiting correlation neglect. It equals R_1. The revision in the repetitive month is $E^{\mathrm{CN}}[p|s,x] - E^{\mathrm{CN}}[p|s]$, which equals R_2. These are positively correlated, as we show above. In the following newsy month, the expectation revision, if there is a positive outcome, is $1-E^{\mathrm{CN}}[p|s,x]$ and $-E^{\mathrm{CN}}[p|s,x]$ otherwise. For the reasons described above, this is negatively correlated with the revision in the repetitive month.

Prediction 2 (Survey data)**.** *Revisions in beliefs in a repetitive month are likely to be in the same direction as those in the previous newsy month, and revisions in a newsy month are likely in the opposite direction to revisions in the previous repetitive month.*

Section 5.2 discusses tests of this hypothesis.

Equations 4 and 5 show that when p_x increases, so do both continuation and reversal. The parameter p_x itself is unobservable. Nonetheless, several tests of the model emerge from these equations, and it is useful to distinguish the parameter from its realization.

The parameter p_x is the probability that the repetitive month repeats the newsy month's information; the realization x determines whether, in a given quarter, it does. The strength of the return pattern can therefore vary for three reasons: the state p_x may vary slowly over time, the realization varies quarter by quarter, and p_x may differ across groups of stocks. The similarity evidence that Section 5.1 presents concerns the realization; Prediction 5 below concerns the cross-section. Here, suppose that p_x is a persistent state variable that varies over time. We would observe periods of time in which continuation is strong to also be those times in which reversal is strong.

Prediction 3 (Association of the two arms)**.** *Continuation and reversal share an underlying mechanism. Therefore when continuation is particularly strong in the data, so is reversal.*

Section 5.3 discusses tests of this hypothesis.

Identifying the model's signals with earnings announcements yields a further prediction: the return pattern should appear only when and where earnings announcements occur.

Prediction 4 (Timing within the quarter)**.** *The return pattern should be absent in the first week of newsy months, before the earnings season begins, and should not concentrate on macroeconomic announcement days.*

Section 5.4 discusses tests of this hypothesis.

To connect the theoretical results from the model to the empirical results, we identify $\hat{s}_1$ and $\hat{s}_2$ with company-level announcements that are informative regarding the underlying state of the economy p. The mistake that investors make lies in assuming that they are independent draws. We could, however, have equally well viewed $\hat{s}_1$ and $\hat{s}_2$ as company-level announcements that are informative of a particular industry, which is a smaller but more tightly connected economy. This prediction is especially interesting, as investors appear keen to learn industry-level information in the newsy months (e.g., Foster (1981), Brochet et al. (2018), Thomas and Zhang (2008), Hilary and Shen (2013), Hann et al. (2019)). We should then observe that industry returns in excess of the market should also exhibit a similar continuation and reversal pattern. Furthermore, a larger industry, like the market itself, has a better diversified aggregate component that can be measured with less noise. In the model,

this could be interpreted as differences in p_x between industries. Specifically, p_x increases with size, as idiosyncratic components of earnings are increasingly diversified.

Prediction 5 (Industry returns)**.** *Industry excess returns should also exhibit continuation and reversal, and the return pattern is stronger for larger industries.*

Section 5.5 discusses tests of this hypothesis. Appendix E extends the analysis to global stock market returns.

5.1 Earnings similarity and the strength of the return pattern

According to our mechanism, investors see a signal in the second month, which they believe is an independent draw from the distribution that produced the signal in the first month. In the data, however, this second signal comprises earnings announcements from a different set of firms. These signals have a range of similarities with the first signal, which we capture with differences between aggregate return on equity in the two months. If the signal that investors see is highly dissimilar, then investors are not likely to interpret it as a draw from the same distribution. Thus, a key test of our underlying mechanism (see Prediction 1) is that overreaction is likely to be especially strong when the earnings in the repetitive month are similar to those announced in the previous newsy month. This prediction lies at the core of our mechanism—it revolves around a repetitive signal, and should therefore amplify with the signal's repetitiveness.

We now test this prediction. Table 5 reports the difference coefficient estimated controlling for the predictor that Guo (2025) proposes, as in Table 3; this coefficient captures our mechanism. It is computed separately for episodes sorted by the distance between the aggregate ROE announced in the newsy month and in the subsequent repetitive month, into four groups: from the lowest decile of distance, in which announcements are most similar, to the highest decile. When the earnings announced in the repetitive month closely resemble those in the preceding newsy month, the pattern is strongest: the coefficient is -0.800 ($t = -4.37$) in the lowest-distance decile, and its magnitude is much smaller in the higher-distance groups. In the highest-distance decile, the coefficient is -0.127 ($t = -0.48$):

the pattern vanishes rather than changes sign. The attenuation distinguishes correlation neglect from an alternative in which investors group months into categories and overestimate the correlation even when announcements are dissimilar. Under that alternative, investors would underweight the genuinely new content of a dissimilar repetitive month, producing a pattern of the opposite sign. Under correlation neglect, a dissimilar announcement leaves little repeated content to neglect, and the pattern disappears. The highest-distance decile contains 42 months, so this comparison has limited power; the point estimates are consistent with correlation neglect.

The repetition channel also requires that the predictor month's return reflect announcements in the first place. Non-fundamental shocks have nothing to do with the earnings cycle and have no reason to be repeated in the second month of a quarter: a predictor month whose return they drive should work less well, whatever the prevailing p_x. To measure the extent to which a month's return is driven by such shocks, we construct a daily series of non-fundamental shocks and, following Da et al. (2014), record the degree to which the days that share the month's direction are outnumbered; Appendix A details the construction.[19] Table 6 splits the 1982–2024 sample, for which the daily classification is available, at the median of this measure. When the predictor month is less driven by non-fundamental shocks, the difference coefficient is -0.378 ($t = -3.69$), estimated controlling for the predictor that Guo (2025) proposes; when it is more driven by them, the coefficient is near zero (0.021, $t = 0.08$). The similarity evidence is the more direct test, conditioning on the announcements themselves; this measure conditions on the return that transmits them. Both point the same way.

5.2 Survey data

An influential branch of the behavioral finance literature uses survey data from IBES to measure investors' expectations of aggregate cash flow growth (Nagel and Xu, 2022; De La O and Myers, 2021; Bordalo et al., 2024). As Prediction 2 states, our model traces the continuation and reversal effects to beliefs. Following Bordalo et al. (2024), we extract beliefs using the

[19] We thank Peter Kelly for inspiring this interpretation.

long-term growth (LTG) measure from the IBES Adjusted Details file.[20] Following Bordalo et al. (2024), we use medians across analysts to represent firm-level consensus, and then aggregate firm-level LTG revisions in each month using the market capitalization weight within the S&P 500 universe to obtain revisions in aggregate cash flow growth expectations. Our model predicts that investors overreact in the repetitive months in the sense that i) revisions in a repetitive month are likely in the same direction as those in the preceding newsy month and ii) revisions in a repetitive month are likely followed by correction in the opposite direction in the subsequent newsy month. These lead to two predictions on the time series of aggregate LTG revisions. First, revisions in a repetitive month are positively correlated with revisions in the previous newsy month. Second, revisions in a newsy month are negatively correlated with revisions in the previous repetitive month.

Column 1 of Table 7 tests prediction i) by running the following regression on the sample in which t is repetitive:

$$Rev_t = \alpha + \beta Rev_{t-1} + \epsilon_t.$$

Here, Rev_t is the aggregate LTG revisions in month t, and as t is repetitive, $t-1$ is newsy. The positive and significant coefficient of 0.543 indicates that revisions in a repetitive month indeed tend to be in the same direction as those in the previous newsy month. However, this number is not in itself convincing. Survey data naturally incorporate latency,[21] resulting in a positive unconditional autocorrelation in revisions. Column 2 controls for this unconditional latency and extracts the extra effect using an interaction term between Rev_{t-1} and a dummy variable indicating whether t is repetitive. The resulting coefficient of 0.323 is again highly significant: the positive relation between the revisions in a repetitive month and those in the previous newsy month is too large to be explained by data latency.

Column 3 tests prediction ii) by running the following regression on the sample in which t is newsy:

$$Rev_t = \alpha + \beta Rev_{t-2} + \epsilon_t.$$

[20] We do not use the Summary files, which contain middle-of-the-month consensuses struck on the Thursday after the 3rd Friday of the month, because we need revisions month-end to month-end.

[21] Barth et al. (2024) and Li et al. (2024) document latency in analyst estimates.

Here, as t is newsy, $t-2$ is repetitive. We do not observe the negative coefficient predicted by our theory, but rather a near-zero coefficient of 0.038. However, this is simply owing to the confounding latency effect mentioned above. Column 4 controls for this effect and extracts the overreaction coefficient using the interaction term between Rev_{t-2} and a dummy variable indicating whether t is newsy. The coefficient is now significantly negative, at -0.174. This implies that reactions in the repetitive months are indeed overreactions that get corrected in the subsequent newsy month.

5.3 The association between the continuation and reversal arms

If p_x is a persistent state that varies over time, both arms of the return pattern strengthen and weaken together: years of strong continuation should also be years of strong reversal. To measure each arm's strength, we create a trading strategy according to our pattern of return predictability, based on the method that Guo and Wachter (2025) develop, which involves no look-ahead bias and targets a constant portfolio volatility of 5%. We separate the strategy's return in each year into the component earned in the newsy months and that earned in the repetitive months. Table 8 regresses the two components on each other, both in raw returns and in ranks (to limit the influence of outliers). The significant positive coefficients confirm the association: a year in which the continuation effect is strong is also likely to see a strong reversal, as Prediction 3 states. The two arms are thus likely driven by the same underlying economic force—in the model, a persistent, time-varying p_x—and call for one story rather than two.

5.4 Timing within the quarter

Besides the earnings cycle, other macroeconomic fluctuations may also operate on a quarterly frequency. Notably, GDP announcements, although occurring monthly, are quarterly in nature; FOMC announcements are disproportionately unlikely to occur in the newsy months. Such macroeconomic announcements are shown to have a significant impact on the stock market (see, e.g., Ai et al. 2024 for a survey of the vast literature studying this impact), and

Hirshleifer and Sheng (2022) document heightened investor attention to earnings news on macroeconomic announcement days. We investigate whether our return pattern can potentially arise from these announcements (see Prediction 4). If macroeconomic announcements drive our return pattern, the pattern should be amplified on those macroeconomic announcement days. The first two bars of panel (b) of Figure 4 show that, if anything, the return pattern appears weaker on those announcement days.

In contrast, our return pattern appears to align with the earnings season, which does not immediately begin on the first day of the newsy months. This is because even the quickest announcers require some time to prepare their statements. Specifically, among companies with fiscal periods aligned with calendar quarters, only 0.13% of the on-cycle earnings announcements occur within the first week of the quarter. This percentage is significantly lower than the average weekly announcement rate of about 8%. In contrast, the second week of each quarter sees 1.87% of the announcements, a 15-fold increase from the previous week. Therefore, our pattern of return predictability, if stemming from these earnings announcements, is not expected to occur in the first week of newsy months. The last two bars of panel (b) of Figure 4 show that our return predictability pattern is indeed non-existent in the pre-season period.

