Publication Effect: Accelerating Anomalies[◊]

Zoran Ivković[♦] Morad Zekhnini[♠]

March 17, 2026

Abstract

Academic publication compresses the timeline of anomaly returns: pre-publication, anomaly signals produced gradual drifts; post-publication, price adjustments occur within days of disclosure, following step-function-like patterns. The effect is heterogeneous: short-term return responses strengthen on average but remain unchanged or even reverse for some anomalies. This heterogeneity matches cash flow predictability—the direction of cash-flow predictability predicts the direction of the publication effect. These patterns are marginally related to standard behavioral frictions and are inconsistent with data mining, Bayesian shrinkage, and conventional risk-factor models. Instead, publication appears to lower the cost of identifying which signals carry fundamental information, enabling faster price discovery.

Keywords: Market efficiency, cross-sectional predictability, anomalies, behavioral finance, rational markets

JEL Codes: G12, G14

[◊] We thank Vikas Aggrawal, Scott Cederburg, Andrew Chen, Alexander Chinco, Charlie Clarke, Lauren Cohen, Tony Cookson, Zhi Da, Joey Engelberg, Chuck Fang, Kenneth French, Caroline Genc, John Hand, Cam Harvey, Charles Hadlock, Ryan Israelsen, Hao Jiang, Bryan Kelly, Naveen Khanna, David McLean, Dmitriy Muravyev, Jeff Pontiff, Jim Poterba, David Schreindorfer, Stanislav Sokolinski, Avanidhar Subrahmanyam, Richard Thakor, Charles Trzcinka, Tengfei Zhang; seminar participants at American University, Drexel University, George Mason University, Georgetown University, McMaster University, Michigan State University, and University of Delaware; and the 12th MSUFCU Conference participants for insightful feedback and comments. All errors are our own.

[♦] Michigan State University, email: ivkovich@bus.msu.edu

[♠] Michigan State University, email: zekhnini@msu.edu (corresponding author)

Publication Effect: Accelerating Anomalies

March 17, 2026

Abstract

Academic publication compresses the timeline of anomaly returns: pre-publication, anomaly signals produced gradual drifts; post-publication, price adjustments occur within days of disclosure, following step-function-like patterns. The effect is heterogeneous: short-term return responses strengthen on average but remain unchanged or even reverse for some anomalies. This heterogeneity matches cash flow predictability—the direction of cash-flow predictability predicts the direction of the publication effect. These patterns are marginally related to standard behavioral frictions and are inconsistent with data mining, Bayesian shrinkage, and conventional risk-factor models. Instead, publication appears to lower the cost of identifying which signals carry fundamental information, enabling faster price discovery.

Keywords: Market efficiency, cross-sectional predictability, anomalies, behavioral finance, rational markets

JEL Codes: G12, G14

1. Introduction

Academic research in financial economics and accounting has documented hundreds of in-sample cross-sectional return predictors, often referred to as "anomalies." The industrious production of these predictors raises a fundamental debate about how and why anomalies arise. Anomalies have been dismissed as statistical artifacts (Schwert, 2003; Harvey, Liu, and Zhu, 2016), linked to Bayesian updating (Jensen, Kelly, and Pedersen, 2023), explained as compensation for systematic risk (Cochrane, 2011; Hou, Xue, and Zhang, 2015), and attributed both to limits to arbitrage (Shleifer and Vishny, 1997) and behavioral mispricing (Hong and Stein, 1999; Hirshleifer, 2001; Barberis and Thaler, 2003). More recently, McLean and Pontiff (2016), Linnainmaa and Roberts (2018), Hou, Xue, and Zhang (2020), and Jensen, Kelly, and Pedersen (2023), among others, reignite the debate about interpreting anomalies and their sources by showing that predictability generally weakens out-of-sample.[1]

McLean and Pontiff (2016) propose and test a learning hypothesis: investors learn about mispricing from academic publications and raising awareness about anomalies weakens them. Accordingly, McLean and Pontiff (2016) highlight a substantial attenuating effect of publications on out-of-sample (post-publication) anomaly strategy profits. They also point out that this decay is more pronounced for anomalies that are less costly to arbitrage, suggesting that the role of publication is intertwined with that of limits to arbitrage.

The attenuation evidence notwithstanding, current discourse regarding role of publication in the context of anomalies remains unsettled. As the learning hypothesis suggests, anomalies may be a source of genuine information that markets, once they become aware of it, learn to process

[1] International evidence appears not to conform to the U.S.-based out-of-sample tests. For example, Lu, Stambaugh, and Yuan (2017) and Jacobs and Müller (2020) show that many anomalies persist internationally.

more efficiently. However, they may also be spurious outcomes of extensive data mining. Understanding how academic research revelations concerning anomalies affect anomaly returns is central for both understanding what anomalies truly are and evaluating how markets adapt to new information.

Using data from Compustat Snapshot and SEC EDGAR filings, we align short-term anomaly returns with disclosure dates of firm accounting variables. This alignment allows us to focus on short-horizon returns and assess how markets process anomaly information as it arrives. In the publication context, we document both a compression in the timeline of reactions to new information and a significant heterogeneity in the post-publication evolution of anomaly returns.

Pre-publication, stock prices responded only gradually, with modest initial movements followed by extended drifts. Capturing those extended drifts constituted the backbone of empirical work that had established the existence of anomalies in pre-publication periods. We find that the post-publication reaction to anomaly signals has transformed into a near-step function. Most of the price adjustment occurs within a few days of disclosure, compressing the temporal pattern into a nearly discrete jump. By contrast, the subsequent drift is much more subdued than it had been pre-publication, as McLean and Pontiff (2016) had documented. This sharp acceleration in information incorporation provides direct evidence of the publication effect: academic research provides revelations concerning anomalies, markets learn from those revelations and, in the process, become more efficient.

A closer inspection of anomaly-by-anomaly returns reveals substantial heterogeneity within publication effects—from decidedly negative, to very modest, to strongly positive. The learning hypothesis, as stated by McLean and Pontiff (2016), explains both the average effects of declining longer-horizon return drift and the return response timeline compression. However, in its present form, it does not address adequately the heterogeneity, that is, the simultaneous presence of both short-term return continuations for some anomalies and short-term return reversals for others. This heterogeneity not only raises important questions about the sources of anomalies but also provides a valuable setting in which to examine competing theories of anomaly origins and mechanisms.

To better understand the mechanism of learning, we appeal to the standard present-value framework, in which prices are functions of expected cash flows and discount rates (e.g., Gordon, 1959; Campbell and Shiller, 1988). Cash flows enter the present-value relation directly and map monotonically into prices. Therefore, we first analyze the predictive relation between subsequent cash flows (standardized unexpected earnings in our empirical tests) and anomaly signals. We find that, on average, the relation is positive and very strong both pre-publication and post-publication; in fact, publication has not altered its intensity. However, we again find substantial anomaly-level heterogeneity in the link between subsequent cash flows and anomaly signals, ranging from decidedly negative, to very modest, to strongly positive.

The correlation between anomaly-level publication effects and anomaly-level cash-flow predictability is 0.8. This strong connection is plausibly based on a substantial combinatorial reduction in the number of promising trading strategies. A significantly reduced viable investment strategies search space allows investors to focus on the signals and the stocks more likely to experience improved cash flows. This interpretation fits both behavioral and rational frameworks. On the behavioral side, not incorporating some potentially useful information can be attributed to either cognitive limitations that prevent investors from processing vast volumes of possible signals, or the belief that such an endeavor is futile and unlikely to yield any value-relevant insights. On the rational side, the cash-flow predictability patterns are consistent with the Grossman and Stiglitz

(1980) paradigm that information acquisition is costly. Before publication, extracting value-relevant information from accounting data required specialized knowledge and skills. It was costly to determine which accounting signals, if any, predict cash flows (and thus returns), so many investors may have rationally remained "ignorant." Academic publication lowers this cost by highlighting a considerably reduced set of valuable signals, prompting a more immediate incorporation of their cash-flow content into prices.

This interpretation is not a generic restatement of the "cash flows matter" adage. The novelty is that academic publication appears to reduce the cost of identifying which accounting signals carry cash-flow information—and the direction of that information—so that investors can react immediately at disclosure. This channel may be viewed as either a reduction in cognitive or search costs (investors now know where to look) or as an updating of higher-level beliefs about the relevance of cash-flow news embedded in accounting signals (as opposed to the more familiar updating about return distributions). Ultimately, our evidence does not distinguish whether this form of inattention is rational or behavioral; we use 'misreaction' as shorthand that encompasses both.

Next, we focus on discount rates. Unlike cash flows, discount rates often pose estimation problems because they depend on the exact model under study. In most empirical implementations, discount rates correspond to factor exposures (Cochrane, 2011; Hou, Xue, and Zhang, 2015). However, priced factor exposures are hard to reconcile with the sudden shifts in return patterns that we detect. This is especially true for the signals that lead to return reversals. To the extent that these signals proxied for some positively compensated dimension of risk before publication, the reversals would imply these risks suddenly changed course drastically and now command a negative premium instead. In direct tests of exposure to the ubiquitous risk factors of Fama and

French (2015), augmented with the Carhart (1997) momentum factor, we find little evidence that the signals we analyze are associated with changes in risk exposure both before and after publication.[2]

Overall, our results underscore the notion that striving toward market efficiency is an ongoing, evolving pursuit. In the context of the anomalies that we analyze, markets appear to ignore or misreact to otherwise value-relevant information during the pre-publication period. Academic publication appears to "correct" these initial misreactions. These patterns are inconsistent with data-mining concerns. By construction, the publication effect is measured out-of-sample because the original research did not consider disclosure-date returns. If the pre-publication anomaly returns were spurious, the post-publication, out-of-sample short-run signal responses would be negligible. They are not.

The average post-publication strengthening of signal responses is also inconsistent with Bayesian shrinkage (e.g., Jensen, Kelly, and Pedersen, 2023). In a Bayesian framework, an empiricist's initial, prior belief is that a signal does not predict returns. Upon observing a sizeable return premium associated with the signal in-sample (pre-publication), the empiricist's updated, posterior belief regarding the expected future (post-publication) return premium associated with the signal attenuates ("shrinks") relative to the in-sample (pre-publication) value. Consequently, the Bayesian framework incorrectly predicts a post-publication attenuation of the signal response rather than what we observe—its amplification.

Our main findings of a misreaction to cash-flow information are broadly consistent with behavioral accounts of anomalies. Within the behavioral framework, however, we find that specific proxies for limits to arbitrage, inattention, and sentiment are only modestly related to the return-

[2] This result is perhaps not surprising given that most anomaly signals were benchmarked against variations of these factor models and still commanded significant alphas.

pattern change. We also find little evidence that idiosyncratic volatility, skewness, and capital gains (all elements of the Barberis, Jin, and Wang's (2021) prospect-theory-based model to explain several anomalies) contributed to the publication effect.

In sum, our results suggest that anomalies are not merely statistical quirks, artifacts of Bayesian shrinkage, or factor exposure. Instead, anomaly signals contain fundamental information about firm performance to which investors initially underreact. Publication accelerates the recognition of this information, helping markets close the gap between disclosure and price adjustment. In this sense, our evidence supports the concept of market-level learning prompted by the revelations contributed by academic research.

Our findings are consistent with prior research that documents the concentration of anomaly returns around subsequent earnings releases (Abarbanell and Bushee, 1997) and firm information days (Engelberg, McLean, and Pontiff, 2018). The association between short-term anomaly returns and cash flows also supports the Daniel and Titman (1997) paradigm that characteristics (i.e., signals), rather than factor exposures, predict returns.

Our paper contributes most directly to the anomalies literature, encompassing the issues of anomaly origins, characterization, and attenuation (e.g., Chordia, Subrahmanyam, and Tong, 2014; Cederburg and O'Doherty, 2015; Akbas, Armstrong, Sorescu, and Subrahmanyam, 2016; Linnainmaa and Roberts, 2018; Chu, Hirshleifer, and Ma, 2020; Hou, Xue, and Zhang, 2020; Muravyev, Pearson, and Pollet, 2022; Chen and Velikov, 2023; and Jensen, Kelly, and Pedersen, 2023). It is also related to the literature on information processing and the speed of price discovery (e.g., DellaVigna and Pollet, 2009; Hirshleifer, Lim, and Teoh, 2009; and Dugast and Foucault, 2018). Finally, it aligns with the emerging literature that highlights the dynamic, evolving nature of financial markets (e.g., Lo, 2017; Lo and Zhang, 2024).

2. Data and empirical design

2.1. Data

We focus on the 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K. Their key feature is that Compustat Snapshot and/or EDGAR filings can identify the precise date as of which variables required to construct accounting anomaly portfolios become available. Because EDGAR filings are available from 1994, our sample spans from 1994 to 2022 and only includes anomalies with an original publication date after 1996.[3] We exclude stocks priced below $5 to mitigate the concern that microcap stocks might influence the results. Finally, some analyses also incorporate signals based on IBES analyst forecasts.

2.2. Empirical design

To avoid look-ahead bias (Fama and French, 1993), prior research routinely introduces a substantial waiting period between the approximate time of information arrival and anomaly portfolio formation. For example, if a firm's fiscal year ends on December 31st and its annual report is released on March 9th, previous studies would form portfolios using the March information at the end of June or the beginning of July. The information will remain incorporated into portfolio formation through June of the following calendar year. By contrast, we focus on price changes around the release date itself (March 9th in our example). This focus deliberately restricts our analysis to only the subset of anomalies that rely on financial reports.

We use two approaches to estimating difference-in-differences models by comparing the long and short legs of the anomalies both before and after academic publication. The first approach

[3] Considering anomalies published in 1996 or later, a necessity because of the EDGAR data availability limitations, precludes the inclusion of some of the well-known anomalies published earlier. For example, among many anomalies published earlier are Total assets to market, Book to market, and Book leverage (Fama and French, 1992).

effectively is a meta-analysis across multiple anomalies and publication dates. We treat each signal-stock as a separate observation in our regression, and we cluster standard errors by calendar quarter. This conservative level of clustering ensures the robustness of our inference, given that a stock and its returns can appear multiple times in the sample, corresponding to various signals.[4] The second approach is a collection of individual anomaly difference-in-differences estimations, yielding as many regression results as there are individual anomalies in our sample.

Under either approach, to estimate a consistent model for all signals, we reorganize each signal so that high values of the signal correspond to the long leg of the anomaly. In addition, we standardize all signals and convert them into the indicator variable *HighSignal*. For binary signals, we set *HighSignal* to 1 for long-leg stocks and to 0 for short-leg stocks. For continuous signals, we use a quintile transformation. Each quarter, we assign *HighSignal* the value of 1 to signal-stock observations in the highest quintile of the signal, and the value of 0 to those in the lowest quintile (the middle three quintiles are omitted). Following McLean and Pontiff (2016), for each anomaly the indicator variable *PostPublication* is set to 1 for all the time periods during or after the year of the first academic paper publication concerning the anomaly, and to 0 otherwise.[5]

The dependent variable y_{iat} in most analyses is the cumulative stock i return over a three-day window (t - 1, t + 1) surrounding the information release date t, identified by anomaly a as of the end of the quarter encompassing date t as either high-signal stock ($HighSignal_{iat}$ = 1) or a low-

---

[4] To confirm the robustness of our inferences, we estimate the baseline results from Table 2 using a broad spectrum of alternative clustering strategies and report the results in Table IA2.2.1 of the Internet Appendix. That table shows that the results are very robust to any even remotely feasible clustering approach.

[5] McLean and Pontiff (2016) use the journal publication year as the defining moment in their analyses. Table IA2.2.2 in the Internet Appendix shows that our results hold if the *PostPublication* indicator variable were defined as starting two years before journal publication (presumably when the first draft was circulated and presented in a few seminars), one year before journal publication (presumably when the already mature working paper appeared in more seminars and a few high-visibility conferences), one year after journal publication (allowing for one year for the diffusion of the information contained in the published paper), and two years after journal publication (allowing for two years for the diffusion of the information contained in the published paper).

signal stock ($HighSignal_{iat} = 0$). In some analyses, the dependent variables are Fama-French factor exposures for stock *i*, estimated using daily returns over the three-calendar month window starting after the information release date *t*. Finally, in some analyses, the dependent variable is future standardized unexpected earnings for stock *i*, based on the first earnings announcement made after the information release date *t*. The independent variables in all analyses are indicator variables *HighSignal*, *PostPublication*, and their interaction. In sum, our empirical model is:

$$y_{iat} = \alpha + \beta \cdot HighSignal_{iat} + \gamma \cdot PostPublication_{at} + \delta \cdot HighSignal_{iat} \times PostPublication_{at} + \varepsilon_{iat} \quad . \qquad (1)$$

2.3. Non-financial anomalies

The key features of the anomalies that we study are that both the signal and its timing are clearly observable to, and unambiguously understood by, both the econometrician and market participants. In the context of market-based signals (i.e., return- and volume-based anomalies), both the innovations themselves and their timing are difficult to identify. For example, consider a pronounced daily positive stock return, a jump as defined in Kapadia and Zekhnini (2019). Such an event constitutes an innovation for short-term reversal strategies and for intermediate-term momentum strategies. Investors implementing either strategy will pursue different positions in the stock, shorting for short-term reversal and buying for momentum. However, implementations of the momentum strategy typically introduce a one-month waiting period in a deliberate attempt to avoid short-term reversal contamination. Assigning an exact date for information arrival in this context requires a judicious, and likely somewhat arbitrary, assessment of when various investors incorporated the jump into their respective portfolios. Further complications associated with strategies such as momentum include overlapping tranches of the portfolios in the (*J*, *K*) sense, as originally articulated by Jegadeesh and Titman (1993).

A plausible extrapolation from our analysis is that non-accounting anomalies may well engender similar estimates of the publication effect as accounting-based anomalies. However, detecting this effect outside of the domain of accounting anomalies is hampered by a lack of clarity regarding precisely what constitutes a signal and when information reaches the market. Investigating market-based anomalies may prove a fruitful avenue for future research, critically hinging on precise signal timing and discerning interpretation of the information embedded in various non-accounting signals.

2.4. Descriptive statistics

Panel A of Table 1 summarizes the number of days firms take to disclose financial information. On average, it takes more than 39 days after a fiscal quarter ends and more than 63 days after a fiscal year ends for financial information to appear in Compustat Snapshot. The corresponding SEC filings (10-Q or a 10-K) on average take about six days longer, 45.6 days and 69.5 days, respectively. Our timing convention, adopted throughout the paper, is that the earliest date on which the information appears in either source is the information-relevant date. The information date occurs on average 36.4 days after fiscal quarters, and 55.9 days after fiscal years.

Panels B and C of Table 1 report the average monthly equal-weighted and value-weighted returns to anomaly strategies, respectively. For binary anomaly signals, all stocks in our sample are placed into long or short legs each month. For continuous anomaly signals, we sort all stocks into five quintiles and retain only the stocks whose returns belong to extreme quintiles. The long leg consists of stocks in the top quintile and the short leg consists of stocks in the bottom quintile. For both equal-weighted and value-weighted portfolios, the returns are significantly larger before publication than they are after publication. The decline of average monthly anomaly returns is from 53.2 to 18.8 bps for equal-weighted returns and from 38.2 to 13.8 bps for value-weighted returns. These attenuations are consistent with the findings reported in McLean and Pontiff (2016).

3. Results

3.1. Cross-sectional predictability acceleration

We focus on accounting anomalies—the anomalies that use financial report-based variables (i.e., Compustat variables). The original studies documenting these anomalies rely almost exclusively on annual Compustat variables. However, given the nature of accounting-based variables, conducting our tests and interpreting their results in the context of the learning hypothesis necessitates that, whenever possible, we transform these anomalies to a quarterly frequency. For all annual Compustat-based signals that can also be derived from the quarterly data, we construct a quarterly analogue and combine this list of anomalies with the quarterly anomalies.

Using the resulting list of 34 accounting anomalies,[6] we estimate the regression in Equation (1). The first two columns of Table 2 summarize the estimates. The first column uses raw returns, and the second column uses Fama-French 3-factor model-adjusted returns. All returns are adjusted for delisting issues (Shumway, 1997). Across both specifications, the coefficient associated with *HighSignal* is positive (18.7 bps and 19.0 bps, respectively), indicating that, before academic publication, signal releases generated a positive spread between long and short anomaly legs. The interaction term *HighSignal* x *PostPublication* also has a positive and significant coefficient in both specifications (48.4 bps and 47.1 bps, respectively). After academic publication, the spread between the long and short anomaly legs immediately after signal release widens further.

[6] The comprehensive list of the 34 accounting anomalies and the corresponding details are provided in the Internet Appendix Table IA1.1.

Fig. 1 shows the difference in the averages of the cumulative returns across stocks in *HighSignal* portfolios and in *LowSignal* portfolios in event time. For each anomaly and each quarter, we construct a cumulative return for each stock in event time. If stock i announces a new value of the signal for anomaly a (high or low) at day t, we build cumulative returns $\boldsymbol{r}_{i,\mathrm{a},\tau}$, $\tau = t$-2, t-1, t, t+1, …, t+63, setting $\boldsymbol{r}_{i,a,t\text{-}2} = 0$. We proceed to calculate subsequent event-time cumulative returns as $\boldsymbol{r}_{i,a,\tau} = (1+\boldsymbol{r}_{i,a,\tau\text{-}1})(1+r_{i,a,\tau}) - 1$, $\tau = t$-1, t, t+1, …, t+63, where $r_{i,a,\tau}$ is the delisting-adjusted return for stock i on day τ. To compare signal reactions before and after publication, we split the sample into observations before academic publication and after academic publication. We then combine all pre-publication and post-publication samples across all anomalies to obtain one combined pre-publication subsample and one post-publication subsample. For each combined subsample, we average cumulative returns for *HighSignal* and *LowSignal* stocks in each quarter, taking the averages of all cumulative returns $r_{i,t,\tau}$ for $\tau = t$-2, t, t+1, …, t+63 for stocks with high signals ($r_{H,\tau}$) and low signals ($r_{L,\tau}$), respectively. The thick lines in the plot show the values $r_{H,\tau}$ - $r_{L,\tau}$ for $\tau = t$-2, t, t+1, …, t+63 both before publication (black) and after publication (blue). The dashed lines are the confidence intervals, calculated for each day in event time using standard errors clustered by calendar quarter.

Fig. 1 demonstrates that, before publication, returns have a very modest short-term reaction to new information, continuing gradually to drift upward over the next few months. After publication, consistent with the learning hypothesis, cumulative returns resemble a step-function: an immediate short-term reaction with hardly any subsequent drift. Investors quickly respond to the signals once informed by academic literature. Arguably, Fig. 1 suggests that the return pattern is no longer anomalous, at least not from the perspective of ubiquitous pre-publication drift patterns.

3.2. Signal dates

Our empirical design relies on the exact dates on which signals reach the market. To determine these dates, we combine information from Compustat Snapshot and the SEC EDGAR filings (10-K/10-Q reports). As with all empirical work, the data we use may introduce some small degree of error in identifying the exact information arrival dates. For instance, a firm might reveal some of the financial information through a press release before the Compustat entry date or the EDGAR filing.[7] In such instances, our empirical setting would be a conservative estimate of the publication effect. Therefore, to the extent that the dates we identify introduce noise to the process, our estimates provide a lower bound on the true value of the publication effect.

To verify this intuition, we conduct a falsification exercise. For each firm-quarter observation, we randomly pick a date between 6 and 30 trading days before or after the actual signal date. We use these randomly generated dates as "placebo" signal dates and repeat the analysis 1,000 times. Fig. 2 shows distributions of pre-publication high-low signal 3-day event returns (top panel), post-publication high-low signal 3-day event returns (middle panel), and the publication effect (bottom panel) from the 1,000 simulation runs. The dependent variables in all simulations are the FF3-adjusted cumulative stock returns over three-day windows (t - 1, t + 1) surrounding the simulated information release dates t. In all panels, green arrows denote the corresponding point estimate based on the actual information release dates, as reported in the second column of Table 2. The tightness of the three simulated estimated coefficient histograms and the remote locations of the green arrows affirm that the reactions we estimate in our baseline

[7] The 2000 passage of Regulation Fair Disclosure, prohibiting selective disclosure to certain individuals or groups, minimizes concerns that our results might be affected by selective disclosure.

specification (Table 2) stand out massively when compared to the coefficients estimated around randomly chosen, non-event dates.[8]

3.3. Individual anomaly-level publication effect heterogeneity

The average publication effect masks substantial heterogeneity across anomalies. We estimate Equation (1) anomaly-by-anomaly and plot the publication effect estimates (the coefficient estimate associated with the interaction term *HighSignal* x *PostPublication*) in Fig. 3. Organizing the estimates from lowest to highest, Fig. 3 suggests that the publication effect is significantly negative for some anomalies, modest for others, and significantly positive and large for many anomalies. This heterogeneity raises important questions about the sources of anomalies. At the same time, it provides a valuable setting in which to examine competing theories of anomaly origins and advance new ones.

4. Cash flow and discount rate news

A growing literature documents that stock returns on news days differ significantly from stock returns on quiet days (e.g., Lucca and Moench (2015) in the context of the Federal Open Market Committee meetings and Jiang and Yao (2013) in the context of stock-return jumps). The release of financial information undoubtedly reveals material news to the market as evidenced by an increase in volatility during the days we study. In the context of financial disclosure, prior studies show that earnings announcement days (also information rich days) increase the firm's

---

[8] Table IA2.2.3 in the Internet Appendix documents the mean values of the 1,000 simulated estimates and their standard deviations for both the raw returns (first column) and the FF3-adjusted returns (second column). Consistent with the histograms from Fig. 2, all the means of the simulated coefficients, particularly those pertaining to the high-low signal pre-publication reactions, the high-low signal post-publication reactions, and the publication effects, are within a few basis points of zero, and all the corresponding sampling distributions are very narrow, as evinced by their low standard deviations (shown in parentheses beneath the means in the table).

market beta (Patton and Verardo, 2012) and reveal market-wide cash flow news (Savor and Wilson, 2016).[9] Therefore, the publication effect can reflect changes in either information about cash flows or risk exposure. In the spirit of the Gordon (1959) and Campbell and Shiller (1988) models, we analyze the association between signals and future cash flows using surprises in earnings and risk premia using factor exposures.

4.1. Information about future cash flows

The signals underlying most anomalies are convenient ways to standardize information about the performance and prospects of firms in the economy. The publication effect represents a change in how this information affects valuations. On the one hand, in a behavioral sense, the publication effect might reflect previously erroneous processing of the new information. The signals, despite containing value-relevant information, did not affect prices immediately, and the act of publication corrected the mispricing by facilitating a quicker incorporation of the signals. On the other hand, in a rational framework, the effects that we document in the previous section might reflect a change in the informational content of the signals. For example, a structural change (e.g., new regulation or accounting standards) renders these signals more salient for valuation. To distinguish between these two hypotheses, we analyze the association between the signals and future cash flows.

Unlike risk exposure, which poses significant measurement issues, cash flows are easy to measure and analyze. To gauge the ability of anomaly signals to predict cash flows, we repeat the analysis in Equation (1) but replace the dependent variable with standardized unexpected earnings

[9] Kapadia and Zekhnini (2019) show that most of the equity premium is realized on a handful of information rich days—days with jumps in stock returns. However, the overlap between stock-level return jump days and earnings announcement days is 13%, modest at best. Therefore, the events we document likely capture some of, but not all of, the valuation-relevant information about stocks.

(SUE), surprises in earnings in the subsequent quarter. SUE is calculated as the difference between the announced earnings excluding special items and expected earnings using a seasonal random walk model.[10]

Table 4 presents the results. On average, anomaly signals positively predict earnings surprises in the pre-publication sample. The surprise in earnings for firms with positive signals is 0.152 standard deviations higher than that of low signal firms. This predictability remained relatively unchanged after publication. The point estimate of the earnings predictability publication effect is very modest, only 0.024 standard deviations, and is not statistically significant. This modest publication effect does not align with a structural change in the relation between signals and fundamentals. Post publication, the signals continue to predict future cash flows.

4.2. Earnings predictability and publication

4.2.1. Heterogeneity in earnings predictability

A crucial insight from earnings prediction is that the average effect, just like the average publication effect, obscures significant heterogeneity across anomaly signals. Fig. 4 presents the earnings predictability estimates, estimated anomaly-by-anomaly and sorted from the lowest to the highest. The figure shows that the subset of anomalies for which signals predict earnings negatively, the subset of anomalies for which the coefficient is small in magnitude (be it mildly negative or mildly positive), and the subset of anomalies for which signals predict earnings positively are about equally represented in our overall sample of anomalies. The range of the individual anomaly-level coefficient estimates is considerable, from -2.09 (Accruals BM) to 2.54

[10] Internet Appendix Table IA2.4 features robustness tests that replicate Table 4 for two alternative measures of standardized earnings, SUE' and SUE''. SUE' is the difference between the announced earnings and expected earnings using a seasonal random walk model (unlike SUE, SUE' does not exclude special items). SUE'' is the difference between the announced earnings and analyst forecast consensus. The results in Internet Appendix Table IA2.4 are very similar to the SUE-based results reported in Table 4.

(PS). Because SUE is standardized, the difference in signal-earnings prediction is more than four standard deviations.

4.2.2. Earnings predictability and the publication effect

The heterogeneity in earnings prediction allows us to sharpen our inferences about the learning hypothesis. Signals differ in their ability to predict cash flows, a key dimension of learning. In other words, cash-flow predictability provides a measure of misreaction for each anomaly. For example, an anomaly signal may predict positive shocks to future cash flows. Because the signals curated in the literature were associated with high returns in the long run, markets likely underreacted in the short run to the information that the signal conveys. Yet the signal predicts positive earnings shocks, so subsequent earnings shocks were more likely to be positive, and the price drifted upward, resulting in high returns. In sum, a positive relation between future cash flows and anomaly signals indicates return underreaction and, analogously, a negative relation between future cash flows and anomaly signals indicates return overreaction. Therefore, academic research, intentionally or not, informs markets about this easily verifiable classification of signals, leading, in turn, to differing publication effects.

For each anomaly, we restrict the sample to the pre-publication period and estimate the relation between the future earnings surprises and the anomaly signal. We use these estimates of predictability of future earnings surprises to organize the anomalies and summarize them in Fig. 5. Conceptually, Fig. 5 restates the anomaly-by-anomaly publication coefficients from Fig. 3 (green bars) by displaying them according to the ascending order of the anomaly-by-anomaly earnings predictability coefficients from Fig. 4 (red bars). Visually, the green bars (anomaly-level publication effect estimates) and the red bars (anomaly-level earnings predictability estimates) follow each other very closely. Indeed, the correlation between anomaly-level publication effects and anomaly-level cash-flow predictability is 0.8. The anomalies represented by the leftmost bars in Fig. 5 can be interpreted as overreactions (because those signals predict reversals) and, correspondingly, the anomalies represented by the bars to the right side of the graph can be interpreted as underreactions (because those signals predict continuations).

The strong visual pattern in Fig. 5 is further confirmed by testing the relation more rigorously and parsimoniously using Equation (1). We organize the anomalies in our sample into five groups based on each anomaly's ability to predict earnings before it appeared in the academic literature. Each group has about seven anomalies, with the first group consisting mostly of signals that predict earnings surprise negatively. We estimate Equation (1) separately for each of the five anomaly groups and summarize the results in Table 5, mimicking the organization of Fig. 5.

Panel A of Table 5, corresponding to the green bars from Figure 5, documents Equation (1) estimation of short-term return responses, ranging from Q1, the quintile encompassing the anomalies characterized by the most negative SUE predictability (indicating reversal, thus overreaction in the pre-publication period, replaced by swift negative step-function like reaction in the post-publication period), to Q5, the quintile encompassing the anomalies characterized by the most positive SUE predictability (indicating continuation, thus underreaction in the pre-publication period, replaced by swift positive step-function like reaction in the post-publication period). The publication effect for the first group of anomalies is -84.2 bps. It increases monotonically to 182.1 bps for the last group of anomalies. This gap across the five groups of anomalies is very wide, amounting to 2.66 percent. A comparison with the average publication effect, estimated across all 34 anomalies (presented in Table 2 and restated in the first column of Panel A, Table 5), suggests that there is a great deal of heterogeneity among the anomalies, ranging from strong overreactions, to mild reactions, to strong underreactions.

Panel B of Table 5, corresponding to the red bars from Figure 5, documents Equation (1) estimation of SUE predictability, ranging from Q1, the quintile encompassing the anomalies characterized by the most negative SUE predictability, to Q5, the quintile encompassing the anomalies characterized by the most positive SUE predictability. As shown in the first column of Panel B (restated from Table 4), the publication effect did not alter the SUE predictability. This largely remains true even when broken down into anomaly quintiles, with but one exception (the sensitivity in Q1 declined by about one-half). The alignment between the publication effect and anomaly short-term returns across quintiles (Panel A of Table 5) and the coefficients in either of the top two rows of Panel B, Table 5 (pertaining to pre-publication SUE sensitivity in the first row and post-publication SUE sensitivity in the second row, respectively) is very close, reaffirming the 0.8 correlation between the two previously calculated from anomaly-by-anomaly estimates.[11]

Finally, we couple the cash-flow predictability anomaly quintiles and trace cumulative returns to long-short strategies by the five groups of anomalies. Fig. 6 further accentuates that cumulative returns are aligned with the direction of cash-flow predictability. It shows what essentially is a disaggregated Fig. 1. As in Fig. 1, before publication (top panel of the figure) each of the five anomaly quintiles, be it the quintile with the lowest estimates of cash-flow predictability (Lowest Q), the next three quintiles, or the quintile with the highest estimates of cash-flow predictability (Highest Q), shows a very modest short-term reaction to new information and a gradual drift thereafter. Once again as in Fig. 1, consistent with the learning hypothesis, after publication (bottom panel of the figure) cumulative returns for each of the five anomaly quintiles,

[11] Internet Appendix Tables IA2.5.1 and IA2.5.2 feature robustness tests that replicate Table 5 for two alternative measures of standardized earnings, SUE' and SUE'', respectively. SUE' is the difference between the announced earnings and expected earnings using a seasonal random walk model (unlike SUE, SUE' does not exclude special items). SUE'' is the difference between the announced earnings and analyst forecast consensus. The results reported in Internet Appendix Tables IA2.5.1 and IA2.5.2 are very similar to the SUE-based results reported in Table 5.

be it the quintile with the lowest estimates of cash-flow predictability (Lowest Q), the next three quintiles, or the quintile with the highest estimates of cash-flow predictability (Highest Q), resemble a step-function, with an immediate short-term reaction and little subsequent drift.

4.2.3. Learning about cash flows

The dispersion shown in Figures 3, 4, and 5 underscores that publication does not merely attenuate anomaly profits on average. Instead, it reshapes the immediate response at disclosure in directions that line up with the signals' cash-flow content. Therefore, the key object of learning is not a return premium per se, but the mapping from disclosed accounting information to expected cash flows (and the sign of that mapping). In this sense, the mechanism is best viewed as signal identification and interpretation rather than a broad "cash-flows-matter" revelation. Before publication, investors face a high-dimensional set of candidate accounting transformations whose links to fundamentals are uncertain; after publication, academic curation narrows the set and clarifies which signals predict both cash flows and returns positively, weakly, or negatively.

There are several non-mutually-exclusive ways to rationalize this learning. First, investors may be cognitively constrained and unable to process thousands of candidate signals; publication effectively focuses attention on a smaller subset. Second, investors may not have held beliefs that cash-flow information embedded in these signals mattered for valuation; publication provides credible evidence that updates this higher-level belief and increases the salience of cash-flow news at disclosure. Third, in a Grossman–Stiglitz (1980) sense, learning reflects a reduction in information acquisition costs: it becomes cheaper to identify and interpret the cash-flow implications of specific accounting signals. All these interpretations focus on information, attention, and beliefs—not preference shifts. It is difficult to reconcile an abrupt post-publication change in reaction speeds and directions with changes in investor tastes, known to be generally stable over short horizons.

4.3. Factor exposures

The previous analyses cast doubt on structural changes in the link between signals and cash flows. This section explores the possibility that publication might have coincided with a structural break in risk factor exposure. Ex ante, it is plausible to assert that, because the fundamental elements of risk can evolve over time, the gradual decline in long-term anomaly returns (McLean and Pontiff, 2016; Hou, Xue, and Zhang, 2018; Jensen, Kelly, and Pedersen, 2023) may be consistent with changes in risk. However, the timeline compression in signal responses to what essentially is a step function appears too sudden to fit any such explanation. Moreover, in a formal test of this hypothesis, we find at best marginal changes in signal predictions for standard risk exposures.

Specifically, using the Fama-French 5-factor model (Fama and French, 2015) augmented with the Carhart (1997) momentum factor, we estimate future factor exposures (i.e., factor betas) at the stock level. In Table 3, we present the results of our analyses of the relation between these exposures over the subsequent three-day windows and the *HighSignal* indicator variable. Before publication, compared to low signals, high signals were associated with statistically significantly lower values of market and SMB exposure, but higher values of HML and RMW exposures. For CMA and UMD, the coefficients are insignificant. For all factor exposures, however, the magnitudes of the associated coefficients are too small to justify the high premia documented in the literature. The absolute values of the exposure changes associated with the high minus low signal pre-publication gap are no larger than 0.053. Over our sample period, the market had the highest average daily premium at 3.4 bps. Multiplying the highest exposure change with this highest average daily premium over a three-day horizon yields only one-half of a single basis point (0.053 x 3.4 bps x 3 = 0.5 bps), a value quite negligible relative to the magnitude of the publication effect.