5.5 The cross section of industries

Baseline return pattern The main intuition behind the theory is as follows. If an industry contains a sizable number of constituents, the group as a whole announces progressively along the earnings cycle, and earnings announced in the second month of a quarter will be predictably similar to those announced in the first month. Failure to recognize this inherent similarity leads to predictable overreaction in the second month that later reverses.

This story applies naturally to all stocks in the US economy, but should additionally apply to industry-level returns (see Prediction 5). Two stocks randomly drawn from an industry are more connected with each other than two random stocks drawn from an economy. Investors understand the tight intra-industry connection and actively learn industry-level information from the early earnings announcements (e.g., Foster (1981)). If the story is indeed about the

early earnings announcements, the dynamic serial predictive relation in the aggregate market returns should also exist in industry returns *in excess of* the aggregate market returns. At the level of individual stocks, Ertan et al. (2022) document continuation and reversal in returns in excess of the market around earnings announcement days; our mechanism concerns the aggregate earnings cycle, in which the newsy months function as aggregate "earnings announcements."

The first two columns of Table 9 test this prediction by conducting the following regression:

$$exret_{i,t} = \alpha + \beta exret_{i,nr(t)} + \epsilon_{i,t}. \tag{8}$$

Here $exret_{i,t}$ is the return of industry i in excess of the market return in month t. The returns are weighted by market capitalization, and the regression is weighted by industry-level market capitalization scaled to sum to 1 in each cross-section. Following Lou et al. (2019), we drop stocks with a smaller market capitalization than the 10th percentile on NYSE. We further drop industries with fewer than 10 constituents to ensure a sufficient number of announcers. Table 9 shows that the same pattern of continuation and reversal exists in industry excess returns as in aggregate market returns.

Two points are worth noting. First, this piece of cross-sectional evidence is important in its own right. The regressions use market capitalization weights and are unlikely to be driven by small stocks, making them substantially more implementable. Second, the consistency between aggregate and cross-sectional evidence is not always seen in the return predictability literature. This consistency supports our hypothesis.

Heterogeneity across industries Variation across industries offers further tests of our theory. As mentioned in the previous subsection, our theory applies to groups of stocks that are sizable in number. A sufficient number of constituents is necessary to set the stage for intra-group learning, generate the newsy and repetitive signals for group-level cash flow in the first and second months of a quarter, and allow for the possibility of correlation neglect. A natural implication of our theory is that the return pattern should be stronger for industries that have more constituents. In terms of the theory, size should increase the parameter p_x.

The greater is p_x, the more predictable are returns within the group of stocks.

Panel A1 of Table 10 tests this prediction. We split each cross-section into industries with high and low numbers of constituents, and run the regression on each half:

$$exret_{i,t} = \alpha + \beta_1 exret_{i,nr(t)} + \beta_2 exret_{i,nr(t)} \times I_t^r + \beta_3 \sum_{j=1}^{4} exret_{i,n(t,j)} \times I_t^r + \beta_4 \sum_{j=1}^{4} exret_{i,n(t,j)} + \gamma I_t^r + \epsilon_{i,t}. \tag{9}$$

The effect of correlation neglect, as measured by the size of the coefficient $-\beta_2$ on the interaction term, is much larger among industries with many constituents, consistent with Prediction 5. Panel A2 conducts an analogous exercise to those in Table 5 and shows that the industry-level effect of correlation neglect is larger when the industry earnings in the repetitive month more closely resemble those in the preceding newsy month. This is consistent with Prediction 1 and the time-series evidence in Table 5.

To further demonstrate that the empirical pattern relates to the economic structure of real industries, we conduct a placebo test in Panel B of Table 10 by randomly generating fake "industries" and conducting regression 9 on the panel of these fake industries. The randomization approaches follow those that Chen et al. (2024) proposes. Panel B1 assigns each stock to the industry of a randomly chosen stock in the same cross-section without replacement. This approach preserves the size distribution of the industries in each cross-section. Panel B2 first sorts each cross-section into deciles by $exret_{s,nr(t)}$, where s is a stock, and then assigns each stock to the industry of a randomly chosen stock in the same cross-section-decile. This approach preserves the size distribution of the industries in each cross-section and further controls for $exret_{i,nr(t)}$. Even the 1st percentiles in Panel B of Table 10 are much less negative than the high-group coefficients for genuine industries in Panel A of Table 10. The results in Panel B2 are especially useful. They tell us that the observed return pattern in Table 10 depends critically on the economic structure of genuine industries and is not the sole result of the regressor $exret_{i,nr(t)}$ itself.

6 Alternative explanations

6.1 Strategic timing of earnings announcements

A large accounting literature examines the relation between the timing of firms' earnings announcements and the news they convey. A clear empirical pattern that emerges is that early announcers tend to announce better news than late announcers (e.g., Kross 1981; Kross and Schroeder 1984; Chambers and Penman 1984; Johnson and So 2018b; Noh et al. 2021). The literature has also extensively studied the reasons behind this pattern and has partially attributed it to firms' endogenous choices of the announcing lag. For instance, deHaan et al. (2015) argue that firms delay earnings announcements with bad results to avoid the early portion of the earnings cycle that receives heightened attention. Givoly and Palmon (1982) argue that they do so to buy time so that they can manipulate their accounting results.

However, it is not clear how this empirical regularity alone speaks to our results. Overall, it implies that early announcers announce better results than late announcers. If investors fail to anticipate that, then early announcers will have higher earnings surprises and higher announcement excess returns. It is not clear why it would cause the market return in the newsy month to positively correlate with that in the next repetitive month. If anything, exogenous variation in the intensity of this self-selection effect seems to lead to a negative correlation of early announcement returns and late announcement returns within the quarter, as bad announcers being moved out of the newsy months makes the newsy months look better, and the subsequent repetitive months look worse than they otherwise would. And even then, it is not clear why a pattern of return reversal would exist between a repetitive month and the upcoming newsy month.

6.2 Time-varying risk

The most important set of explanations for return predictability comes from the predictable resolution of risks, which leads to high expected returns, and from predictably low-risk periods, which lead to low expected returns. The type of variation in risks that is necessary to

explain our time-series results is that, i) after a good newsy month, the stock market is risky in the subsequent repetitive month, and ii) after a good repetitive month, the subsequent newsy month is safe. Individually, both points are not difficult to accommodate in a risk-based framework. In fact, point ii) naturally arises when a repetitive month resolves risks and takes them away from the next newsy month. However, as we show in Table 8, we should ideally have one story that jointly explains both arms of the return pattern. This first requires the risk of the market to vary at a monthly frequency. Standard asset pricing models such as habit, long-run risk, disaster, and intermediary-based asset pricing all involve risks that should be rather persistent at a monthly frequency. Standard macro risk factors therefore appear unlikely to explain our predictability pattern.

Furthermore, even if one can identify a non-conventional, high-frequency source of risk, it is unlikely to account for our return pattern in full owing to the frequency at which the expected stock market excess returns are negative. The top panel of Figure 5 plots the expected stock market return implied by the out-of-sample forecast in Column 3 of Table 4. This time series of expected stock market excess return is negative in 24% of months. This implies that an explanation based on risk alone is unlikely to explain the return pattern in full. In such a risk-based framework, the negative expected returns mean that the stock market is safer than Treasury bills one month out of four. This feature is especially difficult to accommodate, as it is implausible that investors find the stock market to be safer than Treasury bills that often.

Campbell and Thompson (2008) show that the expanding-window regression coefficients used in the out-of-sample expected stock market excess returns incorporate substantial noise, which in turn leads to negative forecasts of expected stock market excess returns. Campbell and Thompson (2008) show that, for the stock market predictors they study, these negative forecasts perform poorly, consistent with the notion that they are driven by noise.

In contrast, negative forecasts in our paper do not appear to be driven by noise more than the positive forecasts. The lower panel of Figure 5 plots the in-sample expected stock market returns. These forecasts do not use the expanding-window coefficients, but are still negative 20% of the time. The first two columns of the top panel of Figure 4 decompose

the returns of the trading strategy in Table 8 by the sign of the out-of-sample excess return forecasts. The next two columns create an analogous in-sample trading strategy and conduct the decomposition based on the in-sample forecasts. The trading strategy does not appear to perform worse when the forecasts are negative. If anything, the performance appears to be better when the out-of-sample forecasts are negative. The negative forecasts in our paper as well as their implications should be taken seriously.

Another candidate is the risk resolution associated with earnings announcements. Savor and Wilson (2016) show that early announcers earn higher excess returns than later announcers (Table IV), which is indeed consistent with the overall importance of the earnings season. This, however, does not explain why the current earnings season being good makes the next repetitive month risky. Neither does it explain why a good repetitive month is followed by a safe newsy month.[22]

6.3 Different investors in newsy and repetitive months

A third possibility is reminiscent of a clientele effect: perhaps different investors see news in different months, just as, as Lou et al. (2019) emphasize, different types of investors may be active at different times of the day.[23] For example, suppose news diffuses gradually, as would occur if investors only pay attention to certain firms that announce in the first versus the second month. This story would imply continuation, but not reversal. That is, the second set of firms would be predictable based on the first set, and there would be no over-reaction. However, we find strong evidence of reversal, ruling out this explanation.

An alternative version is that less sophisticated investors miss the newsy month entirely, and focus on the repetitive month's announcements. These investors do not realize that the market has already incorporated the information. This would explain continuation and reversal, but it is not clear why less sophisticated investors would pay more attention to late-announcing firms.

[22]The mechanism in Savor and Wilson (2016) would however amplify our results in a portfolio of announcing firms, and that appears to be true in the data.

[23]Lou et al. (2019) also document continuation and reversal, but the pattern is very different. In their interpretation, two investor groups exert persistent, opposing price pressures: each group's demand continues in its own trading window and offsets the other's, so that the pattern consists of two continuations that cancel.

A final version is that less sophisticated investors see the announcement in both months, and price the signals on top of what the market has already incorporated. This version produces both continuation and reversal, but it is in fact correlation neglect. Investors see a dependent signal, but nonetheless update their beliefs based on it. Moreover, the focus on "less sophisticated" investors avoids the question of why the more sophisticated investors do not correct the mispricing, unless, as Section 5.2 suggests, they suffer from the bias themselves to some degree.