Post-publication, the associations between future factor exposures and *HighSignal* shrink toward zero for most factors, whereas they increase for CMA and UMD. However, as was the case before publication, the post-publication difference in factor exposures between high and low signal stocks is too small to warrant a difference in premia. The exposure changes are also too small to justify the publication effect we document in Table 2.

5. Data mining and Bayesian updating

Anomalies derive their name from the challenges they often pose to risk-based theories of cross-sectional returns. Under such theories, cross-sectional variation in expected returns should be associated (only) with corresponding variation in risk exposures. Extant anomaly studies use realized returns averaged over large portfolios to proxy for expected returns and document strong return prediction. The findings in these studies can still be consistent with risk-based explanations if realized returns are poor proxies for expected returns or if the signals used to predict returns capture a source of systematic risk. Realized returns, observed ex-post, are subject to data mining concerns that challenge their role as proxies for expected returns, and risk exposures are notoriously difficult to estimate precisely. These considerations lead to econometric challenges in establishing a relation between signals and risk.

5.1. Data mining

Through the lens of long-run returns, anomalies often display strong in-sample performance that weakens or vanishes out-of-sample, raising concerns about data mining (Harvey, 2017). This return pattern is consistent with spurious discovery or deliberate data mining: if in-sample return patterns are indeed spurious, weaker out-of-sample returns are to be expected. In the Appendix, we develop two parsimonious models. The first is a rational model, capturing the basic idea behind data mining. We juxtapose this model with the second, behavioral learning model in which markets are initially unaware of the value of a certain signal but learn about its value over time. The two models provide similar predictions regarding long-term returns but have different predictions about immediate signal responses.

To distinguish between the two theories, we focus on immediate post-signal returns. Perhaps surprisingly, immediate post-signal returns received little attention in the literature. That these short-term returns were not part of the initial studies makes them, by construction, free of data-mining biases. Our baseline findings show that the previously identified signals are indeed associated with returns, both in the pre- and post-publication periods. Consistent with the learning explanation, the publication primarily affects the temporal dynamics and considerably accelerates the return reaction to signals. Put differently, these analyses demonstrate that, on average, anomalies are not mere artifacts of data mining.

The pattern in Fig. 6 further contradicts the data mining hypothesis. Under this hypothesis the signals carried no information content and all associations with returns were spurious. Any out-of-sample analysis should reveal no discernible pattern. Yet the return patterns are very distinct for anomalies with positive and negative associations with cash flows. Focusing on only the lowest and highest quintiles, the pre-publication cumulative returns looked like gradual drifts downwards and upwards, respectively. After publication, the signal release reactions resemble step functions but in opposite directions. It is most unlikely that a spurious relation between signals and values gives rise to these patterns.

5.2. Bayesian updating

Jensen, Kelly, and Pedersen (2023) suggest that the anomaly return attenuation is consistent with Bayesian updating. From a statistical perspective, the researcher's prior belief is that a portfolio built on a given signal should generate an alpha of zero. If a researcher observes a significant alpha in-sample, the updated (posterior) alpha, reflecting all available information, should shrink the sample estimate toward zero (the prior). Without any data-mining efforts, a Bayesian updater should expect the out-of-sample alpha estimates to be lower than the in-sample alpha estimates.

An additional dimension of data mining pertains to the various choices researchers employ to establish return predictability. For example, researchers can select among various empirical proxies for profitability to maximize return predictability in-sample. Jensen, Kelly, and Pedersen (2023) find that, in support of risk-based explanations, most anomalies replicate not only in-sample but also out-of-sample, though with diminished alphas. However, they concede that the attenuation is stronger than what the Bayesian-updating model would predict on its own. This strong attenuation allows for the possibility of an arbitrage/data-mining component to the alpha estimates documented in the literature.

The Bayesian shrinkage argument assumes that the data-generating process remained unchanged across the two samples used to estimate in-sample and out-of-sample anomaly strategy performance. For returns, this assumption is akin to having the source of the original anomaly returns be risk-based. The initial estimates of the risk premiums were likely inflated, and an appropriately shrunk Bayesian estimator should match the out-of-sample estimates more closely. The estimates that we report use a 3-day window of returns and have large magnitudes. It is implausible that the risk inherent in these trading strategies is concentrated around the announcement of financial data. Even if it were, the Bayesian argument ought to apply to our sample as well. The return response to the documented signals should attenuate, just like the premium estimates. Instead, we find that this response increases. Therefore, our baseline results contradict the Bayesian updating explanation.

Moreover, the shrinkage argument cannot explain the heterogeneity in responses. For a few signals, the association with future cash flows was negative. Correspondingly, the immediate reaction to signal releases became very negative post publication. This response contrasts the positive long-run returns estimated by the econometrician. Therefore, the reversal in returns is inconsistent with large in-sample magnitudes.

6. Limits to arbitrage and other behavioral considerations

The strong association between the publication effect and cash flow predictability suggests that many anomalies can be characterized as misreactions to new information. In this section, we discuss several potential explanations for these misreactions, using our setting as a novel laboratory for conducting empirical tests concerning the origin of anomalies. Many extant anomaly explanations focus on a single anomaly or a small subset of anomalies. As such, they are inherently narrow and are unlikely to yield a broad platform through which to enhance our understanding of anomalies at large. In a deliberate effort to circumvent narrow explanations and interpretations, we focus only on the theories that span a large subset of anomalies.

The cash flow predictability analysis from the previous section fits with many of the proposed explanations for anomalies. We established that various anomalies predicted cash flows both before and after academic publications. Absent frictions, be they cognitive or market-structure induced, the reactions should have resembled a step-like function all along, not just after publication.[12] This lack of a pre-publication response indicates that, at the time, markets either ignored relevant information or, at least, could not reflect it in a timely fashion, ostensibly because of some frictions. The extant anomaly explanations catalogue structural and cognitive frictions that could lead to these patterns. Therefore, these frictions can complement our finding of cash flow predictability and enhance our understanding of the fundamentals of anomalies and their evolution.

6.1. Limits to arbitrage

Prices can deviate from fundamental values if frictions prevent "rational" investors from correcting any mispricing. These frictions are often labeled "limits to arbitrage." They encompass implementation costs and capital constraints (Pontiff, 1996, 2006), noise-trader risk that makes arbitrage positions risky (De Long, Shleifer, Summers, and Waldmann, 1990), synchronization risk (Abreu and Brunnermeier, 2002), and liquidity spiral risk (Brunnermeier and Pedersen, 2009), all of which discourage arbitrageurs from correcting the mispricing. Shleifer and Vishny (1997) establish a prominent role for limits to arbitrage in the anomaly literature, arguing that anomalies are likely to appear when limits to arbitrage are high. Stambaugh, Yu, and Yuan (2012) build on the notion that short-selling constraints can limit arbitrage (Miller, 1977) and argue that these constraints, along with variations in investor sentiment, exacerbate mispricing. McLean and Pontiff (2016) show that the extent of the post-publication predictability erosion is particularly pronounced among the more readily tradable anomalies—the anomalies that are easier and less costly to arbitrage.

---

[12] The pre- and post-publication patterns are also congruent with a general attenuation in frictions. In other words, if a new technology or regulation lifted certain frictions, the improvement in signal responses could be attributed to such a wide-sweeping change. However, the publication of anomaly research was scattered throughout our sample period, making the possibility of one such change an unlikely driver of the return patterns we observe.

In the context of anomalies, arbitrage frictions can cause the slow incorporation of signals into prices. The compression of the signal response timelines is consistent with this possibility, but the heterogeneous effect across anomalies is not. Indeed, if limits to arbitrage were the dominant mechanism that prevented timely incorporation of new information into prices, return pattern changes should be similar across all anomalies. Instead, the signal response changes are strongly negative for some anomalies, weak for some of the others, and strongly positive for most. This dispersion challenges the plausibility of limits to arbitrage as the leading mechanism behind accounting anomalies. To leverage our setting further, we use various proxies for the arbitrage limits discussed in the literature and examine the extent to which they relate to short-term anomaly returns. We use market capitalization, illiquidity, and bid-ask spread as proxies for transaction costs, and the percentage of institutional ownership and percentage of short interest as proxies for short-sale constraints. These variables correspond closely to the implementation-cost channel emphasized by Pontiff (1996, 2006) and to short-sale constraints in the spirit of Miller (1977) and Stambaugh, Yu, and Yuan (2012).

Table 6 presents the results. For each stock-level mechanical trading friction proxy (market capitalization, Amihud illiquidity, bid-ask spread, institutional ownership, and short interest), the sample is split into five quintiles, with the quintile cutoffs determined from the sample of all stocks. The ordering of quintiles for each proxy of mechanical trading frictions is such that the leftmost column features the stocks with the lowest friction levels (the easiest stocks to trade), and the rightmost column features stocks with the highest friction levels (the hardest stocks to trade).

6.1.1. Transaction costs

We use market capitalization, Amihud illiquidity (Amihud, 2002), and bid-ask spread (Corwin and Schultz, 2012) as proxies for friction rooted in transaction costs. The results of estimating Equation (1) for quintile subsamples, presented in the first three rows of Table 6, do not show clear evidence of transaction costs inhibiting signal responses. These three rows reveal two key findings. First, across all three transaction cost measures, variation among the publication effect estimate by friction quintiles is very small compared to the very pronounced variation among the quintile estimates based on cash flows from Table 5. Second, whatever small variation there is across friction quintiles does not indicate a monotonic pattern consistent with frictions driving signal reactions. Instead, there is a mild inverted U relation between the publication effect and each of the three transaction cost measures.

6.1.2. Shorting frictions

The (in)ability to short stocks constitutes another friction that might affect the publication effect. Stambaugh, Yu, and Yuan (2012) show that a significant portion of the anomaly spreads can be attributed to the short leg, a pattern consistent with difficulties in shorting overpriced stocks. In other words, short-sale constraints dampen the ability of the market to incorporate new information. To the extent that publication alleviates these constraints, the return reactions should strengthen for constrained stock and remain relatively unchanged for unconstrained stocks.

We use two proxies for short-selling constraints: institutional ownership and short interest. Institutional ownership captures the shares owned by 13-F institutions (reported at the end of the previous calendar quarter), and short interest is the average number of shares reported as short interest (calculated for each stock in each month as the simple average of mid-month and end-of-month observations). Both are expressed in relative terms, that is, scaled by the number of shares outstanding.

The fourth row of Table 6 presents the publication effect estimates for the institutional ownership quintile subsamples. The leftmost column presents the estimates for the least constrained observations—the subsample of stocks most heavily held by institutional owners, typically made available for a variety of trading strategies (e.g., lending; shorting; executing block trades and cross-trades). The rightmost column presents the estimates for the subsample of stocks with the lightest institutional holding, hence, the most constrained set of stocks. The last row of Table 6 presents the publication effect estimates for the short interest quintile subsamples. The leftmost column presents the estimates for the least constrained observations—the subsample of stocks with the highest short interest. The rightmost column presents the estimates for the subsample of stocks with the lowest short interest, hence, the most constrained set of stocks.

The last two rows of Table 6 reveal two key findings. First, there is a declining pattern across the quintile subsamples for each measure. The decline is somewhat lumpy for institutional ownership and is monotonic for short interest. The gaps between the extreme quintiles are 16.4 bps (institutional ownership) to 17.9 bps (short interest). These declining patterns across quintiles are consistent with shorting frictions playing some, albeit modest role in the space of anomalies. However, as was the case for transaction costs, for both shorting friction measures the variation among the publication effect estimate by friction quintiles is very small compared to the very pronounced variation among the quintile estimates based on cash flows from Table 5. Therefore, the relation between the publication effect and shorting frictions is economically very small and likely plays at best a modest role in this context.

6.2. Dimensions of behavioral finance

Barberis (2018) views behavioral finance as a vehicle that brings psychological realism to traditional models. This realism can be viewed along three dimensions: cognitive limits, non-standard beliefs, and non-standard preferences. Traditional models assume perfectly rational agents with perfect Bayesian updating of beliefs and time-consistent preferences. However, mounting psychological evidence is in stark contrast to the traditional assumptions. In this section, we identify candidate empirical measures to analyze each of these dimensions.[13]

6.2.1. Cognitive limits

A prime example of human cognitive limits is (in)attention. Limited attention (Kahneman, 1973; DellaVigna and Pollet, 2009) and gradual information diffusion (e.g., Hong and Stein, 1999; Hirshleifer, 2001) cause sluggish reactions to firm-specific news, leading to medium-horizon return predictability. The post-earnings announcement drift (e.g., Ball and Brown, 1968; Hirshleifer, Lim, and Teoh, 2011), as well as return predictability across economic links (Cohen and Frazzini, 2008), size groups (Lo and Mackinlay, 1988), and industries (Hou, 2007) are all consistent with slow diffusion of information. Our analyses focus on a narrow window around information disclosure, so that inattention can lead to a muted response.

We split the sample into subsamples based on two proxies for limited attention: analyst coverage and news coverage. The first row of Table 7 presents the estimates of Equation (1) in each subsample. Its first two columns focus on analyst coverage, and its second two columns focus on news coverage.[14] The modest economic magnitude of the publication effect gap, documented as about 8 to 9 bps in both pairs of columns, suggests that variation in inattention does not substantially account for the heterogeneity in the publication effect; these gaps are truly miniscule compared to the cash-flow quintile publication effect variation from Table 5.

[13] We recognize that the empirical proxies for the various building blocks of behavioral finance that we employ do not exhaust the rich set of potential proxies for any of the three dimensions of behavioral finance. The proxies that we do use, however, cover many of the major themes and are grounded in extensive literature developed to date.

[14] For both analysts and news, most stocks have zero coverage at any given point. To avoid asymmetric treatment of these stocks, we refrain from quintile grouping and instead use only two groups (no coverage and coverage).

6.2.2. Nonstandard beliefs

We use sentiment proxies as an example of nonstandard beliefs. Many investors and researchers view aggregate investor sentiment as a key driver of mispricing waves across markets (Brown and Cliff, 2005; Baker and Wurgler, 2006, 2007; Huang, Jiang, Tu, and Zhou, 2015). Stambaugh, Yu, and Yuan (2012) argue that, given the difficulty in selling stocks short, periods of high sentiment are more pertinent and lead to more mispricing. Using the Baker and Wurgler (2006) sentiment index, we split the sample into low and high sentiment months to assess the role of sentiment in the context of the publication effect. We then estimate Equation (1) in each subsample and report the corresponding publication effect estimates in the second row of Table 7. We report results for classifications based on the previous month's value of the index (the first two columns) and the contemporaneous value of the index (the second two columns).

Like inattention, sentiment features a very modest economic magnitude of the publication effect gap, about 5.2 to 5.5 bps in both pairs of columns, suggesting that sentiment variation plays at most a modest role in explaining the publication effect. As is the case for inattention, the sentiment gaps of around 5 bps are miniscule compared to the cash-flow quintile publication effect variation from Table 5.

6.2.3. Nonstandard preferences

The most prominent theory regarding nonstandard preferences is prospect theory (Kahneman and Tversky, 1979). It contends that investor preferences exhibit large deviations from the standard expected utility framework. Investors focus on gains and losses relative to a reference point, and these preferences can lead to anomalous valuation patterns relative to an expected utility framework. However, neither prospect theory nor any other theory based on a precise representation of the utility of agents can explain the post-publication change in anomaly returns.

Preferences underlying any such theory are innate and all but immutable; it is implausible that investor preferences could change so drastically so quickly in response to an information-revelation event such as a publication of an anomaly. Nonetheless, literature to date has relied on prospect theory to motivate the prevalence of certain empirical patterns (e.g., Grinblatt and Han, 2005; Barberis and Huang, 2008; Barberis, Jin, and Wang, 2021) by operationalizing prospect-theory channels through a handful of observable variables. Accordingly, we proceed by analyzing several predictability patterns documented in the literature with explanations rooted in prospect theory. We focus on the theories advanced in connection with anomalies and which provide empirical measures that can be tested in our setting.

Prospect-theory preferences have been used to explain several cross-sectional return patterns. Grinblatt and Han (2005) show how reference-dependent preferences and reluctance to realize losses can generate momentum in stocks returns. Barberis and Huang (2008) argue that probability weighting induces investors to overpay for positively skewed, lottery-like stocks and underweight negative skewed stocks. Building on this literature, Barberis, Jin, and Wang (2021) develop a prospect-theory-based model to explain several anomalies. In the model, agents are susceptible to stock idiosyncratic return volatility, idiosyncratic return skewness, and capital gain overhang. The Barberis, Jin, and Wang (2021) model offers a potential mechanism for the patterns we documented around information release. We use our setting to evaluate whether the three characteristics they propose impact information processing and hence substantially drive the publication effect.

Table 8 analyzes the publication effect as a function of stock-level return idiosyncratic volatility, return idiosyncratic skewness, and capital gain overhang. The publication effect and idiosyncratic volatility are related in ways quantitatively similar to those recorded for institutional ownership in Table 6. Once again, the publication effect variation across quintiles is small, and the pattern across the quintiles is largely flat, with the coefficient associated with Q1 somewhat lower, by about 15 to 22 bps; this publication effect gap between Q1 and the other four quintiles is miniscule relative to the publication effect quintile gaps recorded the cash-flow analyses from Table 5.

As shown in the second row of Table 8, the publication effect and idiosyncratic skewness are related in ways that resemble a mirror image of those recorded for idiosyncratic volatility. The publication effect variation across quintiles is small, and the pattern across quintiles is largely flat. This time, however, the coefficient associated with Q1 is higher than the rest by about 20 to 30 bps. Once again, though, this publication effect gap between Q1 and the other four quintiles is miniscule relative to the publication effect quintile gaps in the cash-flow analyses from Table 5.

The last row of Table 8 shows that the pattern of the publication effect magnitudes across the five quintiles, starting from the bottom quintile of the capital gains overhang (intuitively, consisting largely of losing stocks) to the top quintile of the capital gains overhang (intuitively, consisting largely of winning stocks) is monotonic, declining by about 26 bps across the five quintiles. Aside from not having a firm theoretical backing for its existence or magnitude, this variation is smaller by nearly, or about, an order of magnitude than the variation among cash-flow quintiles from Table 5. Overall, these findings suggest that the Barberis, Jin, and Wang (2021) prospect-theory variables do not substantially drive the publication effect.

7. Robustness tests

7.1. Summary of previously reported robustness tests

To ensure the validity of our results, we have conducted a series of robustness tests. We carried out our baseline estimates using numerous alternative standard error clustering approaches

(Internet Appendix Table IA2.2.1). We considered alternative dates for the publication, in the range from two years before journal publication to two years after publication (Internet Appendix Table IA2.2.2). We conducted a thorough falsification exercise in Section 3.2, randomly picking dates between 6 and 30 trading days before or after the actual signal dates and using these randomly generated dates as "placebo" signal dates to repeat our baseline analysis 1,000 times. As discussed in detail in Section 3.2 (Fig. 2, Table IA2.2.3 in the Internet Appendix), return reactions we estimate in our baseline specification (Table 2), grounded in the actual signal release dates, stand out massively when compared to return reactions estimated around randomly chosen, non-event dates. We also considered two alternative measures of standardized unexpected earnings and demonstrated that the results based on them are aligned very closely with those we report in Table 4, Table 5, Fig. 5, and Fig. 6 (the companions tables for the alternative SUE measures are provided in the Internet Appendix as Table IA2.4, IA2.5, Fig. IA2.5.1, Fig. IA2.5.2, Fig. IA2.6.1, and Fig. IA2.6.2).