7 Conclusion

Contrary to prior beliefs, monthly stock market returns in the United States can in fact be predicted with past returns. Specifically, the U.S. stock market's return during the first month of a quarter positively predicts the second month's return, which then negatively predicts the first month's return of the next quarter. The first months of a quarter are "newsy" because they contain fresh earnings news about the aggregate economy, whereas the second months of a quarter contain earnings announcements that produce predictably repetitive "news." We hypothesize that the return predictability pattern arises because investors use announced earnings to predict future earnings, but do not recognize that earnings in the second months of a quarter are inherently repetitive of those in the previous newsy month. Consistent with this hypothesis of correlation neglect, the return pattern amplifies when the earnings in the first and second months of a quarter are more similar. Survey data further support this hypothesis, as does evidence across industries. These results challenge the efficient market hypothesis by documenting a strong and pervasive form of return predictability.

References

Ai, H., Bansal, R., and Guo, H. (2024). Macroeconomic announcement premium. In *Oxford research encyclopedia of economics and finance*. Oxford University Press, Oxford, England.

Akerlof, G. A. and Shiller, R. J. (2009). *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*. Princeton University Press, Princeton, NJ.

Barth, M. E., Li, K., and McClure, C. (2024). Information supporting investor valuations: Evidence from a comparative content analysis of analyst reports and form 10-K. Working Paper, SSRN, https://papers.ssrn.com/abstract=4763765, Accessed 2024-09-01.

Bohren, A. (2016). Informational herding with model misspecification. *Journal of Economic Theory*, 163:222–247.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2024). Belief overreaction and stock market puzzles. *Journal of Political Economy*, 132(5):1450–1484.

Bouchaud, J.-P., Krüger, P., Landier, A., and Thesmar, D. (2019). Sticky expectations and the profitability anomaly. *Journal of Finance*, 74(2):639–674.

Brochet, F., Kolev, K., and Lerman, A. (2018). Information transfer and conference calls. *Review of Accounting Studies*, 23(3):907–957.

Brusa, F., Savor, P., and Wilson, M. (2020). One central bank to rule them all. *Review of Finance*, 24(2):263–304.

Campbell, J. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531.

Chambers, A. and Penman, S. (1984). Timeliness of reporting and the stock-price reaction to earnings announcements. *Journal of Accounting Research*, 22(1):21–47.

Charles, C. (2025). Memory moves markets. *Review of Financial Studies*, 38(6):1641–1686.

Chen, H., Dou, W. W., Guo, H., and Ji, Y. (2024). Industry Distress Anomaly. Working paper, SSRN 4738806.

Cohen, R. T. and Kahana, M. J. (2022). A memory-based theory of emotional disorders. *Psychological Review*, 129(4):742–776.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Da, Z., Gurun, U. G., and Warachka, M. (2014). Frog in the Pan: Continuous Information and Momentum. *The Review of Financial Studies*, 27(7):2171–2218.

De La O, R. and Myers, S. (2021). Subjective cash flow and discount rate expectations. *Journal of Finance*, 76(3):1339–1387.

deHaan, E., Shevlin, T., and Thornock, J. (2015). Market (in)attention and the strategic scheduling and timing of earnings announcements. *Journal of Accounting and Economics*, 60(1):36–55.

Dellavigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *Journal of Finance*, 64(2):709–749.

DeMarzo, P. M., Vayanos, D., and Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118(3):909–968.

Enke, B. and Zimmermann, F. (2019). Correlation neglect in belief formation. *Review of Economic Studies*, 86(1):313–332.

Ertan, A., Karolyi, S. A., Kelly, P. W., and Stoumbos, R. (2022). Earnings announcement return extrapolation. *Review of Accounting Studies*, 27(1):185–230.

Eyster, E. and Rabin, M. (2010). Naïve herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–243.

Fama, E. F. (1965). The behavior of stock-market prices. *Journal of Business*, 38(1):34–105.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417.

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

Fedyk, A. and Hodson, J. (2023). When can the market identify old news? *Journal of Financial Economics*, 149(1):92–113.

Foster, G. (1981). Intra-industry information transfers associated with earnings releases. *Journal of Accounting and Economics*, 3(3):201–232.

Gabaix, X. (2019). Behavioral inattention. In Bernheim, B. D., DellaVigna, S., and Laibson, D., editors, *Handbook of behavioral economics: Applications and foundations 1*, volume 2, pages 261–343. Elsevier, Amsterdam.

Gilbert, T., Kogan, S., Lochstoer, L., and Ozyildirim, A. (2012). Investor inattention and the market impact of summary statistics. *Management Science*, 58(2):336–350.

Givoly, D. and Palmon, D. (1982). Timeliness of annual earnings announcements: Some empirical evidence. *Accounting Review*, 57(3):486–508.

Goetzmann, W. N., Kim, D., and Shiller, R. J. (2024). Crash narratives. NBER Working Paper No. 30195.

Golub, B. and Jackson, M. O. (2010). Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149.

Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508.

Goyal, A., Welch, I., and Zafirov, A. (2024). A comprehensive 2022 look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 37(11):3490–3557.

Guo, H. (2025). Earnings Extrapolation and Predictable Stock Market Returns. *The Review of Financial Studies*, 38(6):1730–1782.

Guo, H. and Wachter, J. A. (2025). Forecast-Agnostic Portfolios. SSRN Scholarly Paper 5808182, Social Science Research Network, Rochester, NY.

Hann, R. N., Kim, H., and Zheng, Y. (2019). Intra-industry information transfers: Evidence from changes in implied volatility around earnings announcements. *Review of Accounting Studies*, 24(3):927–971.

Hilary, G. and Shen, R. (2013). The Role of Analysts in Intra-Industry Information Transfer. *The Accounting Review*, 88(4):1265–1287.

Hirshleifer, D. and Sheng, J. (2022). Macro news and micro news: Complements or substitutes? *Journal of Financial Economics*, 145(3):1006–1024.

Howard, M. W. and Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3):269–299.

Johnson, T. L. and So, E. C. (2018a). Asymmetric trading costs prior to earnings announcements: Implications for price discovery and returns. *Journal of Accounting Research*, 56(1):217–263.

Johnson, T. L. and So, E. C. (2018b). Time will tell: Information in the timing of scheduled earnings news. *Journal of Financial and Quantitative Analysis*, 53(6):2431–2464.

Kahana, M. J., Diamond, N. B., and Aka, A. (2024). Laws of human memory. In Kahana, M. J. and Wagner, A. D., editors, *The Oxford Handbook of Human Memory, Volume 1: Foundations*, pages 29–63. Oxford University Press.

Kahana, M. J., Paron, J., and Wachter, J. (2025). Associative learning and representativeness. Working paper, University of Pennsylvania.

Kendall, M. G. and Hill, A. B. (1953). The analysis of economic time-series, part I: Prices. *Journal of the Royal Statistical Society. Series A (General)*, 116(1):11–34.

Kraft, A. G., Vashishtha, R., and Venkatachalam, M. (2017). Frequent financial reporting and managerial myopia. *Accounting Review*, 93(2):249–275.

Kross, W. (1981). Earnings and announcement time lags. *Journal of Business Research*, 9(3):267–281.

Kross, W. and Schroeder, D. A. (1984). An empirical investigation of the effect of quarterly earnings announcement timing on stock returns. *Journal of Accounting Research*, 22(1):153–176.

Levy, G. and Razin, R. (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review*, 105(4):1634–1645.

Lewellen, J. and Shanken, J. (2002). Learning, asset-pricing tests, and market efficiency. *Journal of Finance*, 57:1113–1145.

Li, K., Mai, F., Shen, R., Yang, C., and Zhang, T. (2024). Dissecting corporate culture using generative AI: Insights from analyst reports. Working Paper, SSRN, https://papers.ssrn.com/abstract=4558295, Accessed 2024-09-01.

Lo, A. W. and MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies*, 1(1):41–66.

Lou, D., Polk, C., and Skouras, S. (2019). A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134:192–213.

Lou, D., Polk, C., and Skouras, S. (2026). The day destroys the night, night extends the day: A clientele perspective on equity premium variation. Working paper, London School of Economics.

Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.

Nagel, S. and Xu, Z. (2022). Asset pricing with fading memory. *Review of Financial Studies*, 35(5):2190–2245.

Noh, S., So, E. C., and Verdi, R. S. (2021). Calendar rotations: A new approach for studying the impact of timing using earnings announcements. *Journal of Financial Economics*, 140(3):865–893.

Ortoleva, P. and Snowberg, E. (2015). Overconfidence in political behavior. *American Economic Review*, 105(2):504–535.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices: Evidence and Implications. *Journal of Financial Economics*, 22(1):27–59.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2010). Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy. *The Review of Financial Studies*, 23(2):821–862.

Rapach, D. E. and Zhou, G. (2022). Asset pricing: Time-series predictability. In *Oxford research encyclopedia of economics and finance*. Oxford University Press, Oxford.

Rees-Jones, A., Shorrer, R., and Tergiman, C. (2024). Correlation Neglect in Student-to-School Matching. *American Economic Journal: Microeconomics*, 16(3):1–42.

Savor, P. and Wilson, M. (2016). Earnings announcements and systematic risk. *Journal of Finance*, 71(1):83–138.

Tetlock, P. C. (2011). All the news that's fit to reprint: Do investors react to stale information? *Review of Financial Studies*, 24(5):1481–1512.

Thomas, J. and Zhang, F. (2008). Overreaction to intra-industry information transfers? *Journal of Accounting Research*, 46(4):909–940.

Wachter, J. A. and Kahana, M. J. (2024). A retrieved-context theory of financial decisions. *The Quarterly Journal of Economics*, 139(2):1095–1147.

Wu, L. (2023). Mental macro-finance models: Evidence and theory. Working paper, New York University.

Table 1: When do firms end their fiscal quarters and when do they announce earnings

	A: FQ-end month		B: Earnings announcement month	
	Count	Percent	Count	Percent
Group 1: Jan/Apr/Jul/Oct	82,257	8.6%	358,128	44.7%
Group 2: Feb/May/Aug/Nov	51,169	5.4%	382,162	47.7%
Group 3: Mar/Jun/Sep/Dec	819,534	86.0%	60,649	7.6%
Total	952,960	100.0%	800,939	100.0%

Panel A counts the company-fiscal quarters in Compustat by the months in which they end. The fiscal quarters always end on the last day of a month, and 86.0% means that 86.0% of the fiscal quarters end on the last day of March, June, September, or December. Panel B counts the company-fiscal quarters by the months in which the fiscal quarters' earnings are announced, requiring that i) the fiscal quarter ends in the Group 3 months and ii) the earnings announcement occurs within 92 days of the fiscal quarter end date. Data are quarterly from 1971 to 2023.

Table 2: Continuation and reversal in the U.S. aggregate market returns

	(1) Reversal	(2) Continuation	(3) Difference
	A: Full sample		
$mkt_{nr(t)}$	-0.277***	0.271***	-0.548***
	[-3.04]	[3.03]	[-4.29]
N	395	395	790
	B: First half		
$mkt_{nr(t)}$	-0.328**	0.358***	-0.686***
	[-2.49]	[2.71]	[-3.68]
N	192	192	384
	C: Second half		
$mkt_{nr(t)}$	-0.191**	0.149**	-0.341***
	[-2.16]	[2.45]	[-3.17]
N	203	203	406
	D: Post-WWII		
$mkt_{nr(t)}$	-0.162**	0.129**	-0.291***
	[-2.34]	[2.42]	[-3.33]
N	311	311	622

The first two columns of this table present estimated βs from the monthly time-series regression $mkt_t = \alpha + \beta mkt_{nr(t)} + \epsilon_t$. Here, mkt_t is the U.S. stock market return in excess of the risk free rate in month t, and $nr(t)$ is the most recent month before t that is newsy or repetitive. Column 1 is on the sample in which the dependent variable month t is a first month of a quarter (newsy); column 2 is on the sample in which the dependent variable month t is a second month of a quarter (repetitive). Column 3 shows their difference, computed as $-\beta_2$ in the regression $mkt_t = \alpha + \beta_1 mkt_{nr(t)} + \beta_2 mkt_{nr(t)} \times I_t^r + \epsilon_t$, where I_t^r is a dummy variable indicating a repetitive t. In Panel A, B, C, and D, data are monthly from 1926 to 2024, from 1926 to 1973, from 1974 to 2024, and 1947 to 2024, respectively. T-statistics computed with White standard errors are reported in square brackets.

Table 3: Controlling for the preceding predictor

	(1) Reversal	(2) Continuation	(3) Difference
$mkt_{nr(t)}$	-0.247***	0.207**	-0.454***
	[-2.73]	[2.15]	[-3.44]
$\sum_{j=1}^{4} mkt_{n(t,j)}$	-0.062*	0.078**	-0.140***
	[-1.87]	[2.38]	[-3.00]
N	392	392	784

The first two columns of this table present estimated βs from the monthly time-series regression $mkt_t = \alpha + \beta mkt_{nr(t)} + \gamma \sum_{j=1}^{4} mkt_{n(t,j)} + \epsilon_t$. Here, mkt_t is the U.S. stock market return in excess of the risk free rate in month t, $nr(t)$ is the most recent month before t that is newsy or repetitive, and $mkt_{n(t,j)}$ is the jth most recent newsy month before t. Column 1 is on the sample in which the dependent variable month t is a first month of a quarter (newsy); column 2 is on the sample in which the dependent variable month t is a second month of a quarter (repetitive). Column 3 shows their difference, computed as $-\beta_2$ and $-\beta_3$ in the regression $mkt_t = \alpha + \beta_1 mkt_{nr(t)} + \beta_2 mkt_{nr(t)} \times I_t^r + \beta_3 \sum_{j=1}^{4} mkt_{n(t,j)} \times I_t^r + \beta_4 \sum_{j=1}^{4} mkt_{n(t,j)} + \gamma I_t^r + \epsilon_t$, where I_t^r is a dummy variable indicating a repetitive t. Data are monthly from 1926 to 2024. T-statistics computed with White standard errors are reported in square brackets.

Table 4: Out-of-sample R^2 in forecasting the U.S. stock market returns

	(1) Rvsl only	(2) Cont only	(3) Rvsl + Cont	(4) Guo 2025	(5) Both
Out-of-sample R^2	2.75%	3.20%	5.95%	3.68%	6.95%

The R^2 in this table are calculated as $1 - \frac{\sum_{t=1}^{n}(mkt_t - \hat{mkt}_{t-1})^2}{\sum_{t=1}^{n}(mkt_t - \overline{mkt}_{t-1})^2}$, where $\overline{mkt}_{t-1}$ is the expanding window mean of past stock market excess returns up to month $t-1$ and $\hat{mkt}_{t-1}$ is the forecast being evaluated, also formed with information available at the end of month $t-1$. The computation is done where the target month t is either newsy or repetitive, and the predictor z_{t-1} underlying $\hat{mkt}_{t-1}$ is $mkt_{nr(t)} - \overline{mkt}_{t-1}$ for a repetitive t, $-(mkt_{nr(t)} - \overline{mkt}_{t-1})$ for a newsy t. For each target month t, 2 coefficients $\hat{\beta}_{t-1}$ and $\hat{\alpha}_{t-1}$—the former for the predictor and latter for the constant term—are extracted from a simple expanding-window OLS regression of historical market excess returns on historical predictor values and a constant using data up to $t-1$. In the first column, the forecast is $\overline{mkt}_{t-1}$ when t is repetitive, and $\overline{mkt}_{t-1} + \hat{\beta}_{t-1} z_{t-1}$ when t is newsy. In the second column, the forecast is $\overline{mkt}_{t-1} + \hat{\beta}_{t-1} z_{t-1}$ when t is repetitive, and $\overline{mkt}_{t-1}$ when t is newsy. In the third column, the forecast is $\overline{mkt}_{t-1} + \hat{\beta}_{t-1} z_{t-1}$ in all cases. The fourth column is the same as the third, except that the predictor is $\sum_{j=1}^{4} mkt_{n(t,j)}$ subtracting its expanding window mean for a repetitive t, and the negative of that for a newsy t. In the last column, the forecast is the average of the two previous columns. Data are monthly from 1926 to 2024.

Table 5: Aggregate earnings distance and the strength of the return pattern

	(1) Lowest 10%	(2) 10%-50%	(3) 50%-90%	(4) Highest 10%	(5) Low - High
Diff. w/ control	-0.800***	-0.281	-0.016	-0.127	-0.673**
	[-4.37]	[-1.39]	[-0.09]	[-0.48]	[-2.11]
N	42	170	170	42	84

This table presents estimated $-\beta_2$s from the monthly time-series regression $mkt_t = \alpha + \beta_1 mkt_{nr(t)} + \beta_2 mkt_{nr(t)} \times I_t^r + \beta_3 \sum_{j=1}^{4} mkt_{n(t,j)} \times I_t^r + \beta_4 \sum_{j=1}^{4} mkt_{n(t,j)} + \gamma I_t^r + \epsilon_t$. Here, mkt_t is the U.S. stock market return in excess of the risk free rate in month t, $nr(t)$ is the most recent month before t that is newsy or repetitive, I_t^r is a dummy variable taking the value of 1 if month t is repetitive and 0 otherwise, and $mkt_{n(t,j)}$ is the jth newsy month before month t. Across columns, the regression sorts on the distance between the aggregate ROE announced in the most recent newsy months before t and the subsequent repetitive month. The aggregate ROE in a given month is defined as the total earnings announced in the month divided by the total book value of equity of these announcing firms. The ROE distance is computed as the absolute value of the difference between the newsy and the repetitive months' ROE relative to the median value of the difference across time. The distance increases from column 1 to 4 and is between the 0-10th, 10-50th, 50-90th, and 90-100th percentiles of the sample, respectively. Data are monthly from 1971-2024. T-statistics computed with White standard errors are reported in square brackets.

Table 6: Return pattern and non-fundamental drivers in the predictor

	Difference		Diff w/ control	
Non-funda. drivers	Small	Large	Small	Large
$-\beta_2$	-0.383***	-0.113	-0.378***	0.021
	[-3.68]	[-0.41]	[-3.69]	[0.08]
N	172	167	172	167

The first two columns of this table present estimated $-\beta_2$s from the monthly time-series regression $mkt_t = \alpha + \beta_1 mkt_{nr(t)} + \beta_2 mkt_{nr(t)} \times I_t^r + \gamma I_t^r + \epsilon_t$. Here, mkt_t is the U.S. stock market return in excess of the risk free rate in month t, $nr(t)$ is the most recent month before t that is newsy or repetitive, and I_t^r is a dummy variable indicating whether t is repetitive. The last two columns additionally control for $\sum_{j=1}^{4} mkt_{n(t,j)} \times I_t^r$ and $\sum_{j=1}^{4} mkt_{n(t,j)}$, where $\sum_{j=1}^{4} mkt_{n(t,j)}$ is the sum of returns of the past four newsy months strictly before t. The first and the third columns are on the sub-sample in which $mkt_{nr(t)}$ is less driven by non-fundamental forces than a median month, whereas the second and the fourth columns are on the other half of the sample in which $mkt_{nr(t)}$ is more driven by non-fundamental forces than a median month. The exact measurement of non-fundamental drivers broadly follows the method in Da et al. (2014) and is described in the last paragraph of Appendix A. Data are monthly from 1982 to 2024. T-statistics computed with White standard errors are reported in square brackets.

Table 7: Evidence from survey data

	(1) Rev_t	(2) Rev_t	(3) Rev_t	(4) Rev_t
Rev_{t-1}	0.543***	0.220***		
	[4.91]	[2.94]		
I_t^r		0.078		
		[1.15]		
$Rev_{t-1} \times I_t^r$		0.323***		
		[3.42]		
Rev_{t-2}			0.038	0.211***
			[1.02]	[3.92]
I_t^n				-0.052
				[-1.05]
$Rev_{t-2} \times I_t^n$				-0.174***
				[-2.72]
const	-0.032	-0.110**	-0.143***	-0.091*
	[-0.81]	[-2.03]	[-3.04]	[-1.93]
N	170	509	170	509

Column 1 runs the following monthly time-series regression on the sample in which t is repetitive: $Rev_t = \alpha + \beta Rev_{t-1} + \epsilon_t$. Here, Rev_t is the aggregate revision in the IBES long-term growth (LTG) measure for month t within the S&P 500 universe using the market capitalization weight. Column 2 adds in a dummy variable indicating whether t is repetitive and its interaction with Rev_{t-1}, and expands the sample to all months. Columns 3 and 4 conduct analogous regressions to columns 1 and 2, but for newsy months rather than repetitive months. Data are monthly from 1982 to 2024. T-statistics computed with Newey-West standard errors are reported in square brackets.

Table 8: The link between the continuation and reversal effects

	(1) TSP_t^n	(2) TSP_t^r	(3) $Rank(TSP_t^n)$	(4) $Rank(TSP_t^r)$
TSP_t^r	0.957***			
	[15.16]			
TSP_t^n		0.893***		
		[6.87]		
$Rank(TSP_t^r)$			0.361***	
			[3.42]	
$Rank(TSP_t^n)$				0.361***
				[3.40]
N	97	97	97	97

Column 1 shows results from the following annual time-series regression: $TSP_t^n = \alpha + \beta TSP_t^r + \epsilon_t$. Here, TSP_t is our time-series portfolio return in year t. This portfolio takes long or short positions in the aggregate market for a given month, and the weight is proportional to the previous-month-end predictor value s_{t-1}, which is $mkt_{nr(t)} - \overline{mkt_{t-1}}$ for a repetitive t, $-(mkt_{nr(t)} - \overline{mkt_{t-1}})$ for a newsy t, and 0 otherwise. The market weight is then scaled so that the strategy has a historical volatility of 5% as of $t-1$, and then truncated at 3 and -3. TSP_t^n is the portfolio's return in the newsy months of year t, and TSP_t^r is the return in the repetitive months of year t. The $Rank(\cdot)$ function ranks the TSP_t^r and TSP_t^n across years, thus transforming them into integers between 1 and 97. T-statistics computed with White standard errors are reported in the square brackets.