7.2. Annual signals

In this section, we provide an additional robustness check, pertaining to the frequency of the anomaly signal (considering annual rather than quarterly signals). The learning hypothesis has different implications for annual and quarterly signals. Accounting-based variables can be classified as flow variables (e.g., revenue, profit, or investment) or stock variables (e.g., assets, debt, or equity). For flow variables, the annual values are merely the sum of the corresponding four quarterly values. Stock variables tend to evolve slowly, so their value at the end of the fiscal year is relatively close to what their value was at the end of the third fiscal quarter. Both flow and stock signals derived from annual accounting data are easily predicted using the information revealed after the third quarter. Therefore, annual signals, at least in their raw form, are unlikely to constitute new information for market participants. Consequently, the learning hypothesis implies that pure annual signals should not exhibit a publication effect. We test this hypothesis by re-estimating Equation (1) using annual signal data instead of quarterly signal data.

Table 9 reports model estimates for all the annual anomalies that rely on Compustat data. The two columns of the table correspond to raw returns and FF3-adjusted returns. Like quarterly signals, annual signals are associated with a gap between long and short leg stocks in the pre-publication sample. This gap is 8.8 bps and 4.9 bps for raw returns and FF3-adjusted returns, respectively. However, the gap remains essentially unchanged after publication. It shrinks minimally by 3.6 bps for raw returns and grows minimally by 2.3 bps for FF3-adjusted returns; both estimates are statistically insignificant. These results are consistent with Bowles, Reed, Ringgenberg, and Thornock (2024). In Figure A2 from their Internet Appendix G (Section G.3), Bowles, Reed, Ringgenberg, and Thornock (2024) report that the event time returns surrounding information releases before publication and after publication do not differ substantially. The framework presented in this paper provides a nuanced interpretation of their result—the publication effect is not inherently absent; rather, it is not detectable using annual signals.[15] Quarterly signals, as shown in this paper, tell a very different story of a very pronounced role of the publication event in the space of anomaly returns.

Fig. 7 uses annual signals to mimic the analyses reported for quarterly signals in Fig. 1. The cumulative returns for the post-publication period do not resemble a step-function. In the days immediately after the release of the signal information, the pre-publication and post-publication

[15] Further evidence may be gleaned from Internet Appendix Table IA2.9. It mimics the format of Table 9, featuring instead of the 34 accounting anomalies, those that rely (or can be transformed to rely) on quarterly signals from 10-Q and 10-K filings, the six accounting anomalies from McLean and Pontiff (2016) that can only be constructed from annual data (because the pertinent signal information only appears annually in 10-K filings). These six anomalies are listed and described in Internet Appendix Table IA1.2. As Internet Appendix Table IA2.9 demonstrates, the findings reported in connection with these six anomalies line up with expectations, that is, they are congruent with those presented in Table 9.

long-short returns are indistinguishable. After about a month, the two sets of returns start diverging as the pre-publication long-short portfolio continues to earn positive returns, but the post-publication portfolio does not.

8. Conclusion

This paper revisits how academic publication affects the dynamics of stock market anomalies. We focus on short-horizon reactions to anomaly signals by aligning returns with the precise disclosure dates of firm accounting variables. This alignment allows for sharper inferences about the effect of information releases on anomaly returns.

We report two central results. First, academic publications compress the timeline of anomaly returns. During the pre-publication periods, information releases led to a gradual drift. Post-publication, the return pattern morphs into a sharp, near-instantaneous price adjustment. This acceleration in information incorporation is consistent with a learning mechanism whereby markets become more efficient once academic research highlights anomaly patterns. Second, the publication effect is heterogeneous across anomalies. On average, the short-term signal responses strengthen, but some anomalies show little change, and others even reverse.

Based on fundamental valuation models, we decompose learning into cash flow and discount rate components. We find a strong and persistent link between anomaly signals and subsequent earnings surprises. We further show that the anomaly-level publication effect is tightly linked to this cash-flow predictability. On the other hand, traditional measures of risk exposure have little to no connection both with the anomaly signals and with the changes in return responses to these signals brought on by academic publication.

The publication effect magnitudes and its heterogeneity across anomalies are inconsistent with explanations rooted in data mining or Bayesian shrinkage. Instead, markets appear to have misreacted to the relevant valuation information contained in anomaly signals. These misreactions are in part corrected after the dissemination of academic research. This evolution of anomalies suggests that behavioral frictions might have played a role in their initial presence. However, the specific proxies we use for limits to arbitrage, inattention, sentiment, and prospect-theory variables account for relatively little of the cross-anomaly variation in the publication effect, especially compared to cash-flow predictability. Fig. 8 provides a visual summary by simultaneously reporting all the empirical findings reported in Sections 4 and 5. The black bars, representing cash-flow predictability, show a very strong, positive relation between the publication effect and cash-flow predictability: anomalies with relatively negative (positive) cash-flow predictability generate negative (positive) publication effects. Cash flow predictability stretches the reactions across the five quintiles to the range of 2-3% over a three-day horizon, a very pronounced magnitude over such a short period. By contrast, both limits to arbitrage and behavioral explanations, represented in the graph with a collection of grey bars across the five quintiles, are relatively flat.

Our results advance several debates in financial economics. They demonstrate that anomalies are not rooted in spurious statistical quirks or risk exposure shifts, and that the role of standard market frictions appears limited compared to the cash-flow information channel. Instead, anomaly signals carry information, and markets learn to incorporate that information into prices more efficiently. Therefore, our findings underscore the evolving nature of return predictability and contribute to a broader understanding of how markets assimilate information over time.

References

Abarbanell, Jeffrey S., and Brian J. Bushee, 1997, Fundamental Analysis, Future Earnings, and Stock Prices, *Journal of Accounting Research* 35, 1-24.

Abreu, D., and M.K. Brunnermeier, 2002, Synchronization risk and delayed arbitrage. *Journal of Financial Economics* 66 (2–3), 341–360.

Akbas, Ferhat, William J. Armstrong, Sorin M. Sorescu, and Avanidhar Subrahmanyam, 2016, Capital Market Efficiency and Arbitrage Efficacy, *Journal of Financial and Quantitative Analysis* 51, 387-413.

Amihud, Y. (2002), Illiquidity and stock returns: cross-section and time series effects, *Journal of Financial Markets* 5, 31-56.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor Sentiment and the Cross-Section of Stock Returns, *Journal of Finance* 61, 1645-1680.

Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor Sentiment in the Stock Market, *Journal of Economic Perspectives* 21, 129-152.

Ball, Ray, and Philip Brown, 1968, An Empirical Evaluation of Accounting Income Numbers, *Journal of Accounting Research* 6, 159-178.

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

Barberis, Nicholas, Lawrence. J. Jin, and Baolin Wang, 2021, Prospect theory and stock market anomalies. *Journal of Finance* 76 (5), 2639–2687.

Barberis, Nicholas, and Richard H. Thaler, 2003, A Survey of Behavioral Finance, in Constantinides, G. M., M., and Stulz, R. M. (Eds.): *Handbook of the Economics of Finance*, Vol. 1, Chapter 18, 1053-1128, Elsevier, Amsterdam.

Bowles, Boone, Adam V. Reed, Matthew C. Ringgenberg, and Jacob R. Thornock, 2024, Anomaly time, *Journal of Finance* 79 (5): 3543-3579.

Brown, Gregory W., and Michael T. Cliff, 2005, Investor Sentiment and Asset Valuation, *Journal of Business* 78, 405-440.

Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity. *Review of Financial Studies* 22 (6), 2201–2238.

Campbell, John Y., and Robert J. Shiller, 1988, Stock Prices, Earnings, and Expected Dividends, *Journal of Finance* 43, 661-676.

Carhart, Mark. M., 1997, On Persistence in Mutual Fund Performance. *Journal of Finance* 52(1), 57-82.

Cederburg, Scott, and Michael S. O'Doherty, 2015, Asset-Pricing Anomalies at the Firm Level, *Journal of Econometrics* 186, 113-128.

Chen, Andrew Y., and Mihail Velikov, 2023, Zeroing In on the Expected Returns of Anomalies, *Journal of Financial and Quantitative Analysis* 58, 968-1004.

Chen, A. Y., and Zimmermann, T. (2022). Open source cross-sectional asset pricing. *Critical Finance Review* 11(2), 207-264. https://doi.org/10.1561/104.00000112

Chordia, Tarun, Avanidhar Subrahmanyam, and Qing Tong, 2014, Have Capital Market Anomalies Attenuated in the Recent Era of High Liquidity and Trading Activity? *Journal of Accounting and Economics* 58, 41-58.

Chu, Yongqiang, David Hirshleifer, and Liang Ma, 2020, The Causal Effect of Limits to Arbitrage on Asset Pricing Anomalies, *Journal of Finance* 75, 2631-2672.

Cochrane, John H., 2011, Presidential Address: Discount Rates, *Journal of Finance* 66, 1047-1108.

Cohen, Lauren, and Andrea Frazzini, 2008, Economic links and predictable returns, *Journal of Finance* 63 (4): 1977-2011.

Corwin, Shane A., and Paul Schultz, 2012, A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices. *Journal of Finance* 67(2), 719-760.

Daniel, Kent, and Sheridan Titman, 1997, Evidence on the Characteristics of Cross-Sectional Variation in Stock Returns, *Journal of Finance* 52, 1-33.

DellaVigna, Stefano, and Joshua M. Pollet, 2009, Investor Inattention and Friday Earnings Announcements, *Journal of Finance* 64, 709-749.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98, 703-738.

Dugast, Jérôme, and Thierry Foucault, 2018, Data Abundance and Asset Price Informativeness, *Journal of Financial Economics* 130, 367-391.

Engelberg, Joseph, R. David McLean, and Jeffrey Pontiff, 2018, Anomalies and News, *Journal of Finance* 73, 1971-2001.

Fama, E. F., and French, K. R., 1992, The Cross-Section of Expected Stock Returns, *Journal of Finance* 47(2), 427-465.

Fama, Eugene F., and Kenneth R. French, 1993, Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics* 33, 3-56.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116 (1): 1-22.

Gordon, Myron J., 1959, Dividends, Earnings, and Stock Prices, *Review of Economics and Statistics* 41, 99-105.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70 (3), 393–408.

Grinblatt, Mark, and Bing Han, 2005, Prospect theory, mental accounting, and momentum, *Journal of Financial Economics* 78 (2), 311–339.

Harvey, Campbell R., 2017, Presidential Address: The Scientific Outlook in Financial Economics, *Journal of Finance* 72, 1399-1440.

Harvey, Campbell R., Yan Liu, and Heqing Zhu, 2016, … and the Cross-Section of Expected Returns, *Review of Financial Studies* 29, 5-68.

Hirshleifer, David, 2001, Investor Psychology and Asset Pricing, *Journal of Finance* 56, 1533-1597.

Hirshleifer, David, Sonya S. Lim, and Siew Hong Teoh, 2009, Driven to Distraction: Extraneous Events and Underreaction to Earnings News, *Journal of Finance* 64, 2289-2325.

Hirshleifer, David, Sonya S. Lim, and Siew Hong Teoh, 2011, Limited investor attention and stock market misreactions to accounting information, *Review of Asset Pricing Studies* 1(1), 35–73.

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54 (6), 2143–2184.

Hou, Kewei, 2007, Industry Information Diffusion and the Lead-Lag Effect in Stock Returns, *Review of Financial Studies* 20, 1113-1138.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting Anomalies: An Investment Approach, *Review of Financial Studies* 28, 650-705.

Hou, Kewei, Chen Xue, and Lu Zhang, 2020, Replicating Anomalies, *Review of Financial Studies* 33, 2019-2133.

Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou, 2015, Investor Sentiment Aligned: A Powerful Predictor of Stock Returns, *Review of Financial Studies* 28, 791-837.

Jacobs, Heiko, and Sebastian Müller, 2020, Anomalies across the Globe: Once Public, No Longer Existent?, *Journal of Financial Economics* 135, 213-230.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48 (1): 65-91.

Jensen, Theis Ingerslev, Bryan T. Kelly, and Lasse Heje Pedersen, 2023, Is There a Replication Crisis in Finance? *Journal of Finance* 78, 2465-2518.

Jiang, George J., and Tong Yao, 2013, Stock Price Jumps and Cross-Sectional Return Predictability, *Journal of Financial and Quantitative Analysis* 48, 1519-1544.

Kahneman, Daniel, 1973, Attention and Effort, Englewood Cliffs, NJ: Prentice-Hall.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263-291.

Kapadia, Nishad, and Morad Zekhnini, 2019, Do Idiosyncratic Jumps Matter?, *Journal of Financial Economics* 131, 666-692.

Linnainmaa, Juhani T., and Michael R. Roberts, 2018, The History of the Cross-Section of Stock Returns, *Review of Financial Studies* 31, 2606-2649.

Lo, Andrew W., 2017, *Adaptive Markets: Financial Evolution at the Speed of Thought*, Princeton, NJ: Princeton University Press.

Lo, Andrew W., and A. Craig MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* 1, 41-66.

Lo, Andrew W., and Ruixun Zhang, 2024, *The Adaptive Markets Hypothesis: An Evolutionary Approach to Understanding Financial System Dynamics*, Oxford: Oxford University Press.

Lu, Xiaomeng, Robert F. Stambaugh, and Yu Yuan, 2017, Anomalies Abroad: Beyond Data Mining, NBER Working Paper 23809.

Lucca, David O., and Emanuel Moench, 2015, The Pre-FOMC Announcement Drift, *Journal of Finance* 70, 329-371.

McLean, R. David, and Jeffrey Pontiff, 2016, Does Academic Research Destroy Stock Return Predictability? *Journal of Finance* 71, 5-32.

Miller, Edward M., 1977, Risk, Uncertainty, and Divergence of Opinion, *Journal of Finance* 32, 1151-1168.

Muravyev, Dmitriy, Neil D. Pearson, and Joshua M. Pollet, 2022, Is There a Risk Premium in the Stock Lending Market? Evidence from Equity Options, *Journal of Finance* 77, 1787-1828.

Odean, Terrance, 1998, Are Investors Reluctant to Realize Their Losses? *Journal of Finance* 53(5), 1775-1798.

Patton, Andrew J., and Michela Verardo, 2012, Does Beta Move with News? Firm-Specific Information Flows and Learning about Profitability, *Review of Financial Studies*, 2789-2839.

Pontiff, J., 1996, Costly arbitrage: Evidence from closed-end funds. *Quarterly Journal of Economics* 111 (4), 1135–1151.

Pontiff, J., 2006. Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics* 42 (1–2), 35–52.

Savor, Pavel, and Mungo Wilson, 2016, Earnings announcements and systematic risk, *Journal of Finance* 71(1), 83-138.

Schwert, G. William, 2003, Anomalies and Market Efficiency, in G. M. Constantinides, M. Harris, and R. M. Stulz, eds.: Handbook of the Economics of Finance, Vol. 1, Chapter 15, 939-974, Elsevier.

Shleifer, Andrei, and Robert W. Vishny, 1997, The Limits of Arbitrage, *Journal of Finance* 52, 35-55.

Shumway, Tyler, 1997, The delisting bias in CRSP data. *Journal of Finance* 52(1), 327-340.

Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies. *Journal of Financial Economics* 104 (2), 288–302.

Table 1. Summary statistics

The table summarizes the calendar days between the fiscal period end and the announcement or publication of the corresponding financial information (Panel A), and the monthly returns to long-short anomaly strategies (Panel B: equal-weighted; Panel C: value-weighted). The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). Panel A summarizes the calendar days until the reporting of the information in Compustat Snapshot or the release of the 10-Q/10-K report (quarterly signals) or of the 10-K report (annual signals). Panel B reports equal-weighted quarterly and annual accounting signal returns, and Panel C reports value-weighted quarterly and annual accounting signal returns. Each anomaly signal is used to split the CRSP universe into five quintiles such that the top quintile has historically earned higher returns than the bottom quintile. The long-short return is the difference between the top and the bottom quintile returns. In each of the four segments of Panels B and C, the first line considers the sample period before the publication of the corresponding academic paper, and the second line considers the sample period after the publication of the corresponding academic paper. Each panel summarizes all monthly observations for all included anomalies. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

	Mean	**Std. dev.**	**Q1**	**Median**	**Q3**	**N**
Panel A: Days from fiscal period end to disclosure						
Days from fiscal quarter end						
Days until Compustat Snapshot update	39.3	26.7	28.0	35.0	45.0	378,696
Days until 10-K/10-Q release	45.6	20.1	36.0	41.0	46.0	327,364
Days until earlier information date	36.4	15.8	27.0	34.0	41.0	378,696
Days from fiscal year end						
Days until Compustat Snapshot update	63.1	40.7	37.0	60.0	88.0	94,214
Days until 10-K release	69.5	17.2	58.0	70.0	83.0	75,638
Days until earlier information date	55.9	27.9	37.0	57.0	74.0	94,214
Panel B: Average monthly returns: Equal-weighted (%)						
Before publication	0.532	2.941	-0.763	0.425	1.693	5,496
After publication	0.188	2.856	-1.015	0.102	1.311	7,944
Panel C: Average monthly returns: Value-weighted (%)						
Before publication	0.382	3.604	-1.415	0.242	2.077	5,496
After publication	0.138	3.611	-1.487	0.138	1.669	7,944

Table 2. Publication effect for quarterly signals

Panel A provides regression estimates of Equation (1). The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variable in the first column is the raw cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. Dependent variable in the second column is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. *HighSignal* is an indicator variable, set to 1 for firms in the top quintile of the signal and set to 0 for firms in the bottom quintile (the middle three quintiles are dropped from the sample). *PostPublication* is an indicator variable, set to 1 for all the time periods during or after the year of the first academic paper publication concerning the anomaly and set to 0 for all the time periods before the year of the first academic publication concerning the anomaly. Panel A provides baseline regression estimates. Panel B provides tests of high-signal post-publication performance, low-signal post-publication performance, and their difference; high-signal pre-publication performance, low-signal pre-publication performance, and their difference; and difference in differences—the publication effect. All estimates exclude penny stocks (stocks priced below $5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Raw returns	**FF3-adjusted returns**
Panel A: Regression coefficients		
Intercept	0.070	-0.126***
	(0.101)	(0.030)
HighSignal	0.187***	0.190***
	(0.027)	(0.023)
PostPublication	-0.282**	-0.181***
	(0.142)	(0.053)
HighSignal x *PostPublication*	0.484***	0.471***
(Difference-in-differences: Publication effect)	(0.048)	(0.041)
Number of observations	4,240,488	4,152,716
Number of clusters	116	116