Table 9: Continuation and reversal in the U.S. industry excess returns

	Reversal	Continuation	Difference	Diff w/ control
		A: Full sample		
$exret_{i,nr(t)}$	-0.079**	0.086***	-0.165***	-0.146***
	[-2.37]	[2.82]	[-3.65]	[-2.97]
N	16,647	16,631	33,278	33,077
		B: First half		
$exret_{i,nr(t)}$	-0.067	0.085*	-0.152**	-0.150*
	[-1.36]	[1.73]	[-2.19]	[-1.91]
	4,107	4,049	8,156	7,969
		C: Second half		
$exret_{i,nr(t)}$	-0.085*	0.086**	-0.172***	-0.142**
	[-1.96]	[2.23]	[-2.95]	[-2.28]
	12,540	12,582	25,122	25,108
		D: Post-WWII		
$exret_{i,nr(t)}$	-0.065*	0.075**	-0.141***	-0.116**
	[-1.74]	[2.33]	[-2.84]	[-2.19]
	15,466	15,461	30,927	30,781

The first two columns of this table present estimated βs from the industry-monthly panel regression $exret_{i,t} = \alpha + \beta exret_{i,nr(t)} + \epsilon_{i,t}$. Here, $exret_{i,t}$ is the return of industry i in excess of the market return in month t, and $nr(t)$ is the most recent month before t that is newsy or repetitive. Column 1 is on the sample in which the dependent variable month t is a first month of a quarter (newsy); column 2 is on the sample in which the dependent variable month t is a second month of a quarter (repetitive). Column 3 shows the difference, extracted as $-\beta_2$ from the regression of $exret_{i,t} = \alpha + \beta_1 exret_{i,nr(t)} + \beta_2 exret_{i,nr(t)} \times I_t^r + \gamma I_t^r + \epsilon_t$ on the combined sample in columns 1 and 2, where I_t^r is a dummy variable taking the value of 1 if the month t is repetitive, and 0 otherwise. The last column additionally controls for the sum of industry excess returns of the past four newsy months strictly before t and its interaction with I_t^r. The returns are weighted by market capitalization, and the regression is weighted by industry-level market capitalization scaled to sum to 1 in each cross-section. The regressions require at least 10 firms in an industry. In Panels A, B, C, and D, data are monthly from 1926 to 2024, from 1926 to 1973, from 1974 to 2024, and from 1947 to 2024, respectively. T-statistics computed with standard errors clustered by month are reported in square brackets.

Table 10: Further analyses across industries

	A: Heterogeneity across industries			B: Random "industries"				
	Low	High	High-Low	1%	5%	50%	95%	99%
	A1: Size			B1: Preserve industry sizes				
Diff w/ control	-0.048*	-0.126***	-0.078**	-0.056	-0.041	-0.002	0.035	0.051
	[-1.67]	[-3.01]	[-2.38]					
	A2: Earnings similarity			B2: Control for $exret_{i,nr(t)}$				
Diff w/ control	-0.079	-0.205***	-0.126**	-0.067	-0.049	-0.014	0.024	0.040
	[-1.03]	[-10.50]	[-2.04]					

Panel A of this table presents estimated $-\beta_2$s from the industry-monthly panel regression: $exret_{i,t} = \alpha + \beta_1 exret_{i,nr(t)} + \beta_2 exret_{i,nr(t)} \times I_t^r + \beta_3 \sum_{j=1}^{4} exret_{i,n(t,j)} \times I_t^r + \beta_4 \sum_{j=1}^{4} exret_{i,n(t,j)} + \gamma I_t^r + \epsilon_t$. Here, $exret_{i,t}$ is the return of industry i in excess of market return in month t, $nr(t)$ is the most recent month before t that is newsy or repetitive, $\sum_{j=1}^{4} exret_{i,n(t,j)}$ is the sum of the past four newsy month excess returns of industry i, and I_t^r is a dummy variable taking the value of 1 if month t is repetitive, and 0 otherwise. The returns are weighted by market capitalization, and the regression is weighted by industry-level market capitalization scaled to sum to 1 in each cross-section. The regressions require at least 10 firms in an industry. Panel A1 separates the regression by whether an industry has below- or above-average number of constituents and shows the difference in β_2 in the 3rd column. Panel A2 does so for earnings similarity, computed as the absolute value of the distance between the industry's ROE announced in the newsy month before t and the subsequent repetitive month. Panel B conducts 2 placebo tests on randomly generated, fake "industries," and reports the 1st, 5th, 50th, 95th, and 99th percentiles of the distribution of $-\beta_2$ generated in 1000 simulations. The simulation approaches follow those in Chen et al. (2024). Data are monthly from 1971-2024. T-statistics computed with standard errors clustered at monthly level are reported in square brackets.

Figure 2: Repetitive month returns' loadings on past newsy months

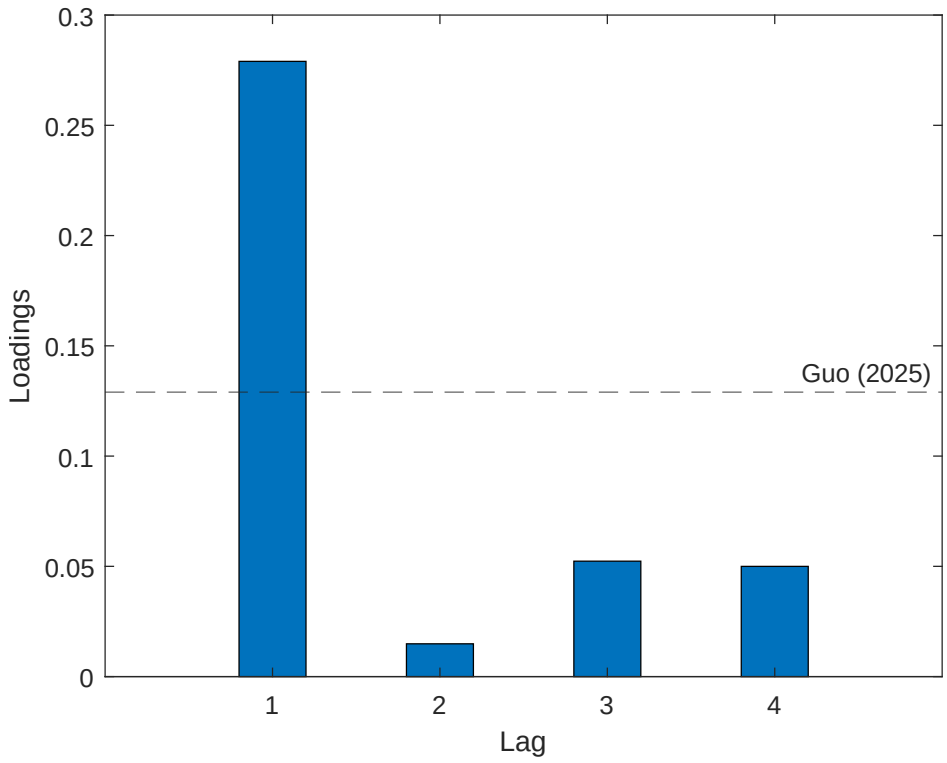

This figure plots the coefficients in a regression of repetitive month returns on the past four newsy months and a constant. The coefficient on the constant is omitted.

Figure 3: OOS R^2s in predicting monthly stock market returns

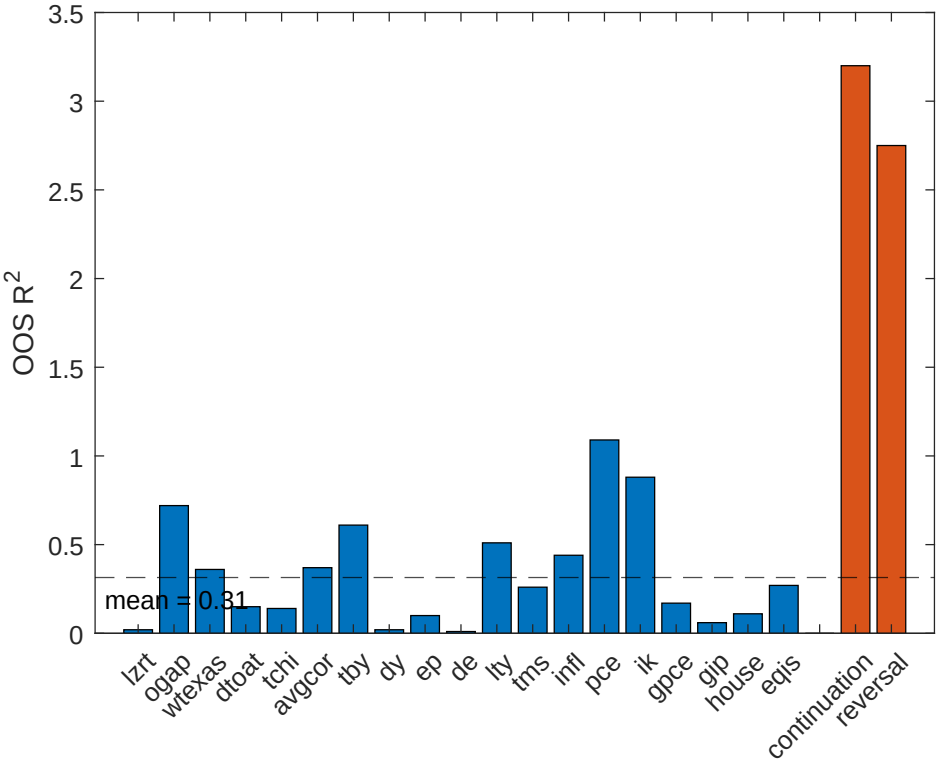

This figure plots the 19 positive monthly OOS R^2s generated by the comprehensive list of predictors in Goyal et al. (2024) as well as those generated by the continuation and reversal arms of our return pattern.