Table 2. (*Continued*)

	Raw returns	**FF3-adjusted returns**
Panel B: Linear combinations of regression coefficients		
High signal pre-publication	0.257***	0.065**
(*Intercept* + *HighSignal*)	(0.095)	(0.026)
Low signal pre-publication	0.070	-0.126***
(*Intercept*)	(0.101)	(0.030)
(1) *High - Low signal pre-publication*	0.187***	0.190***
(*HighSignal*)	(0.027)	(0.023)
High signal post-publication	0.458***	0.355***
(sum of all four coefficients from Panel A)	(0.112)	(0.046)
Low signal post-publication	-0.213*	-0.306***
(*Intercept* + *PostPublication*)	(0.119)	(0.048)
(2) *High - Low signal post-publication*	0.671***	0.661***
(*HighSignal* + *HighSignal* x *PostPublication*)	(0.035)	(0.031)
(2) - (1): Publication effect	**0.484*** **	**0.471*** **
(*HighSignal* x *PostPublication*)	**(0.048)**	**(0.041)**
Number of observations	4,240,488	4,152,716
Number of clusters	116	116

Table 3: Predicting factor exposures using anomaly signals

The table provides Equation (1) tests for each risk factor of **(1)** differences between high-signal and low-signal pre-publication risk factor exposure, **(2)** differences between high-signal and low-signal post-publication risk factor exposure, and their difference **(2) - (1)**—the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variables are the stock-level factor exposures (betas) from the Fama-French five-factor models augmented with the Carhart (1997) momentum factor. The factors are the market (MKT - RF), size (SMB), value (HML), profitability (RMW), investment (CMA), and momentum (UMD). Factor exposures are estimated using daily returns over the three-calendar month window starting after the information release date. All estimates exclude penny stocks (stocks priced below $5). Standard errors, reported in parentheses, are clustered by calendar quarter. The sample period is from 1994 to 2022. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Beta					
	MKT - RF	**SMB**	**HML**	**RMW**	**CMA**	**UMD**
High signal pre-publication	0.885***	0.673***	0.064***	-0.188***	0.005	-0.062***
	(0.009)	(0.010)	(0.011)	(0.014)	(0.013)	(0.011)
Low signal pre-publication	0.860***	0.645***	0.075***	-0.135***	0.009	-0.061***
	(0.007)	(0.009)	(0.011)	(0.012)	(0.011)	(0.009)
(1) *High - Low signal pre-publication*	-0.025***	-0.028***	0.011**	0.053***	0.004	0.001
	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)	(0.004)
High signal post-publication	0.921***	0.691***	0.010	-0.143***	0.000	-0.072***
	(0.005)	(0.006)	(0.008)	(0.011)	(0.010)	(0.008)
Low signal post-publication	0.913***	0.687***	0.025***	-0.135***	0.019**	-0.056***
	(0.005)	(0.006)	(0.009)	(0.009)	(0.009)	(0.007)
(2) *High - Low signal post-publication*	-0.008***	-0.004	0.015***	0.008	0.019***	0.016***
	(0.002)	(0.003)	(0.006)	(0.006)	(0.006)	(0.005)
(2) - (1): Publication effect	0.017***	0.024***	0.004	-0.045***	0.015*	0.015**
	(0.005)	(0.005)	(0.007)	(0.008)	(0.008)	(0.006)
Number of Observations	4,267,190	4,267,190	4,267,190	4,267,190	4,267,190	4,267,190
Number of Clusters	116	116	116	116	116	116

Table 4: Predicting standardized unexpected earnings (SUE) using anomaly signals

The table provides Equation (1) tests of **(1)** difference between high-signal and low-signal pre-publication performance, **(2)** difference between high-signal and low-signal post-publication performance, and their difference **(2) - (1)**—the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. The dependent variable is the stock-level future standardized unexpected earnings (SUE), based on the first earnings announcement made after the information release date. SUE is the difference between the announced earnings excluding special items and expected earnings using a seasonal random walk model. All models exclude penny stocks (stocks priced below \$5). Standard errors, reported in parentheses, are clustered by calendar quarter. The sample period is from 1994 to 2022. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	SUE
(1) *High - Low signal pre-publication*	0.152***
	(0.022)
(2) *High - Low signal post-publication*	0.176***
	(0.027)
(2) - (1): Publication effect	**0.024**
	(0.032)
Number of observations	2,873,905
Number of clusters	116

Table 5: Earnings predictability and the publication effect

The table provides Equation (1) tests of **(1)** differences between high-signal and low-signal pre-publication performance, **(2)** differences between high-signal and low-signal post-publication performance, and their differences **(2) - (1)**—the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. The dependent variable in Panel A is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. The dependent variable in Panel B is the stock-level future standardized unexpected earnings (SUE), based on the first earnings announcement made after the information release date. SUE is the difference between the announced earnings excluding special items and expected earnings using a seasonal random walk model. The first column of each panel restates the full anomaly sample estimates from Tables 2 and 4. In the remaining columns of each panel, anomalies are split into quintiles based on the way they predict SUE before publication. The five columns correspond to these anomaly quintiles, ordered by the coefficient associated with *HighSignal* in regressions analogous to those in Table 4 (estimated separately for each anomaly and plotted in Figure 4, from the lowest (Q1) to the highest (Q5)). All estimates exclude penny stocks (stocks priced below \$5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Panel A: FF3-adjusted returns (quintiles according to SUE)	**All anomalies (Table 2)**	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
(1) *High - Low pre-publication*	0.190***	-0.401***	0.149***	0.075***	0.253***	0.736***
	(0.023)	(0.053)	(0.033)	(0.021)	(0.040)	(0.082)
(2) *High - Low post-publication*	0.661***	-1.242***	0.351***	0.318***	1.126***	2.557***
	(0.031)	(0.059)	(0.053)	(0.037)	(0.057)	(0.089)
(2) - (1): Publication effect	**0.471*** **	**-0.842*** **	**0.202*** **	**0.243*** **	**0.874*** **	**1.821*** **
	(0.041)	**(0.080)**	**(0.061)**	**(0.038)**	**(0.074)**	**(0.120)**
Number of observations	4,152,716	659,693	887,853	922,028	887,653	795,489

Panel B: SUE (quintiles according to SUE)	**All anomalies (Table 4)**	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
(1) *High - Low pre-publication*	0.152***	-0.833***	-0.158***	0.144***	0.381***	1.008***
	(0.022)	(0.109)	(0.040)	(0.053)	(0.051)	(0.116)
(2) *High - Low post-publication*	0.176***	-0.437***	-0.131**	0.170**	0.309***	0.941***
	(0.027)	(0.076)	(0.064)	(0.086)	(0.044)	(0.097)
(2) - (1): Publication effect	**0.024**	**0.396*** **	**0.027**	**0.026**	**-0.072**	**-0.067**
	(0.032)	**(0.110)**	**(0.057)**	**(0.072)**	**(0.066)**	**(0.127)**
Number of observations	2,873,905	462,772	620,516	632,262	624,604	552,807

Table 6. Mechanical trading frictions and the publication effect

The table provides Equation (1) estimates of the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variable is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. For each stock-level mechanical trading friction proxy (market capitalization, Amihud illiquidity, bid-ask spread, institutional ownership, and short interest), the sample is split into five quintiles, with the quintile cutoffs determined from the sample of all stocks. The ordering of quintiles for each proxy of mechanical trading frictions is such that the leftmost column features the easiest quintile—the quintile with the highest predicted ability to exploit the anomaly—and the rightmost column features the most challenging (hardest) quintile—the quintile with the lowest predicted ability to exploit the anomaly. All estimates exclude penny stocks (stocks priced below \$5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Q_Easiest	**Q2**	**Q3**	**Q4**	**Q_Hardest**
Transaction costs					
Market capitalization	0.319***	0.466***	0.624***	0.537***	0.348***
(Easiest = Largest)	(0.036)	(0.050)	(0.065)	(0.061)	(0.096)
Amihud illiquidity	0.321***	0.551***	0.613***	0.449***	0.290***
(Easiest = Least illiquid)	(0.038)	(0.053)	(0.068)	(0.057)	(0.066)
Bid-ask spread	0.323***	0.573***	0.633***	0.393***	0.274***
(Easiest = Smallest spread)	(0.042)	(0.052)	(0.064)	(0.066)	(0.071)
Shorting frictions					
Institutional ownership	0.497***	0.507***	0.440***	0.476***	0.333***
(Easiest = Highest)	(0.050)	(0.057)	(0.054)	(0.061)	(0.072)
Short interest	0.479***	0.436***	0.409***	0.354***	0.300***
(Easiest = Highest)	(0.071)	(0.061)	(0.055)	(0.047)	(0.064)

Table 7. Inattention, sentiment, and the publication effect

The table provides Equation (1) estimates of the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variable is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. For inattention, subsamples are based on analyst coverage (the first two columns) and on media coverage (the last two columns). Observations are split into those featuring stocks with and without coverage—those that have higher and lower predicted ability to exploit the anomaly, respectively. For sentiment, subsamples are based on the sentiment index of Baker and Wurgler (2006). The subsamples correspond to months with above-median and below-median sentiment. The split reported in the first two columns relies on lagged sentiment (month t - 1) and the split reported in the last two columns relies on contemporaneous sentiment (month t). All estimates exclude penny stocks (stocks priced below \$5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Analyst coverage		**Media coverage**	
	No coverage	Coverage	No coverage	Coverage
Inattention	0.481***	0.400***	0.358***	0.271***
(Easier = Coverage)	(0.043)	(0.050)	(0.050)	(0.056)
	Orthogonalized sentiment (*t* - 1)		**Orthogonalized sentiment (*t*)**	
	Below median	Above median	Below median	Above median
Sentiment	0.439***	0.384***	0.447***	0.395***
(Easier = Below median)	(0.049)	(0.053)	(0.049)	(0.054)

Table 8. Prospect theory and the publication effect

The table provides Equation (1) estimates of the publication effect. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variable is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. The table focuses on idiosyncratic volatility, idiosyncratic skewness, and capital gains overhang, the three Prospect Theory variables that Barberis, Jin, and Wang (2021) argue are related to anomalies: idiosyncratic volatility, idiosyncratic skewness, and capital gains. For each of these stock-level friction proxies, the sample is split into five quintiles. Idiosyncratic volatility is estimated using daily FF3-adjusted returns over the previous month, idiosyncratic skewness is estimated using daily FF3-adjusted returns over the previous quarter, and capital gains are estimated using daily data from the previous year (260 trading days). All estimates exclude penny stocks (stocks priced below $5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Q1	Q2	Q3	Q4	Q5
Idiosyncratic volatility	0.337***	0.505***	0.557***	0.488***	0.511***
	(0.038)	(0.042)	(0.059)	(0.073)	(0.105)
Idiosyncratic skewness	0.655***	0.397***	0.349***	0.449***	0.485***
	(0.063)	(0.048)	(0.052)	(0.053)	(0.068)
Capital gains overhang	0.668***	0.559***	0.457***	0.424***	0.410***
	(0.107)	(0.065)	(0.055)	(0.047)	(0.055)

Table 9. Publication effect for annual signals

Panel A provides regression estimates of Equation (1). The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-K (listed and described in Internet Appendix Table IA1.1). The event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K report. Dependent variable in the first column is the raw cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. Dependent variable in the second column is the FF3-adjusted cumulative stock return over a three-day window (t - 1, t + 1) surrounding the information release date t. *HighSignal* is an indicator variable, set to 1 for firms in the top quintile of the signal and set to 0 for firms in the bottom quintile (the middle three quintiles are dropped from the sample). *PostPublication* is an indicator variable, set to 1 for all the time periods during or after the year of the first academic paper publication concerning the anomaly and set to 0 for all the time periods before the year of the first academic publication concerning the anomaly. Panel A provides baseline regression estimates. Panel B provides tests of high-signal post-publication performance, low-signal post-publication performance, and their difference; high-signal pre-publication performance, low-signal pre-publication performance, and their difference; and difference in differences—the publication effect. All estimates exclude penny stocks (stocks priced below $5). The sample period is from 1994 to 2022. Standard errors, reported in parentheses, are clustered by calendar quarter. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Raw returns	**FF3-adjusted returns**
Panel A: Regression		
Intercept	-0.203	0.001
	(0.173)	(0.045)
HighSignal	0.088***	0.049**
	(0.031)	(0.020)
PostPublication	0.227	0.037
	(0.283)	(0.099)
HighSignal x *PostPublication*	-0.036	0.023
(Difference-in-differences: Publication effect)	(0.058)	(0.041)
Number of observations	1,020,214	994,193
Number of clusters	29	29

Table 9. (*Continued*)

	Raw returns	**FF3-adjusted returns**
Panel B: Post-estimation		
High signal pre-publication	-0.203	0.001
(*Intercept* + *HighSignal*)	(0.173)	(0.045)
Low signal pre-publication	-0.115	0.050
(*Intercept*)	(0.167)	(0.032)
(1) *High - Low signal pre-publication*	0.088***	0.049**
(*HighSignal*)	(0.031)	(0.020)
High signal post-publication	0.024	0.038
(sum of all four coefficients from Panel A)	(0.234)	(0.094)
Low signal post-publication	0.076	0.110
(*Intercept* + *PostPublication*)	(0.196)	(0.076)
(2) *High - Low signal post-publication*	0.052	0.071*
(*HighSignal* + *HighSignal* x *PostPublication*)	(0.048)	(0.037)
(2) - (1): Publication effect	**-0.036**	**0.023**
(*HighSignal* x *PostPublication*)	**(0.058)**	**(0.041)**
Number of observations	1,020,214	994,193
Number of clusters	29	29

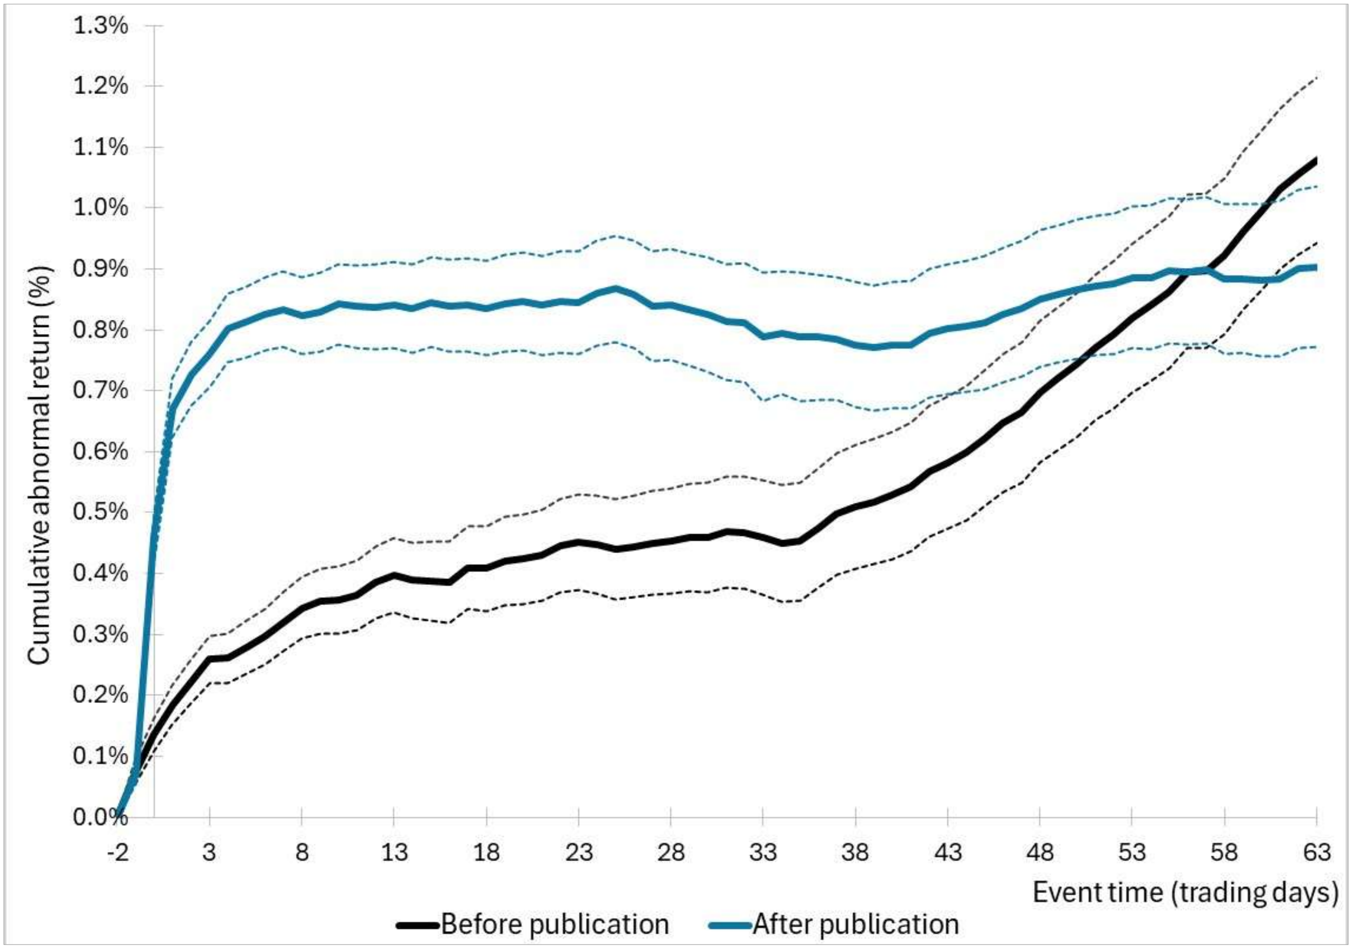

Fig. 1. This figure displays the average of the cumulative returns across all 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The figure displays the cumulative return to a long-short strategy in event time using a 63-trading-day window. For each firm/signal, day zero is the earlier of the release of the financial statement via SEC's EDGAR and the signal's availability via Compustat Snapshot. The figure depicts the difference in the averages of the cumulative returns across stocks in *HighSignal* portfolios and in *LowSignal* portfolios in event time. For each anomaly and each quarter, we construct a cumulative return for each stock in event time. If stock *i* announces a new value of the signal for anomaly *a* (high or low) at day *t*, we build cumulative returns $\boldsymbol{r}_{i,a,\tau}$, $\tau = t$-2, t-1, t, t+1, …, t+63, setting $\boldsymbol{r}_{i,a,t-2} = 0$. We proceed to calculate subsequent event-time cumulative returns as $\boldsymbol{r}_{i,a,\tau} = (1+\boldsymbol{r}_{i,a,\tau-1})(1+r_{i,a,\tau}) - 1$, $\tau = t$-1, t, t+1, …, t+63, where $r_{i,a,\tau}$ is the delisting-adjusted return for stock *i* on day τ. To compare signal reactions before and after publication, we split the sample into observations before academic publication and after academic publication. We then combine all pre-publication and post-publication samples across all anomalies to obtain one combined pre-publication subsample and one post-publication subsample. For each combined subsample, we average cumulative returns for *HighSignal* and *LowSignal* stocks in each quarter, taking the averages of all cumulative returns $r_{i,t,\tau}$ for $\tau = t$-2, t, t+1, …, t+63 for stocks with high signals ($r_{H,\tau}$) and low signals ($r_{L,\tau}$), respectively. The plot shows the values $r_{H,\tau}$ - $r_{L,\tau}$ for $\tau = t$-2, t, t+1, …, t+63 both before and after publication. The black line uses anomaly observations over the sample period before or during the year the academic research regarding the specific anomaly was published. The blue line uses anomaly observations over the sample period after the academic research was published. The dotted lines provide 5% confidence intervals, calculated for each day in event time using standard errors clustered by calendar quarter. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

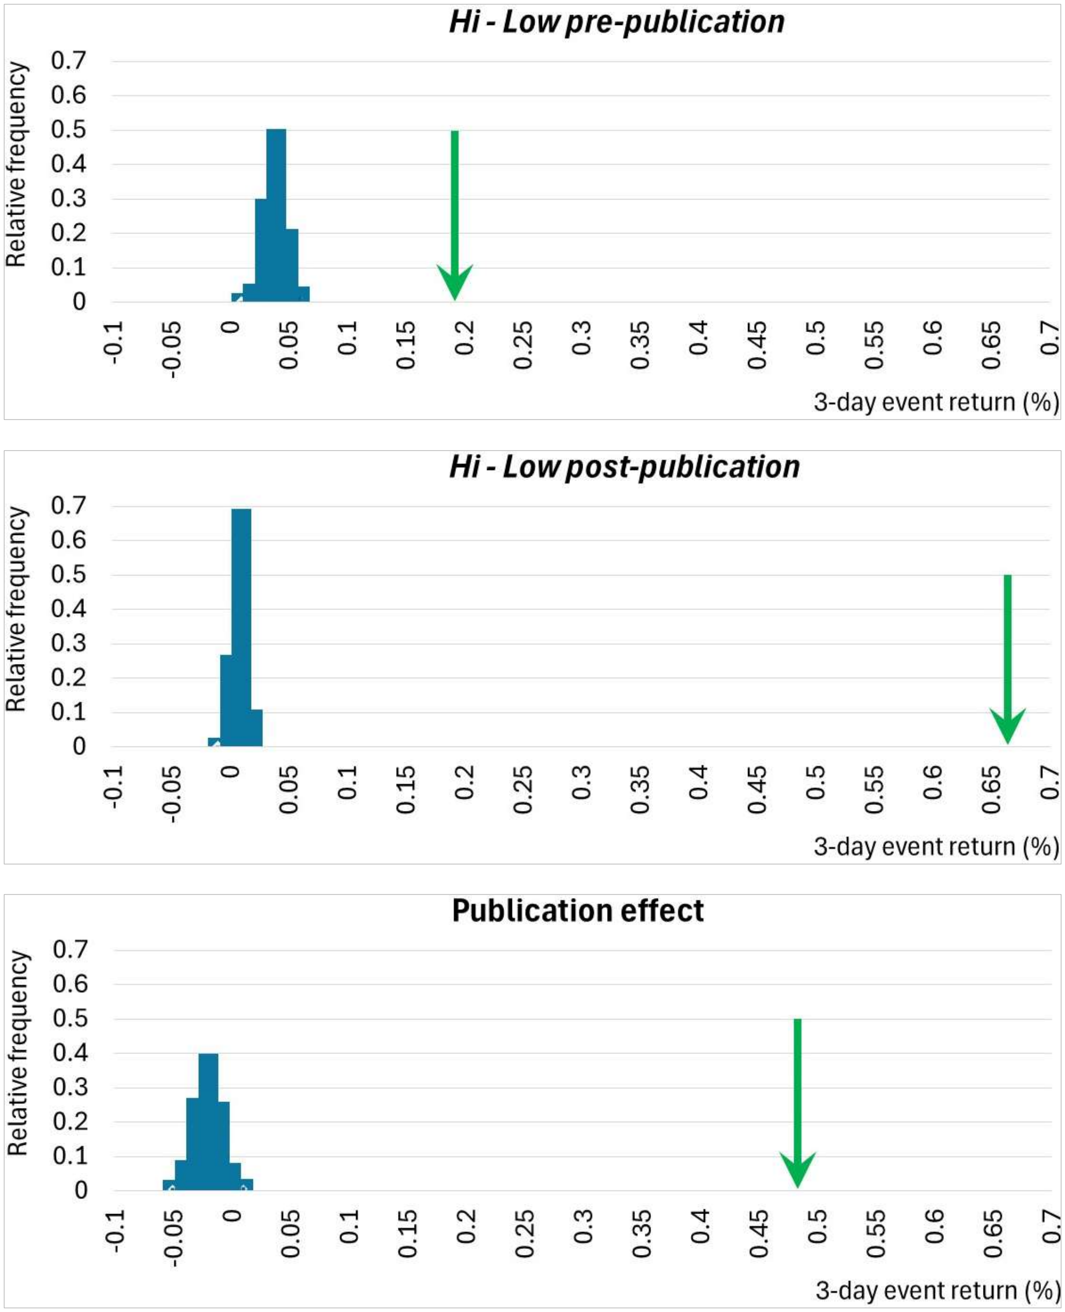

Fig. 2: This figure shows distributions of pre-publication high-low signal 3-day event returns (top panel), post-publication high-low signal 3-day event returns (middle panel), and the publication effect (bottom panel) from 1,000 simulation runs. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). Each simulation run picks for each firm-quarter observation a date at random between 6 and 30 trading days before or after the actual signal date and estimates the regression from Equation (1). The actual event date is defined as the earlier of (a) the earliest date on which all necessary variables for signal calculation become available from Compustat Snapshot, and (b) the release date of the 10-K/10-Q report. Dependent variables in all simulations are the FF3-adjusted cumulative stock returns over three-day windows (t - 1, t + 1) surrounding the simulated information release dates t. In each panel, the green arrow denotes the corresponding point estimate based on the actual information release dates, as reported in the second column of Table 2. The sample period is from 1994 to 2022.

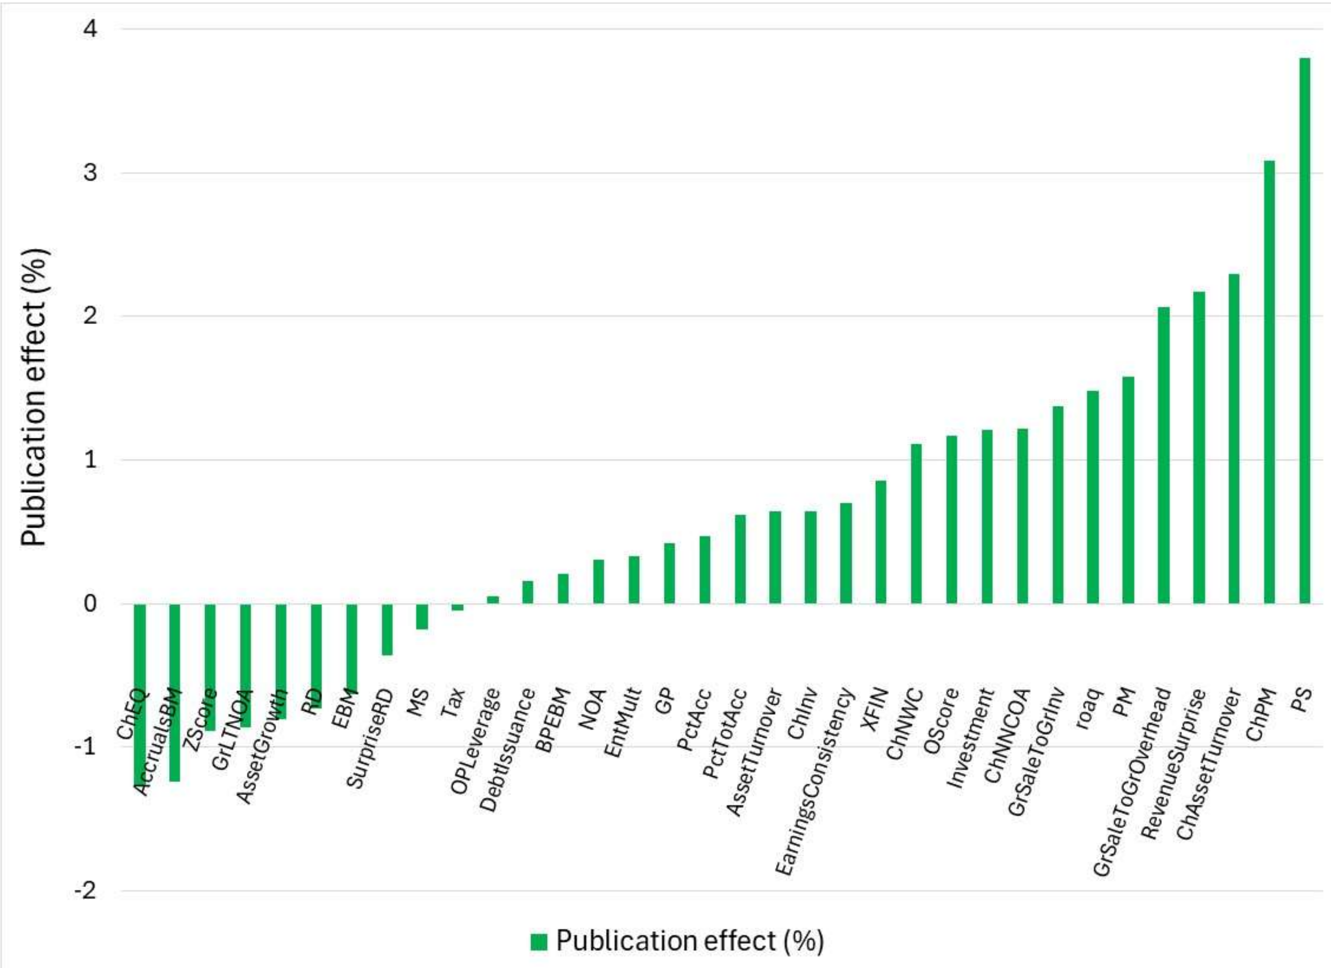

Fig. 3. This figure displays estimates of the publication effect for each anomaly signal, estimated from Equation (1) anomaly-by-anomaly. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The sample excludes penny stocks (price below \$5). The sample period is from 1994 to 2022.

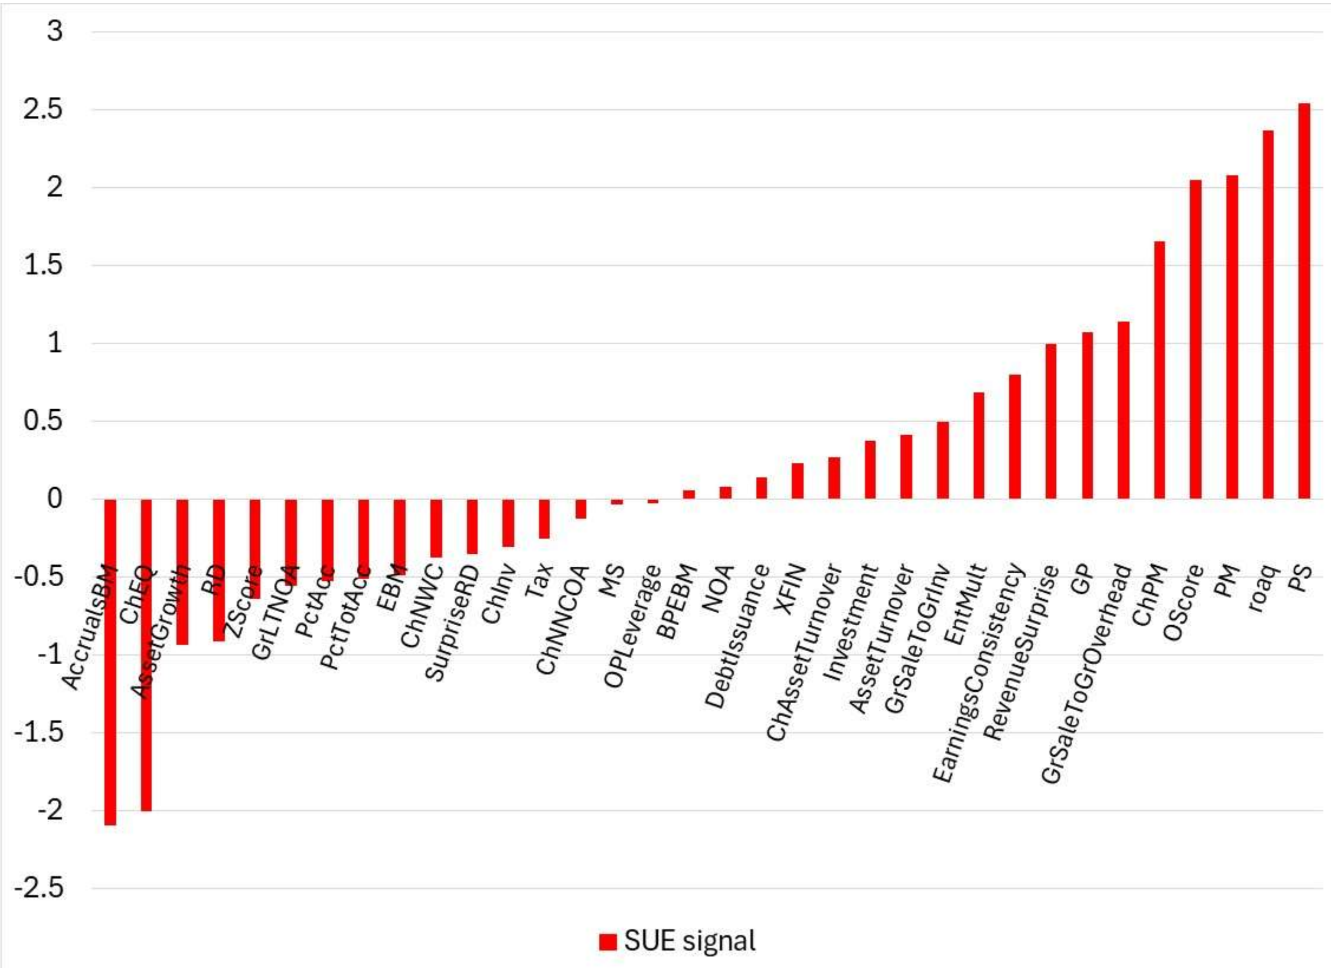

Fig. 4. This figure displays the earnings predictability estimates, estimated anomaly-by-anomaly and sorted from the lowest to the highest. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). All estimations are rooted in Equation (1), where estimations concerning earnings predictability feature as the dependent variable standardized unexpected earnings SUE over the subsequent quarter. SUE is the difference between the announced earnings excluding special items and expected earnings using a seasonal random walk model. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

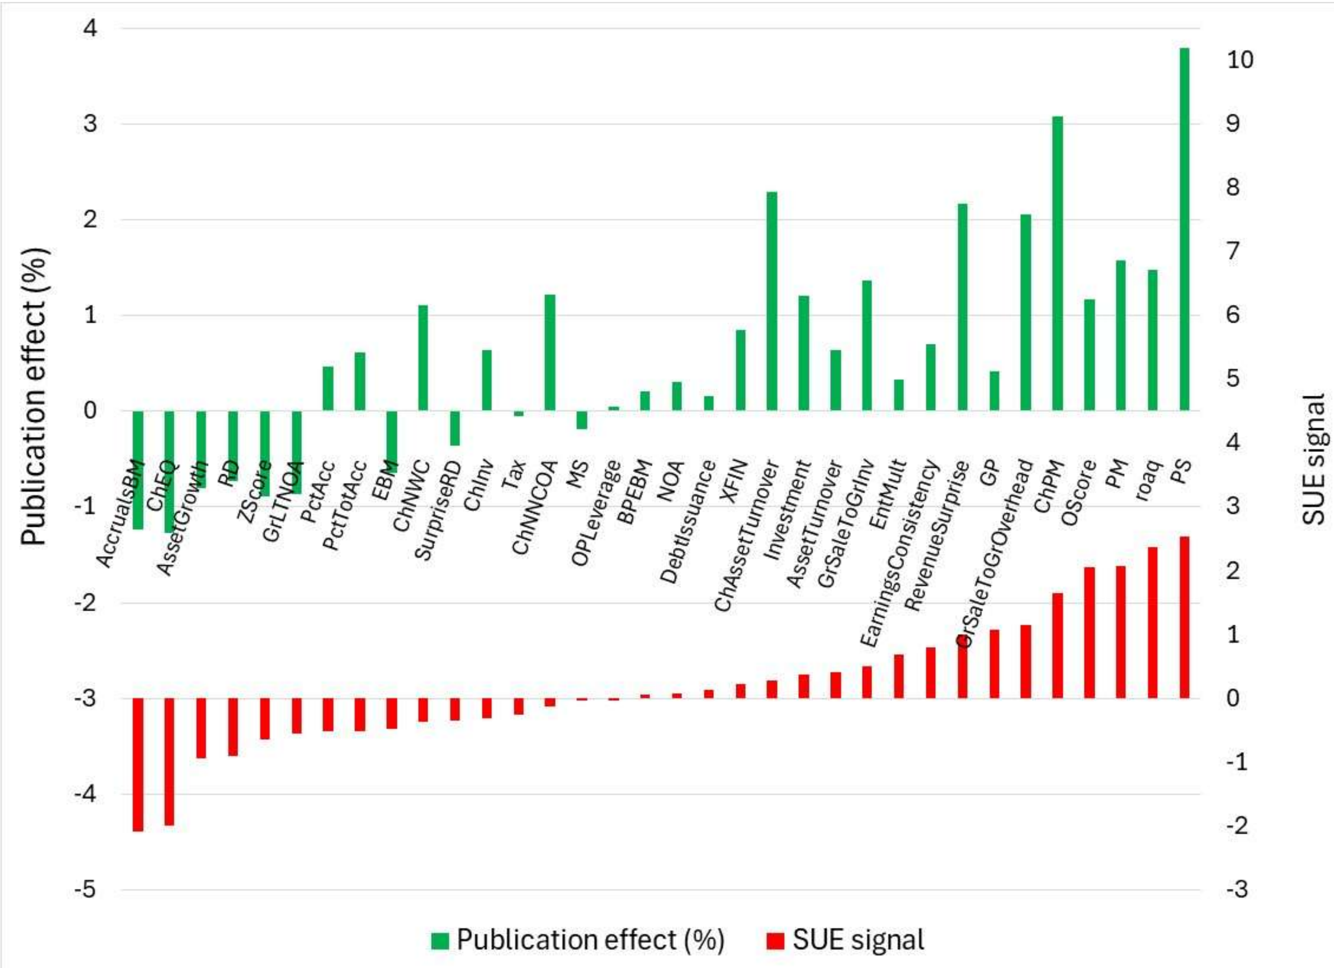

Fig. 5. This figure displays estimates of earnings predictability and the publication effect for each anomaly. The analysis encompasses 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). Earnings predictability estimates are shown at the bottom of the figure (red bars), using the right y-axis, and the publication effects are shown at the top of the figure (green bars), using the left y-axis. All estimations are rooted in Equation (1), where estimations concerning earnings predictability feature as the dependent variable the standardized unexpected earnings (SUE) over the subsequent quarter. SUE is the difference between the announced earnings excluding special items and expected earnings using a seasonal random walk model. The dependent variable for the publication effect regressions, as usual, is the three-day return surrounding the information announcement date. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