Figure 4: Trading strategy performance decomposition

(a) Signs of the forecasts

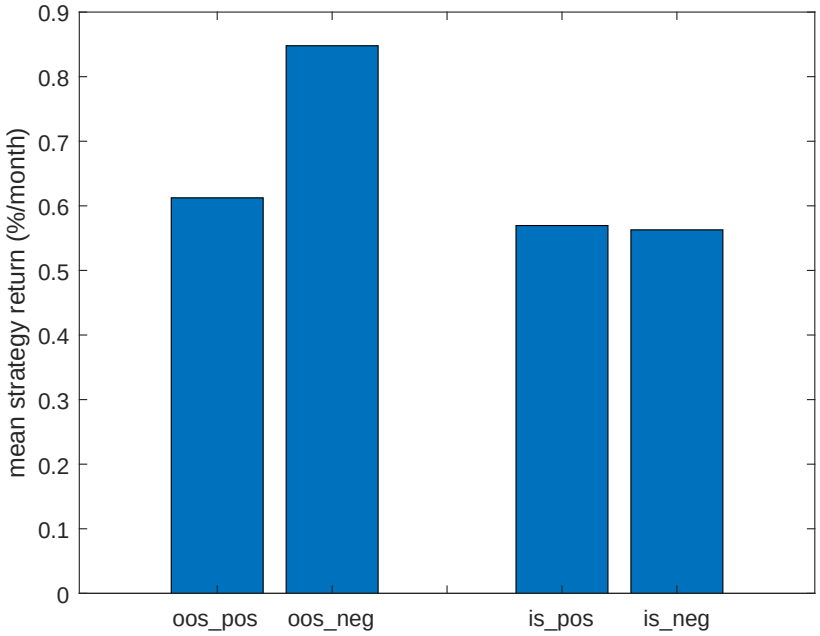

(b) Macro announcement days and the pre-season periods

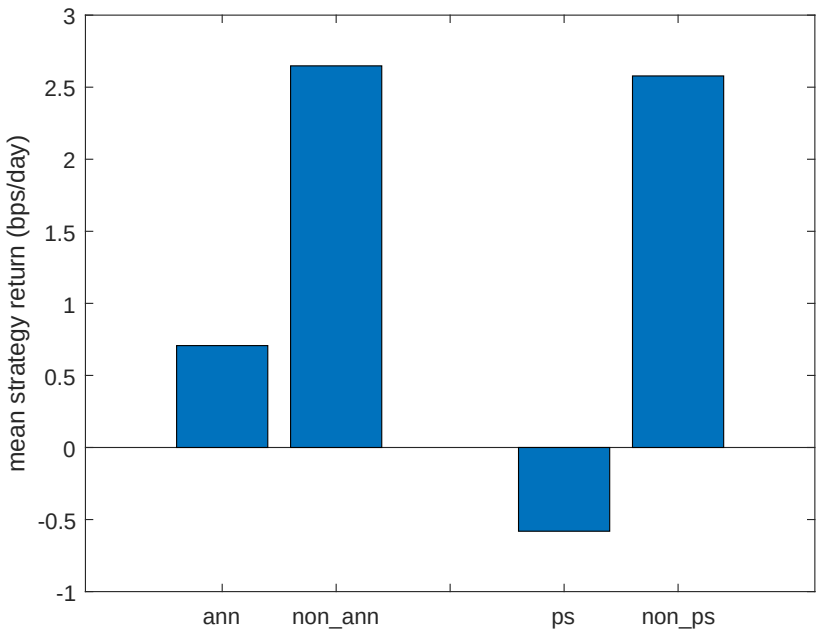

The top panel of the figure decomposes the average monthly return of the out-of-sample and in-sample trading strategies (as in Table 8 and created based on the in-sample return predictability pattern in Table 2) by the sign of the out-of-sample and in-sample stock market excess return forecasts, respectively. The bottom panel of the figure plots the average daily return of a trading strategy constructed based on the in-sample return predictability pattern in Table 2, normalized to have a volatility of 5% per month. The average returns are computed on macroeconomic announcement days, non-announcement days, the pre-earnings season periods (i.e., the first week of a quarter), and the non-preseason periods.

Figure 5: Implied expected stock market excess returns over time

(a) Out-of-sample

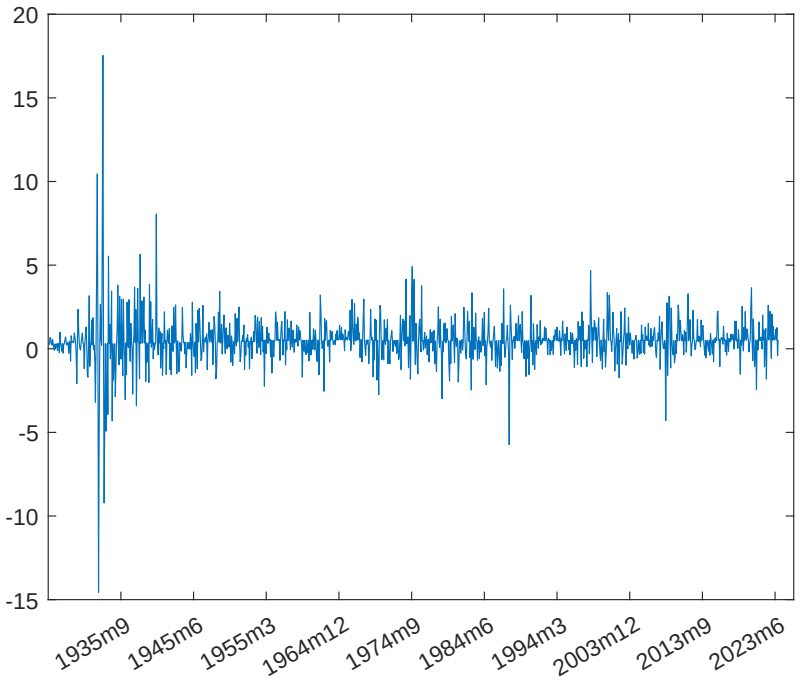

(b) In-sample

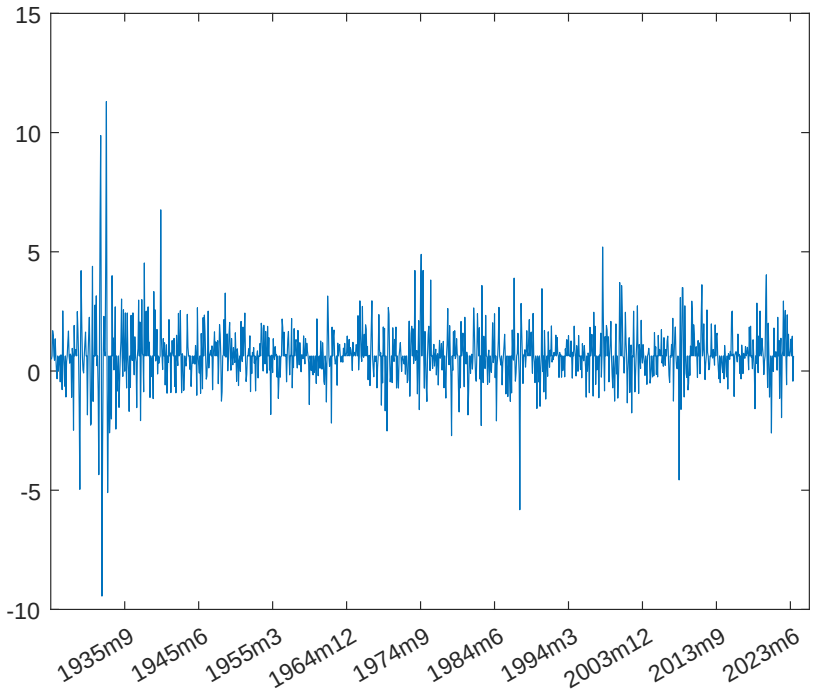

The top panel of this figure plots the expected stock market excess returns from the out-of-sample forecast in Column 3 of Table 4. The bottom panel of this figure plots the in-sample expected market excess returns. The unit is percent per month.

Appendix

A Measurement of non-fundamental shocks

This appendix describes the construction of the daily non-fundamental shock series and the monthly measure that Table 6 uses. In the model, the return pattern operates through announced earnings; a month whose return is driven by shocks outside earnings announcements carries little of the information that the repetitive month repeats.

To measure whether our predictors (i.e., market returns in a given newsy or repetitive month) are driven by non-fundamental forces, we follow the approach in Da et al. (2014). We start by creating a tangible-earnings-news-implied[24] return series using the learning model in Savor and Wilson (2016). We first aggregate earnings revisions on each day, weighting by market capitalization of the underlying firms, to create a measure of daily earnings news. We then aggregate these daily revisions to quarters. Quarterly earnings are the shared systematic component of these daily revisions, which can be gradually learned over time. This systematic component should enter aggregate returns approximately one-for-one, by the Campbell–Shiller decomposition.[25] The learning coefficient in a given day M of the quarter is $\frac{1}{N}\frac{N\sigma_\eta^2+\sigma_v^2}{M\sigma_\eta^2+\sigma_v^2}$, where N is the total number of trading days in the quarter, σ_η^2 is the variance of the systematic component, and σ_v^2 is that of the idiosyncratic component of earnings news. The tangible-earnings-news-implied return is this coefficient multiplied by the daily earnings revision.

We then take the difference between the market return and the earnings-implied return to arrive at a daily non-fundamental shock series nfs_d. Then, following Da et al. (2014), we measure the degree to which monthly market return mkt_t is driven by these non-fundamental shocks as:

$$sign(mkt_t)\frac{\sum_{d\in t} 1(nfs_d < 0) - \sum_{d\in t} 1(nfs_d > 0))}{\sum_{d\in t} 1}$$

This value is between -1 and 1. It is high when mkt_t is directionally driven by a few large outliers. For example, if two large negative days outweighed the other 18 positive days in a given month to make the monthly return negative, the measure will take a high value of 0.8. Our assumption in this exercise is that such daily outliers are unlikely to be fundamental-related. They lower p_x by confounding with fundamental shocks. If mkt_t is driven by them, it is likely to work less well as a predictor.

[24] This is an incomplete measure of fundamental driven returns. Bordalo et al. (2024), for example, argue that most earnings shocks are likely intangible.

[25] $r_{t+1} - E_t[r_{t+1}] = (E_{t+1} - E_t)\sum_{j=0}^{\infty}\rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t)\sum_{j=1}^{\infty}\rho^j r_{t+1+j}$. Since LTG measures growth 3–5 years away, ρ^j for the corresponding horizons ought to be close to 1.

B Memory model

This section shows how the model in Section 4 arises naturally from a cognitive model of contextual encoding and retrieval: the retrieved-context theory of Howard and Kahana (2002), as applied to financial decisions in Wachter and Kahana (2024). Agents view features f_t, which are basis vectors in a high-dimensional vector space. These features become embedded in memory together with a temporal context, an internal mental state. Context, through the memory matrix, governs what comes to mind. We consider a simplified version of the model, in which contexts are basis vectors. Contexts decay sufficiently rapidly so that each month's context can be viewed as independent of prior months; however, firms that announce within a month are all aggregated into the same temporal context. The approach of assuming contexts are either independent or the same is also adopted by Charles (2025), and is convenient in that it enables contextual encoding and retrieval without the analytical difficulties of modeling the contextual decay process.

It suffices to consider only part of the features space: the part corresponding to firm-level announcements that the agent interprets as good or bad news for the aggregate market. Thus contexts c_t are a series of independent basis vectors, and features are a series of vectors which, without loss of generality, we can assume to equal $[1, 0]$ for a positive announcement, and $[0, 1]$ for a negative one.[26]

Contexts and features form an outer product, whose sums make up the memory matrix. The matrix M_0 captures prior information: the uniform prior corresponds to two prior episodes, a positive announcement and a negative announcement, each encoded in its own context. Then memory evolves recursively as:

$$M_t = M_{t-1} + c_t f_t^\top.$$

Under the assumption that $\{c_t\}$ are basis vectors, independent of each other and of the prior-episode contexts, and positive and negative features are as specified above, we find:

$$M_0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

[26] More broadly, these features form a subspace of a very high dimensional space on the order of 10^7 or greater, 10^7 being a lower bound on the number of neurons in the brain involved in memory storage. As in Section 4 we consider a series of independent, but overlapping economies of length of three periods. Thus the features that are observed in the subsequent quarter would be independent and equal, for example, to $[0, 0, 1, 0, ...]$ and $[0, 0, 0, 1, ...]$.

Then, given an observation of either positive or negative news in the first period,

$$M_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \hat{s}_1 & 1-\hat{s}_1 \\ 0 & 0 \end{bmatrix}$$

and again in the second period

$$M_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \hat{s}_1 & 1-\hat{s}_1 \\ \hat{s}_2 & 1-\hat{s}_2 \end{bmatrix}$$

Note that what is stored in memory is the observed announcement itself, with no adjustment for the fact that it might be correlated with a previous announcement. Thus the neglect of correlation does not require a specific assumption in the memory model; rather, it falls out of the encoding process.

The assignment of each month's announcements to a distinct context matters for what repetition does to memory. A repeated item is not, by itself, easier to remember. When repetitions occur close together, in effectively the same context, the later occurrences add little; memory strengthens when repetitions are spaced, so that each occurrence is encoded in a different context (Kahana et al., 2024). The repetition in our setting is spaced: the repetitive month's announcements come from different firms, in a different month, and hence in a new context, and each is encoded as a fresh episode. The experimental findings that Fedyk and Hodson (2023) report draw the same line. Finance professionals do not react to a news article that is literally reprinted, a repetition within the same context that is easily recognized as old; they do react when old information is recombined across sources, so that the repetition arrives in a new form. In our model, the aggregate component of the repetitive month's earnings is a repetition of exactly this kind.

Agents at time t form prices based on their memory. They assess the state of the economy through free recall, prompted by a neutral context c_{neutral} that places equal weight on each encoded context. The neutral context parallels the one Wachter and Kahana (2024) use to assess portfolio choice, and the one Cohen and Kahana (2022) use to assess the effect of mood on recall. In this simplified setting, the theory's feature retrieval rule reduces to

$$f_t^{in} = M_t^\top c_{\text{neutral}},$$

whose elements are the probabilities that the agent assigns to the good and bad states.

Retrieval occurs at the end of the month, after that month's announcements are encoded. For example, in period 2, $c_{\text{neutral}} = [1/4, 1/4, 1/4, 1/4]$ and

$$f_2^{in} = M_2^\top c_{\text{neutral}} = \frac{1}{4}\begin{bmatrix} 1 + \hat{s}_1 + \hat{s}_2 \\ 3 - \hat{s}_1 - \hat{s}_2 \end{bmatrix}.$$

The first element of f_2^{in} equals P_2, the probability that the agent with correlation neglect assigns to the good state; the same holds in every period.

The memory model also clarifies the relation to diagnostic expectations. Under diagnostic expectations (Bordalo et al., 2019), beliefs overweight outcomes made representative by recent news; the distortion operates through the updating rule. Here, the retrieval rule is Bayesian given what memory contains, and beliefs err because repeated episodes accumulate in the memory matrix as if they were independent. Kahana et al. (2025) provide a foundation for the representativeness heuristic based on contextual encoding and retrieval. The present model relies only on repeated items being reinforced in memory, a broad property of memory models. Representativeness, in contrast, relies on contextually determined similarity, which is specific to retrieved-context theory.

C Alternative model of correlation neglect

In this section, we consider an alternative model that is based more closely on Enke and Zimmermann (2019). As in the model in Section 4, assume there are three periods, corresponding to the newsy month, the repetitive month, and the first month of the next quarter, when the value of the market portfolio $\tilde{P}$ is drawn from the normal distribution with mean μ. As in Enke and Zimmermann (2019), assume there are latent signals $s_j, j = 1, 2$, jointly normally distributed, independent, and with unknown mean μ. Investors directly observe s_1 in period 1 and $\hat{s}_2 = 1/2(s_1 + s_2)$ in period 2. While Enke and Zimmermann (2019) do not specify a prior distribution, their results are consistent with assuming a normal prior distribution for μ, a known variance for s_j, and then taking the limit as the prior becomes uninformative. We will also make these assumptions, along with specifying μ_0 as the prior mean of μ.

Equilibrium with rational investors

In period 1, $P_1 = s_1$, and $R_1 = P_1 - \mu_0$. In period 2, $P_2 = \hat{s}_2$ and $R_2 = \frac{1}{2}(s_2 - s_1)$. The investor observes s_1 and evaluates the expected period-2 return:

$$E[R_2|s_1] = \frac{1}{2}(E[s_2|s_1] - s_1) = 0,$$

because $E[s_2|s_1] = E[\mu|s_1] = s_1$ under the uninformative prior. Therefore the investor cannot profit from knowledge of R_1, and there is no continuation. The investor evaluates the expected period-3 return:

$$E[R_3|s_1, s_2] = E[\tilde{P}|s_1, s_2] - P_2 = \frac{s_1 + s_2}{2} - \frac{s_1 + s_2}{2} = 0.$$

There is no reversal in the case of a rational investor.

Equilibrium with correlation neglect

In period 1, again $P_1 = s_1$ and $R_1 = s_1 - \mu_0$. However, in period 2, investors form an inefficient forecast that over-weights s_1, $P_2 = \frac{1}{2}(s_1 + \hat{s}_2) = \frac{3}{4}s_1 + \frac{1}{4}s_2$, because they mistakenly believe that $\hat{s}_2$ is an independent draw. The period-2 return equals $R_2 = P_2 - P_1 = \frac{1}{4}(s_2 - s_1)$. Because R_2 differs from its value under the rational investor only by a multiplicative constant,

$$E[R_2|s_1] = \frac{1}{4}(E[s_2|s_1] - s_1) = 0.$$

The model does not generate continuation from the investor's perspective. For reversal, the investor evaluates:

$$E[R_3|s_1, s_2] = E[\tilde{P}|s_1, s_2] - P_2 = \frac{s_1 + s_2}{2} - \left(\frac{3}{4}s_1 + \frac{1}{4}s_2\right) = \frac{1}{4}(s_2 - s_1) = R_2.$$

Because $E[R_3|s_1, s_2] = R_2$, a positive period-2 return predicts a positive period-3 return, whereas the data calls for a negative relation. Thus the model generates the wrong sign for reversal.

This model differs from the one in the main text in two respects. Most obviously, investors here learn about the mean of a normal distribution, whereas in the main model, investors learn about a probability. When learning about the mean, receiving the same signal twice and incorporating it does not change the investor's estimate of the price (though it does shrink the variance). However, when one is learning about a probability, receiving a signal always causes the mean to update (see Equation 2), though eventually by a vanishingly small amount. In a model with normally distributed risks, over-reaction only occurs in the second moment, not in the first, as in the data. A second difference lies in the source of randomness in the second period. In this model, the second-period signal always contains new information (s_2), and the CN investor's price under-weights s_2 relative to the efficient price. The correction in period 3 therefore moves toward s_2—in the same direction as R_2. In the model in the main text, the source of randomness is whether the first signal is brought to investors' attention or not. When $x = 1$, there is no new information, and the CN investor's reaction is entirely spurious, leading to reversal ($E[R_3|x = 1, s_1] = -R_2$). Despite these differences, the models have an important source of similarity in that incorrect inference arises from neglecting correlations.

D The common time component in quarterly earnings

In this section, we use four approaches to detect the common time component in quarterly corporate earnings in the United States. We use income before extraordinary items (ibq) from Compustat to measure quarterly firm earnings, and compute return-on-equity (ROE), or earnings divided by book value of equity, as a stationary representation of earnings news. We group each quarterly cross-section of U.S. firms by the first digit of their SIC code, which broadly represents sectors. We then conduct a principal component analysis (PCA) on the quarterly panel of sector-level return on equity, computed as total quarterly earnings divided by book value of equity within the sector. Note that using a broad measure of sector rather than a granular measure of industries avoids missing values and enables PCA without any imputation.

The upper left panel of Figure D1 shows that the first PC on this panel of earnings is effectively their time component—it has a correlation of 0.93 with aggregate quarterly ROE. The upper right figure shows that the first PC explains about 30% of the total variation on the panel, which is economically significant. The lower left panel is a histogram of the R^2s from regressions of firm-level ROE on aggregate ROE among firms that are ever in the S&P 500 universe. They have a mean of 8.4%. These three panels show that there is clearly a strong common time component in quarterly earnings in the United States.

As a result of this common time component, earnings in a newsy month predict those in the subsequent repetitive month especially strongly. The lower right panel uses three bars to represent predictive coefficients of i) quarter q's newsy month earnings on repetitive month earnings, ii) q's newsy month earnings on $q+1$'s newsy month earnings, and iii) q's repetitive month earnings on $q+1$'s repetitive month earnings. The common time component makes i) significantly larger than ii) and iii), despite the significant overlap in the firms announcing in two repetitive or newsy months from two adjacent quarters.