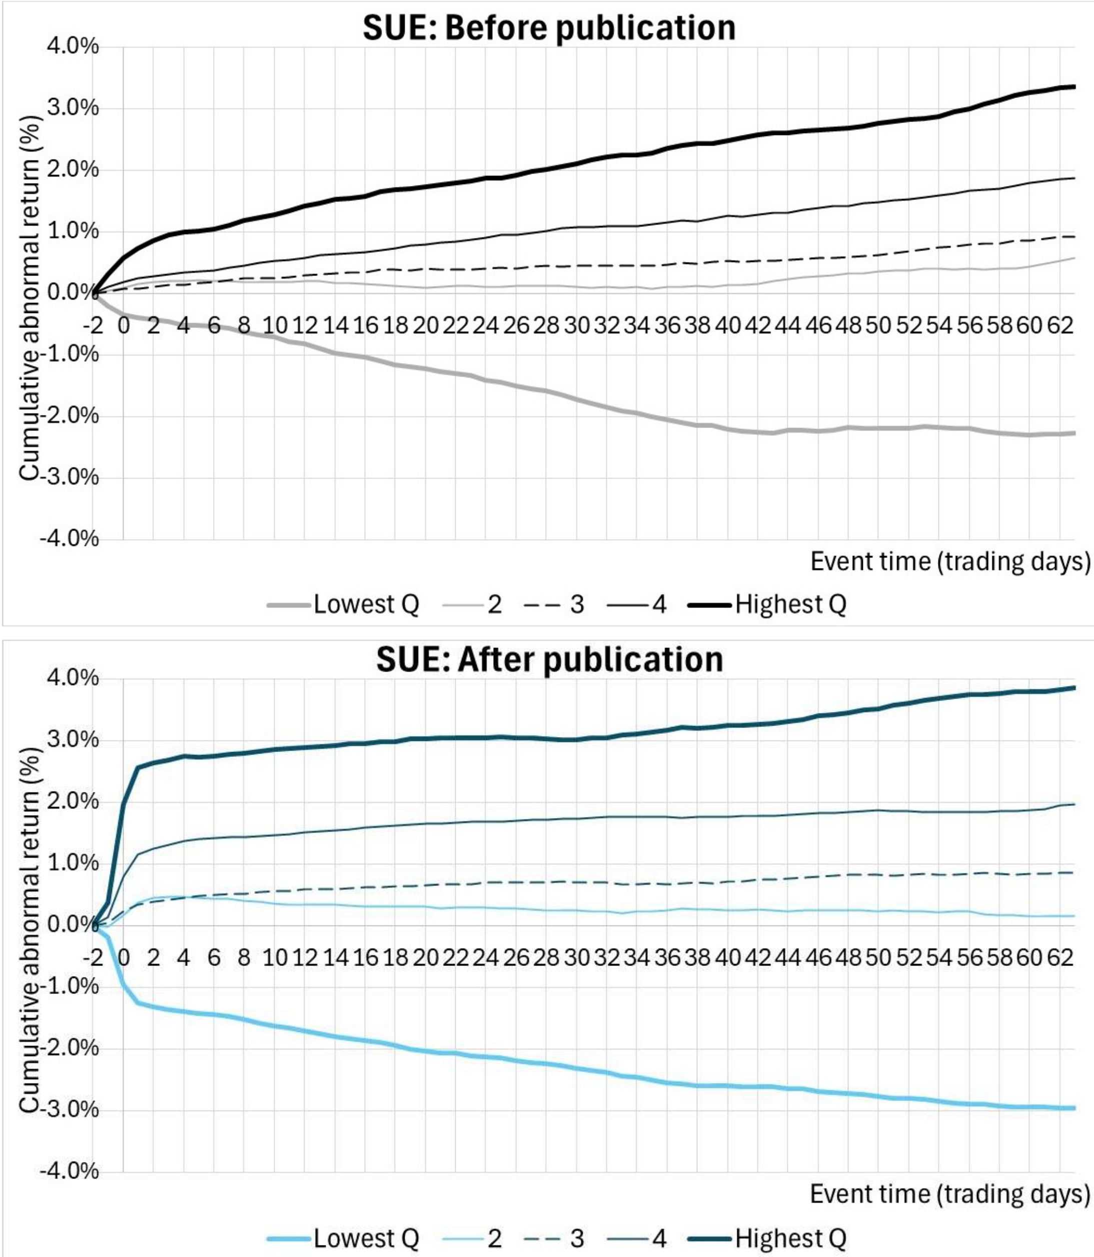

Fig. 6. This figure displays the average of the cumulative returns for anomaly quintiles from the set of 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). As in Table 5, anomalies are split into quintiles based on the way they predict SUE before publication. The five columns correspond to these anomaly quintiles, ordered by the coefficient associated with *HighSignal* in regressions analogous to those in Table 4 (estimated separately for each anomaly and plotted in Figure 4, from the lowest (Q1) to the highest (Q5)). Top panel focuses on pre-publication cumulative returns, whereas the bottom panel focuses on post-publication cumulative returns. Each line in the figure (black and grey in the top, pre-publication panel, and dark and light blue in the bottom, post-publication panel) represents the cumulative return to a long-short strategy in event time using a 63-trading-day window based on the corresponding anomaly quintile. For each firm/signal, day zero is the earlier of the release of the financial statement via SEC's EDGAR and the signal's availability via Compustat Snapshot. The figure depicts the difference in the averages of the cumulative returns across stocks in *HighSignal* portfolios and in *LowSignal* portfolios in event time. Cumulative return construction is described in Section 3.1 and in the caption to Fig. 1. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

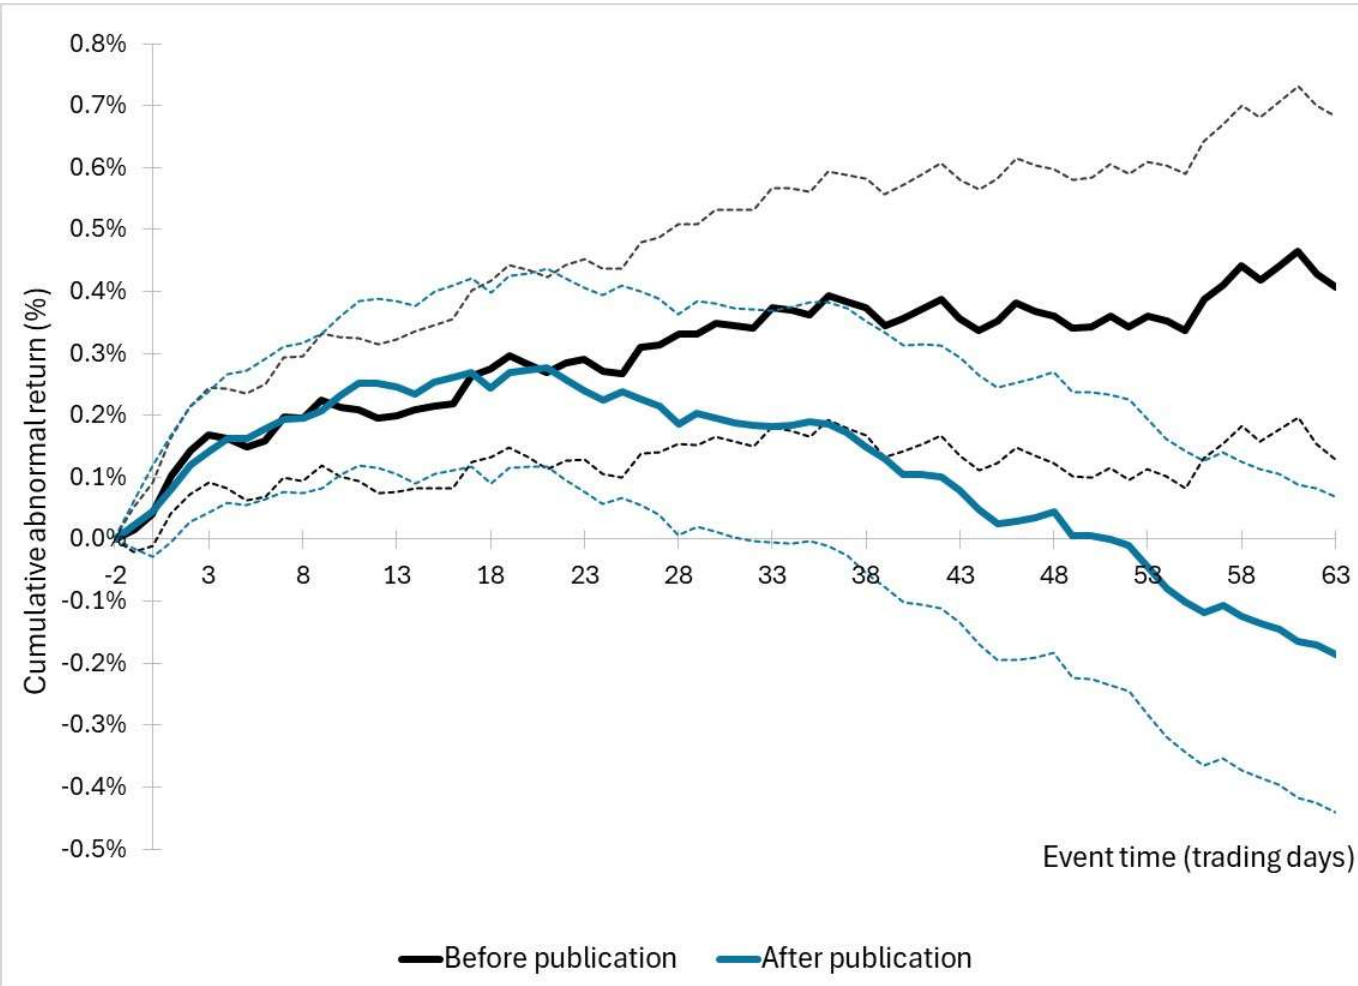

Fig. 7. This figure displays the average of the cumulative returns across all 34 accounting anomalies from McLean and Pontiff (2016) that rely on publicly traded companies' financial information released through Forms 10-Q and 10-K (listed and described in Internet Appendix Table IA1.1). The figure displays the cumulative return to a long-short strategy in event time using a 63-trading-day window. For each firm/signal, day zero is the earlier of the release of the 10-K via SEC's EDGAR and the signal's availability via Compustat Snapshot. The figure depicts the difference in the averages of the cumulative returns across stocks in *HighSignal* portfolios and in *LowSignal* portfolios in event time. The black line uses anomaly observations over the sample period before or during the year the academic research regarding the specific anomaly was published. The blue line uses anomaly observations over the sample period after the academic research was published. The dotted lines provide 5% confidence intervals, calculated for each day in event time using standard errors clustered by calendar year. Cumulative return construction is described in Section 3.1 and in the caption to Fig. 1. The sample excludes penny stocks (price below $5). The sample period is from 1994 to 2022.

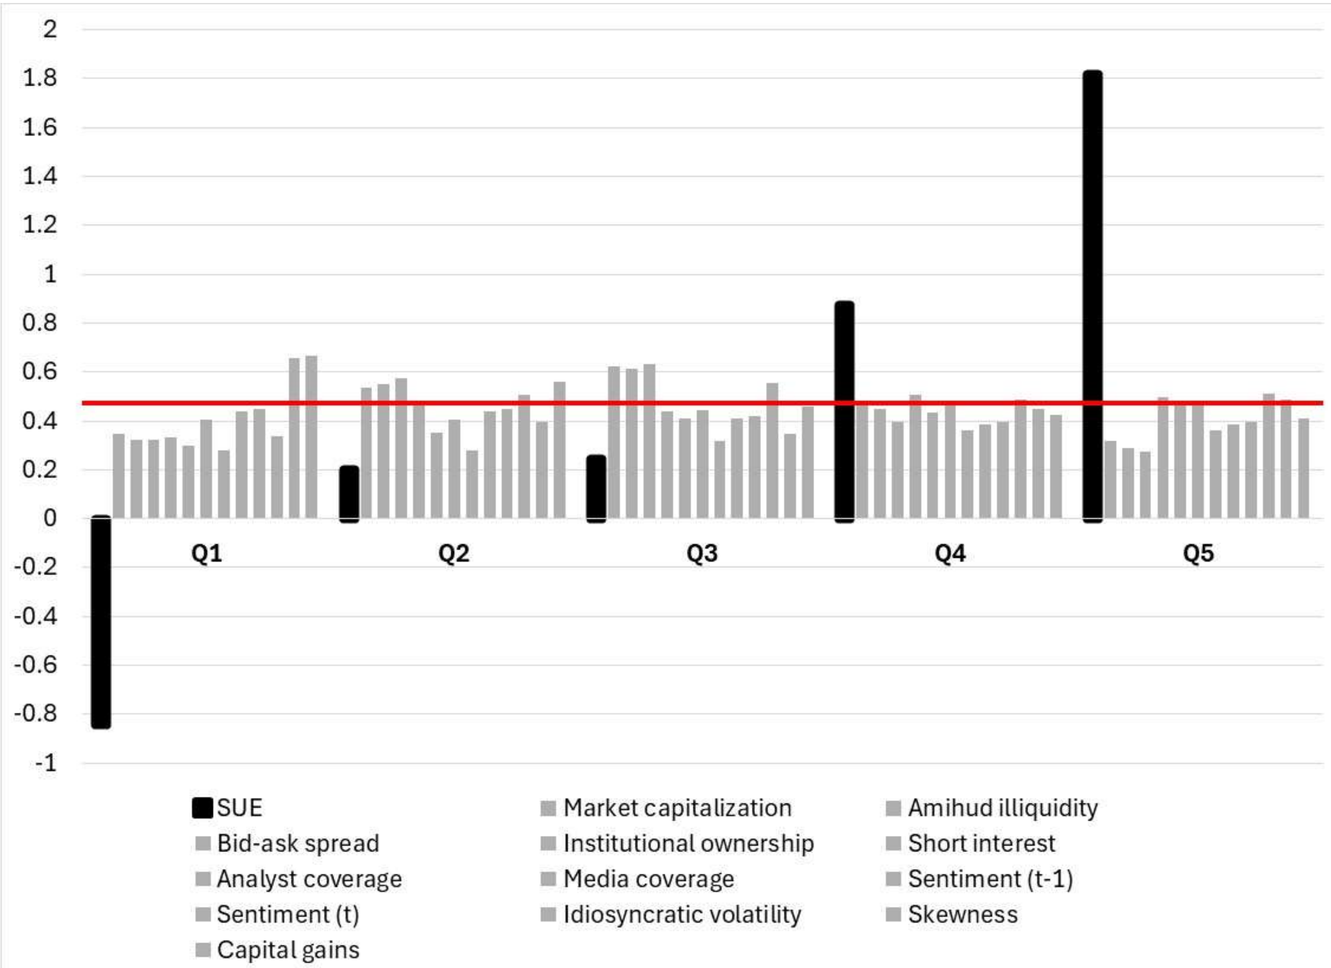

Fig. 8. This figure summarizes the publication effect estimates reported in Tables 2, 5, 6, 7, and 8. The red line depicts the baseline publication effect from Table 2. The black bar represents the quintile publication effects associated with standardized unexpected earnings SUE, as reported in Table 5 by quintile subsamples. Grey bars represent the quintile publication effects associated with all the friction measures reported in Tables 6, 7, and 8 by quintile subsamples (for the purposes of the figure, inattention and sentiment publication effects from Table 7 are spread across the quintiles so that the estimates for the below median/low subsamples are assigned as the values for the bottom two quintiles, the estimates for the above median/high subsamples are assigned as the values for the top two quintiles, and the middle quintiles are populated with their respective averages).

Appendix

A.1. Theoretical motivation

In the context of anomalies, one strand of explanations for return predictability appeals to some combination of behavioral biases and market frictions that, ultimately, prevent prices from reflecting new information. Another strand contends that prices reflect new information in a timely fashion, but statistical issues lead researchers to conclude erroneously that the return patterns are anomalous. Whereas modeling each explanation precisely is beyond the scope of this paper, we consider the two sets of explanations side by side, focusing on the common elements of the mechanism underlying them.

To that end, we develop two parsimonious models to derive implications for the behavioral and rational explanations. Common to the two models is the assumption of a risk-neutral representative investor with a discount rate of 0. This assumption, while unrealistic, allows us to focus on the effect of new information on asset prices without the confounding effects of risk considerations. Also common to the two models is the meaning of "publication," corresponding to the first academic publication date, one that defines the *PostPublication* indicator in Equation (1). The models distinguish between immediate, disclosure-date returns when the signal arrives, and longer-horizon returns, realized as fundamentals are learned over time. These returns map into three-day ($t-1$, $t+1$) disclosure-window returns and the delayed-formation portfolio returns, respectively.

Throughout the paper, we refer to the accounting-based return predictors as anomaly signals, and we align them with the exact disclosure dates of the underlying accounting information. In this Appendix, we use the generic term "signal" for the same object; "signal arrival" should be read as the information-relevant disclosure date used in the empirical design.

In the first model, markets initially misreact—specifically, underreact—to an otherwise valid signal about the fundamental value (cash flows) of the firm. Market participants are initially unaware of the signal's association with fundamentals, leading to the underreaction, that is, the sluggish price adjustment. Academic research highlights this misreaction and induces investors to update their beliefs about the validity of the signal as a predictor of fundamental asset value. This Bayesian updating eliminates the underreaction and, post-publication, asset prices change immediately upon the observation of the signal. Under this model, documenting the original misreaction in academic research and the associated exploitation of the return patterns by practitioners lead to the eventual disappearance of the anomaly.

The second model is a rational model in which the candidate signal is uninformative about fundamentals. Multiple testing and data mining in a finite sample lead researchers erroneously to attribute predictability to such a signal, even though the true data-generating process implies no link to value. However, rational market participants ignore this spurious relation and, over a sufficiently long period, it becomes evident that the signal was uninformative. Under these circumstances, the dissemination of academic research does not alter investor beliefs, so the signal fails to predict returns out-of-sample.

A.2. The behavioral model

The representative investor in this model uses a Bayesian updating process to learn about the correlation of a signal with the fundamental value of the asset. The representative investor behaves rationally given the investor's information set, but the asset price deviates from that of the rational framework. This deviation is more pronounced when the representative investor has diffuse priors about the value of the observed signal, that is, before the representative investor observes any data. We first consider a single asset and define the timing of information and the investor pricing function. We then expand the model to consider multiple assets and the learning process that characterizes the dynamic evolution of the price function.