Figure D1: The common component of quarterly earnings

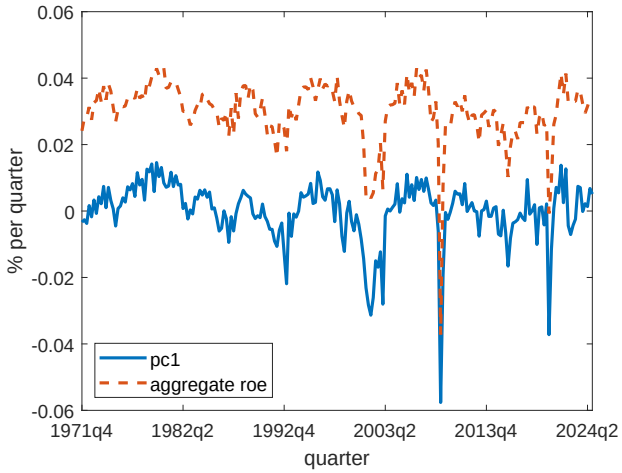

(a) PC1 and aggregate ROE

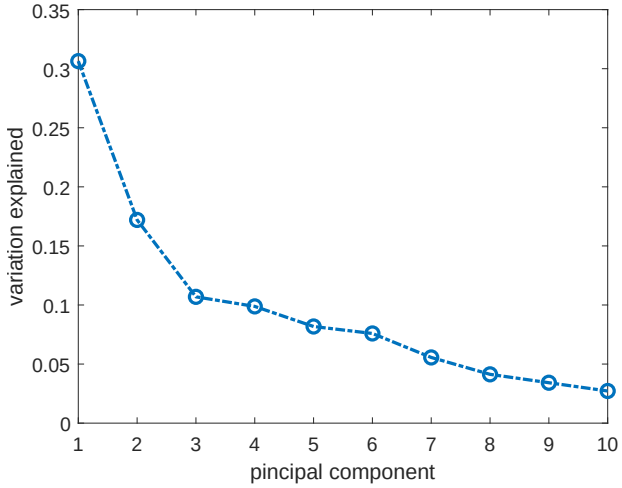

(b) Variation explained by PCs

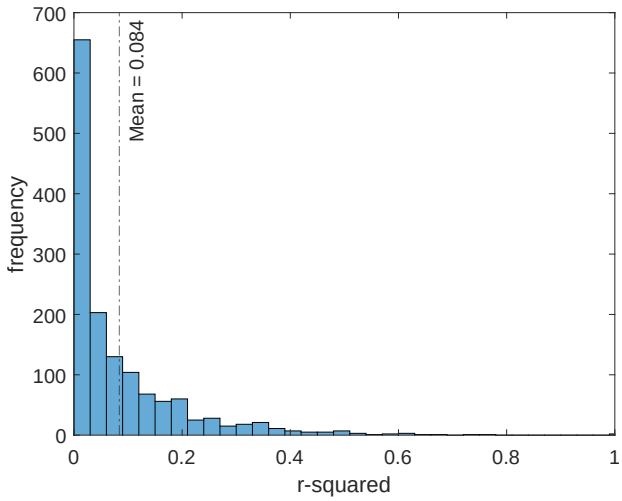

(c) Firm level ROE on aggregate ROE

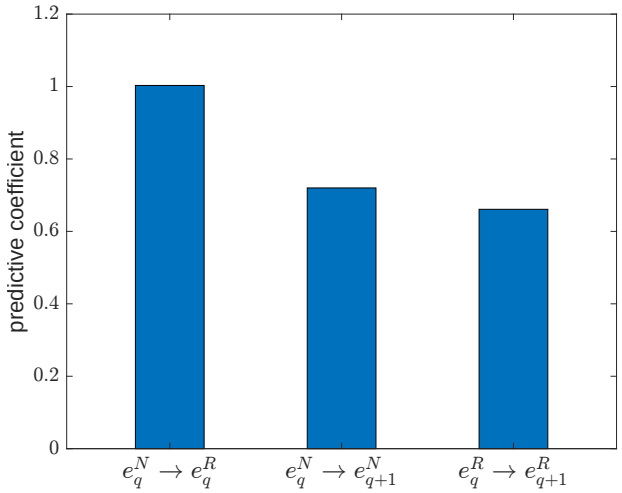

(d) Predicting monthly ROE

Figure (a) conducts a PCA on the quarterly-sector (represented using the first digit of a firm's SIC code) panel of ROE, and plots the first PC against the aggregate quarterly ROE. Figure (b) plots the fraction of total variation explained by each PC on the panel of ROE. Figure (c) conducts regressions of firm level ROE on aggregate ROE among firms that historically have been part of the S&P 500 universe, and plots a histogram of the R^2s. Figure (d) shows predictive coefficients of i) quarter q's newsy month earnings on repetitive month earnings, ii) q's newsy month earnings on $q+1$'s newsy month earnings, and iii) q's repetitive month earnings on $q+1$'s repetitive month earnings, from left to right.

E Evidence from global markets

The U.S. stock market serves as the testing ground for financial theories due to its size, liquidity, and high-quality data. As our theory revolves around the earnings seasons, it is applicable to the U.S. market for two additional reasons. First, the earnings announcements in the United States occur much faster than in other major countries. Second, the U.S. market exerts an asymmetric influence on other markets (e.g., Brusa et al. 2020). These two factors make the newsy months in the United States uniquely important, as they take information content away from the newsy months in other countries.

Despite the reduced information content, newsy months in countries other than the U.S. do contain valuable information about these countries' aggregate economies. Earnings seasons are primarily a U.S. feature, but not exclusively so. We therefore find it valuable to carefully investigate the international markets, as they offer another out-of-sample testing ground for our theory. We study the same set of global stock markets as do Moskowitz et al. (2012)—Austria, Australia, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, the United Kingdom, Greece, Hong Kong, Ireland, Israel, Italy, Japan, Netherlands, Norway, New Zealand, Portugal, Sweden, and Singapore. These are developed markets classified by Morgan Stanley Capital International (MSCI), with high liquidity and accessibility to foreign investors.

We first characterize the global earnings cycle in Table E1. Our data come from Worldscope. Panel A shows that, as in the United States, the fiscal periods of the global markets are aligned with the calendar quarters, with 87.8% of the fiscal periods closing at the end of a calendar quarter. Many countries mandate semi-annual, rather than quarterly, financial reporting, a difference that is not as significant as it seems, given that, in countries requiring semi-annual reporting, comparable numbers of firms have Mar/Sep and Jun/Dec as the end of fiscal periods. Hence, the country as a whole still operates on a quarterly calendar, and the first months of each quarter remain the earliest time at which fresh news arrives. Our theory's prediction therefore directly applies to those countries as well.[27]

Panel B of Table E1 shows that the earnings announcements in the global markets are slower than those in the United States. However, about 20% of the announcements occur in the first month of a quarter, and they are more than sufficient for investors to draw lessons about the aggregate cash flow. In other words, the first months of quarters are also newsy in

[27]The only exception is Australia, where Jun/Dec fiscal periods are overwhelmingly more common and the Mar/Sep firms only represent a trivial portion (less than 4%) of the economy. In Australia, the return pattern indeed appears to operate at a semiannual frequency, and adjusting the predictor construction accordingly improves its forecasting performance. As this source of variation applies to only one country, we hesitate to read too much into it, and keep the predictor construction recipe the same for simplicity.

the global markets. The second months contain about 60% of the earnings announcements, similar to those in the United States. We therefore continue to set the first months of quarters as newsy and the second months as repetitive for the global markets.[28]

As the previous paragraph discusses, we can interpret the same months as newsy, despite the variation in the regulations between the US and other countries. However, our explanation for the return pattern relies not only on the current reporting patterns but also on these reporting patterns remaining stable over the history of the returns that we study. In both the United States and the global sample, the return data start substantially earlier than the accounting data from which the newsy and repetitive months are inferred. However, the history of accounting data is much longer in the United States, and the aggregate earnings cycle remained stable over the course of this longer history. Due to NYSE requirements (Kraft et al. 2017), quarterly reporting predates the CRSP data, which begin in 1926. Such systematic evidence is lacking outside the United States.

Given this context, we focus on the financially developed countries, where financial regulation is sounder and has a longer history. Even with this restriction, results prior to 1974 should be interpreted with caution. With this caveat in mind, we test the predicted return pattern on the global markets in Table E2. The pattern in Table E2 is similar to that in Table 2, providing out-of-sample support for our theory. The pattern of reversal appears weaker in column 1. This is likely a result of unconditional return continuation being much stronger outside the United States (Moskowitz et al. (2012)). The key test is again the differences in columns 3 and 4, which appear statistically significant across the panels.

[28]There is also considerable heterogeneity in reporting lags across countries, although it is rare for the first months of quarters to contain so few earnings announcements that it is impossible for investors to reliably learn anything. Australia is again the most probable exception to it: January and July contain about 1.5% of the on-cycle earnings announcements in Australia, and setting February and August as newsy instead appears to improve our predictor's forecasting performance. We, however, choose to keep our predictor construction logic uniform across countries for simplicity.

Table E1: **The global earnings cycle: When do firms end their fiscal periods, and when do they announce earnings?**

	A: FP end month		B: Earnings announcement month			
			All announcements		Aligned & timely	
	Count	Percent	Count	Percent	Count	Percent
Group 1: Jan/Apr/Jul/Oct	56,951	5.7%	228,720	23.1%	168,568	19.8%
Group 2: Feb/May/Aug/Nov	64,303	6.5%	559,349	56.4%	534,837	62.7%
Group 3: Mar/Jun/Sep/Dec	870,263	87.8%	203,448	20.5%	149,654	17.5%
Total	991,517	100.0%	991,517	100.0%	853,059	100.0%

Panel A counts the company-fiscal periods in Worldscope by the months in which they ended; 87.8% means that 87.8% of the fiscal periods ended on the last day of March, June, September, or December. Panel B counts the company-fiscal periods by the months in which the fiscal periods' earnings were announced. The "All announcements" columns include all data present in panel A. The "Aligned & timely" columns apply two filters: i) the fiscal quarter ended in the Group 3 months and ii) the earnings announcement occurred within 183 days of the fiscal quarter end date. Data are quarterly and semi-annually from 1998 to 2021.

Table E2: Continuation and reversal in global market returns

	(1) Reversal	(2) Continuation	(3) Difference	(4) Diff. w/ control
		A: Full sample		
$mkt_{c,nr(t)}$	-0.046	0.113***	-0.159***	-0.156***
	[-1.15]	[3.62]	[-3.15]	[-2.75]
N	6,549	6,547	13,096	13,001
		B: First half		
$mkt_{c,nr(t)}$	-0.005	0.153***	-0.157**	-0.147**
	[-0.10]	[3.71]	[-2.51]	[-2.15]
N	2,432	2,431	4,863	4,813
		C: Second half		
$mkt_{c,nr(t)}$	-0.068	0.094**	-0.162**	-0.163**
	[-1.23]	[2.15]	[-2.30]	[-2.05]
N	4,117	4,116	8,233	8,188
		D: Post-WWII		
$mkt_{c,nr(t)}$	-0.053	0.100***	-0.154***	-0.149**
	[-1.16]	[2.83]	[-2.64]	[-2.29]
N	5,601	5,599	11,200	11,118

This table presents estimated βs from the monthly time-series regression $mkt_{c,t} = \alpha + \beta mkt_{c,nr(t)} + \epsilon_{c,t}$. Here, $mkt_{c,t}$ is country c's stock market return in excess of the risk-free rate in month t, and $nr(t)$ is the most recent month before t that is newsy or repetitive. Column 1 is on the sample in which the dependent variable month t is a first month of a quarter (newsy); column 2 is on the sample in which the dependent variable month t is a second month of a quarter (repetitive). Column 3 shows their difference, extracted as $-\beta_2$ from the regression $mkt_{c,t} = \alpha + \beta_1 mkt_{c,nr(t)} + \beta_2 mkt_{c,nr(t)} \times I_t^r + \gamma I_t^r + \epsilon_{c,t}$ on the combined sample in column 1 and 2, where I_t^r is a dummy variable taking the value of 1 if month t is repetitive, and 0 otherwise. Column 4 additionally controls for the sum of returns of the past four newsy months strictly before t and its interaction with I_t^r. In Panels A, B, C, and D, data are monthly from 1926 to 2021, from 1926 to 1973, from 1974 to 2021, and from 1947 to 2021, respectively. T-statistics computed with standard errors clustered by month are reported in square brackets.