This Bayesian-learning setup is intended as a stylized representation of the learning mechanism discussed in the body of the paper. In the main text, we use "misreaction" as shorthand for sluggish incorporation of value-relevant information that can arise from limited attention or costly information acquisition; here, that sluggishness is captured by diffuse priors about the signal's mapping into fundamentals. Academic publication (and diffusion of the associated evidence) narrows that uncertainty and reduces interpretation costs, shifting price discovery toward the disclosure date (the timeline compression and near step-function patterns in Figures 1 and 6 in the main text).

A.2.1 One asset

In this section, we outline the timeline for a single asset with a stochastic payoff and derive the returns to the asset at two points in time: the arrival of a signal ($t = s_0$) and the realization of the payout ($t = l_0$).[16]

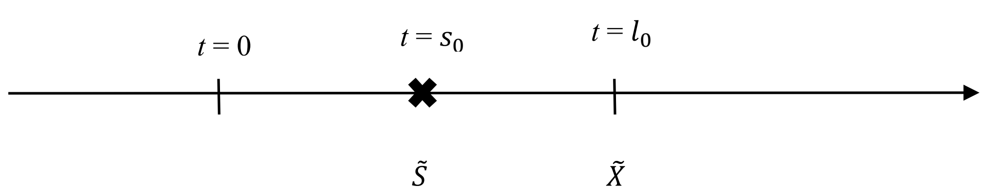

Figure A.1. The timeline for a single asset. The asset is created at time $t = 0$. The risk-neutral representative investor with a discount rate of 0 observes a signal $\tilde{S}$ about the asset at $t = s_0 \in (0, l_0)$. The risky asset payout at $t = l_0$ is $\tilde{X}$.

---

[16] Subscripts s and l are intended to convey a short-term effect and a long-term effect, respectively. The model's "short-run" return (s) at signal arrival corresponds to the disclosure-window return around the information-relevant date, measured over the (t−1, t+1) window. The model's "long-run" return (l) corresponds to the drift/holding-period component captured by traditional anomaly portfolio formation with delayed implementation.

As shown in Figure 1, the asset is created at time $t = 0$. The risk-neutral representative investor with a discount rate of 0 observes a signal $\tilde{S}$ about the asset at $t = s_0 \in (0, l_0)$. Without loss of generality, let $\tilde{S} \sim N(0, 1)$. The risky asset payout at $t = l_0$ is $\tilde{X}$. We assume that the representative investor knows that the relation between the payout $\tilde{X}$ and the signal $\tilde{S}$ is linear:

$$\tilde{X} = \alpha + \beta \times \tilde{S} + \tilde{\epsilon}, \qquad \tilde{S} \perp \tilde{\epsilon}, \qquad \tilde{\epsilon} \sim N(0,1).$$

We further assume that the representative investor knows the true value of the parameter $\alpha = \bar{x}$. At any point in time t, the representative investor prices the asset as $P_t = E_t[\tilde{X}]$. Therefore:

$$P_t = \begin{cases} \bar{X} & \text{if } \; t < s_0 \\ \bar{X} + E_t[\beta] \times \tilde{S} & \text{if } \; t \in [t_0, l_0) \\ \tilde{X} & \text{if } \; t = l_0 \end{cases}$$

Based on this price function, the returns at $t = s_0$ and $t = l_0$ are:

$$R_s = \left.\frac{\partial P_t}{P_t}\right|_{t=s_0} = \frac{E_t[\beta] \times \tilde{S}}{\bar{x}}$$

and

$$R_l = \left.\frac{\partial P_t}{P_t}\right|_{t=l_o} = \frac{\tilde{X} - \bar{X} - E_t[\beta] \times \tilde{S}}{\bar{X} + E_t[\beta] \times \tilde{S}}.$$

A.2.2 Sequentially created multiple assets

The key to the learning hypothesis is that investors update their beliefs about the informativeness of the signal upon observing multiple assets. This updating of beliefs results in an evolution of prices that eliminates the signal's predictive ability. To capture this evolution, we consider a series of risky assets, denoted by the subscript n, $n \in \{1,2, \dots, N\}$. As shown in Figure 2, each asset n is created at time $t = n$. The risk-neutral representative investor receives a signal S_n about asset n at time $t = s_n$, $n < s_n < l_n$. Each asset n delivers a risky payout $\tilde{X}_n$ at $t = l_n$.

In the context of the paper, the sequence of assets can be interpreted as repeated firm-quarter realizations of a given signal (and, more broadly, repeated realizations across many firms). The posterior updating captures the idea that, as evidence accumulates or is disseminated through academic publication, investors become more certain about the sign and magnitude of the signal's mapping into fundamentals. Interpreting the model's payoff as a reduced-form proxy for fundamentals (cash flows), the parameter linking signals to payoffs corresponds to our cash-flow predictability results in Section 4.1, including the fact that some signals predict subsequent cash flows negatively, whereas others predict them positively.

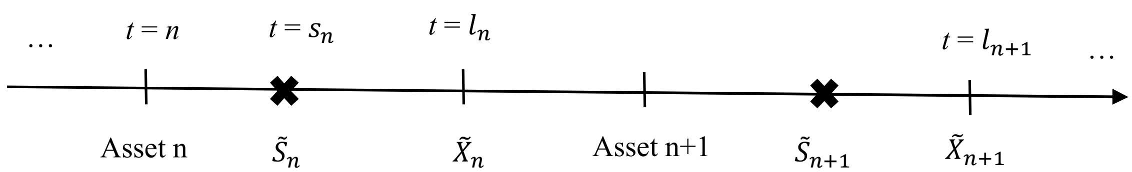

Figure 2. The timeline for multiple assets. There is a series of risky assets, denoted by the subscript n, $n \in \{1,2,\dots,N\}$. Each asset n is created at time $t = n$. The risk-neutral representative investor receives a signal S_n about asset n at $t = s_n$, $n < s_n < l_n$. Each asset n delivers a risky payout $\tilde{X}_n$ at $t = l_n$.

The price function obtained for one asset applies, except that the value of $E_t[\beta]$ in this instance depends on the historical information received from observing previous assets. Under the assumption of noninformative priors about the parameter β, after n observations, the posterior distribution will depend on the sample correlation between $\tilde{X}$ and $\tilde{S}$.

Proposition: Let the prior distribution of the parameter β be Jeffreys' improper prior $p_0(\beta) = 1, \forall\beta$, and let $E_0[\beta] = 0$ under the improper prior. The posterior distribution after observing n independent assets has the proper posterior distribution $\beta|X_1,\dots,X_n \sim N(\hat{\beta},\sigma_{\beta}^2)$ where:

$$\hat{\beta} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})S_i}{\sum_{i=1}^{n} S_i^2} \quad \text{and} \quad \frac{1}{\sigma_\beta^2} = \sum_i S_i^2.$$

Proof: Using Bayes theorem and the fact that $\alpha = \bar{X}$, the posterior probability density function is

$$f(\beta|X_1, \dots, X_n) \propto p_0(\beta) \times f(X_1, \dots, X_n|\beta)$$

$$= p_0(\beta) \times \phi(X - \alpha 1_n - \beta S; 0_n, I_n)$$

$$= p_0(\beta) \times \phi(X - \bar{X} 1_n - \beta S; 0_n, I_n),$$

where $X = (X_1, \dots, X_n)'$ and $S = (S_1, \dots, S_n)'$ are vectors corresponding to the values of $\tilde{X}$ and $\tilde{S}$, 1_n and 0_n are *n*-vectors of ones and zeroes, $I_n$is the $n \times n$ identity matrix, and $\phi(. ; 0_n, I_n)$ is the joint normal density function with mean 0_n and covariance matrix I_n.

Focusing on the log-likelihood, we have:

$$\ln f(\beta|X_1, \dots, X_n) = const. - \frac{1}{2}\sum_{i=1}^{n}(X_i - \bar{X} - \beta S_i)^2$$

$$= const. - \frac{1}{2}\,[\sum_{i=1}^{n} S_i^2 \beta^2 - 2\sum_{i=1}^{n}(X_i - \bar{X})S\beta + \sum_{i=1}^{n}(X_i - \bar{X})^2]$$

$$= const.' - \frac{1}{2\sigma_\beta^2}\left(\beta - \hat{\beta}\right)^2,$$

where the last equality follows from observing that the previous expression is a quadratic polynomial in β, and completing the square with an appropriate choice of constant. This log-likelihood corresponds to the normal distribution $N\left(\hat{\beta}, \sigma_\beta^2\right)$.■

To characterize the investor behavior before and after the estimation of the parameter β, we consider the case in which the true parameter in the data generating process is $\beta^* > 0$. Because $\tilde{S} \perp \tilde{\epsilon}$ and $E_t[\tilde{\epsilon}] = 0$, we have:

$$E_n[\hat{\beta}] = E\left[\frac{\sum_{i=1}^{n}(X_i - \bar{X})S_i}{\sum_{i=1}^{n} S_i^2}\right] = E\left[\frac{\sum_{i=1}^{n}(\beta^* S_i + \epsilon_i)S_i}{\sum_{i=1}^{n} S_i^2}\right] = \beta^*.$$

We next consider the expected returns in the short and long run for the first asset (i.e., at $t = s_1$ and $t = l_1$, $R_{s,1}$ and $R_{l,1}$), upon observing the first, unity signal $\widetilde{S_1} = 1$. Under improper priors, $E_0[\beta] = 0$:

$$E[R_{s,1}|\tilde{S}_1 = 1] = E\left(\frac{E_0[\beta] \times \tilde{S}_1}{\bar{X}}|\tilde{S}_1 = 1\right) = \frac{1}{\bar{X}}E_0[\beta] = 0$$

However,

$$E[R_{l,1}|\tilde{S}_1 = 1] = E\left(\frac{\widetilde{X_1} - \bar{X} - E_0[\beta] \times \tilde{S}}{\bar{X} + E_0[\beta] \times \tilde{S}}|\tilde{S}_1 = 1\right)$$

$$= \frac{1}{\bar{X}}E\left(\widetilde{X_1} - \bar{X}|\tilde{S}_1 = 1\right)$$

$$= \frac{1}{\bar{X}}E\left(\bar{X} + \beta \times \widetilde{S_1} + \tilde{\epsilon} - \bar{X}|\tilde{S}_1 = 1\right)$$

$$= \frac{1}{\bar{X}}E\left(\beta \times \widetilde{S_1} + \tilde{\epsilon}|\tilde{S}_1 = 1\right)$$

$$= \frac{1}{\bar{X}}E(\beta)$$

$$= \frac{\beta^*}{\bar{X}} > 0.$$

Finally, we derive the same set of returns for the n^{th} asset. The representative investor now uses own estimate of $\hat{\beta}$ in deriving asset prices, so that $P_{s_n} = \bar{X} + \hat{\beta} \times \tilde{S}_n$. Noting that $E[\hat{\beta}] = \beta^*$, we have:

$$E[R_{s,n+1}|\tilde{S}_{n+1} = 1] = E\left(\frac{\widehat{\beta} \times \tilde{S}_1}{\bar{X}}|\tilde{S}_{n+1} = 1\right) = \frac{\beta^*}{\bar{X}} > 0 \quad \text{and} \quad E[R_{l,n+1}|\tilde{S}_{n+1} = 1] = 0.$$

A.2.3 Bayesian learning and Bayesian shrinkage

The model in the previous section resembles the framework advanced by Jensen, Kelly, and Pedersen (2023, JKP hereafter). We view a JKP-style Bayesian-shrinkage benchmark as a counterpoint for our learning model. Considering only traditional anomaly returns, formed with delayed implementation, makes the two frameworks appear observationally similar. Because both can rationalize a weaker post-publication expected return, attenuation in long-horizon anomaly returns alone does not distinguish between learning and shrinkage. The distinction arises at the disclosure date. In the JKP-style benchmark, the signal proxies for an expected-return component earned over time, so the model does not generate an immediate disclosure-window reaction. In our learning model, publication changes investors' understanding of the signal's mapping into fundamentals, so prices react more immediately when the signal is released.

Our framework differs from JKP's along several material dimensions. In our model, neither the econometrician nor investors know the fundamental value of the signal (β^*), whereas for JKP only the econometrician lacks this knowledge. Moreover, instead of our assumption of diffuse priors, JKP assume priors centered around $\beta = 0$. Lastly, we assume zero risk, while JKP implicitly assume that risk compensation will accrue over the entire period. This last assumption is crucial in characterizing the return dynamics because the JKP assumption precludes any immediate reactions to the signal.

Maintaining the notation from the previous section, we change the prior distribution of β to $\beta \sim N(0, \tau^2)$ and use the sampling distribution of the estimate (after n observations) $\beta \sim N\left(\hat{\beta}, \sigma_{\beta}^2\right)$. The posterior estimate of β, given the information set that includes all n observations Ω_n, is $E[\beta|\Omega_n] = \frac{\tau^2}{\tau^2+\sigma_{\beta}^2}\hat{\beta}$. This estimate shrinks $\hat{\beta}$ towards zero; the extent of the shrinkage is related both to the distribution parameter σ_{β}^2 and to how confident the econometrician is about own priors (i.e., how small the prior variance τ^2 is). From the viewpoint of the econometrician, because the signal proxies for risk exposure and the premium can only be realized over time, the immediate returns will be zero:

$$E\left[R_{s,1}\middle|\tilde{S}_1 = 1\right] = E\left[R_{s,n+1}\middle|\tilde{S}_{n+1} = 1\right] = 0.$$

For long-term returns, the perspective changes slightly in terms of its focus on the ex-post average returns. That is, once the econometrician observed n observations with a sample mean of $\hat{\beta}$, the expected long-run return is:

$$E\left[R_{s,t}\middle|\hat{\beta}, \tilde{S}_t = 1\right] = \frac{\hat{\beta}}{\bar{X}} > 0\,, \forall t \leq n.$$

In other words, after observing n assets, the expected return for one of these assets is simply the sample mean. However, the best estimate for the $n + 1^{st}$asset uses $E[\beta|\Omega_n]$, that is, a zero-centered Bayesian shrinkage of $\hat{\beta}$ in the spirit of JKP:

$$E_n\left[R_{s,n+1}\middle|\, \tilde{S}_{n+1} = 1\right] = \frac{\tau^2}{\tau^2 + \sigma_\beta^2}\frac{\hat{\beta}}{\bar{X}} < \frac{\hat{\beta}}{\bar{X}}.$$

A.3 The rational model

As in the behavioral model, we assume a risk-neutral representative investor with a discount rate of 0. Therefore, the price of the asset at any time t is given by $E_t\left[\widetilde{X_t}\right]$. However, in this case, the representative investor knows the true value of $\beta^* = 0$. Therefore, the rational price of the assets remains at $\bar{X}, \forall t < l_1$. At $t = l_1$, the price will become $P_l = \tilde{X}$.

This model corresponds to the data-mining interpretation discussed in Section 5.1: researchers search over a large set of candidate signals (a high-dimensional set of possible accounting transformations) and, because of publication selection, a spurious signal may appear to earn in-sample returns. Because the signal is truly uninformative about fundamentals, a rational market has no reason to react at the disclosure date—either before or after academic publication—

so the model predicts zero short-horizon (disclosure-window) returns in both regimes. This contrast sharpens the interpretation of the empirical publication effect: nonzero disclosure-date reactions, and especially their post-publication strengthening, are difficult to reconcile with pure data mining.

To model multiple hypothesis testing, we assume that there are $K \gg 1$ signals $\tilde{S}_k$, $k \in \{1,2,\dots,K\}$. Upon observing N assets and the associated prices and signals, researchers (independently) evaluate the statistical significance of the relation between $\tilde{S}_k$ and $\tilde{X}$. The researchers estimate the empirical model:

$$P_{l_n} = \alpha + \beta \times \tilde{S}_{k,n} + \tilde{\epsilon}_{k,n},$$

where $\tilde{S}_{k,n}$ refers to the k^{th} signal for the n^{th} asset. Under the assumptions of normality, the parameter β is distributed according to $\beta \sim N\left(0, \frac{1}{N}\right)$.

The multiple hypothesis testing problem arises when, motivated by publication preferences, only the researcher with the highest estimate of β gets to publish. Under this assumption, the signal corresponding to the article published is the highest order statistic $\beta^{(K)}$ from the set of the estimated parameters β. The probability density function of $\beta^{(K)}$ is (Casella and Berger, 1990):

$$f_K(\beta) = \left(1 - \Phi\left(\frac{\beta}{\sqrt{N}}\right)\right)^{K-1} \phi(\beta/\sqrt{N}),$$

where the functions $\phi \; and \; \Phi$ denote the standard normal probability and cumulative density functions, respectively.

The expected value of $\beta^{(K)}$ may be approximated as (Blom, 1958):

$$E\left[\beta^{(K)}\right] \approx \frac{1}{N}\Phi^{-1}\left(\frac{K-0.375}{K+0.25}\right).$$

As K becomes large, the value $\frac{K-0.375}{K+0.25} \to 1$ and $\Phi^{-1}\left(\frac{K-0.375}{K+0.25}\right) \to \infty$. Therefore, a large set of potential signals yields a large estimate of $\beta^{(K)}$.

Let K be sufficiently large so that the value $E\left[\beta^{(K)}\right] \gg 0$. We characterize returns R_s and R_l both in-sample (i.e., subject to the multiple hypothesis testing) and out-of-sample. To do so, we fix k such that the k^{th} signal $\tilde{S}_k$ corresponds to the highest ordered statistic $\beta^{(K)}$. For this signal, we have:

$$E_1\left[R_{s,1}|\tilde{S}_{k,1} = 1\right] = 0 \quad \text{and} \quad E_1\left[R_{l,1}|\tilde{S}_{k,1} = 1\right] = \frac{E\left[\beta^{(K)}\right]}{\bar{x}} > 0.$$

However, the $(n+1)^{st}$ observation is not subject to the multiple hypothesis testing problem. In fact, $\tilde{X}_{n+1}$ and $\tilde{S}_{k,n+1}$ are independent. Therefore,

$$E_{n+1}\left[R_{s,n+1}|\tilde{S}_{k,n+1} = 1\right] = 0 \quad \text{and} \quad E_{n+1}\left[R_{l,n+1}|\tilde{S}_{k,n+1} = 1\right] = 0.$$

A.4 Rational and behavioral model comparison

The models capture the differences between the learning model and the rational model. Given that researchers have used the long-run returns in their analyses of anomalies, the returns immediately upon the release of information offer an out-of-sample test. Using the samples used by prior research, that is, the long-run returns both pre-and post-academic publication, the two models have the same prediction. Using the expressions in our models, the two models have the same predictions regarding $E_1\left[R_{l,1}\right]$ and $E_{n+1}\left[R_{l,n+1}\right]$. To protect against look-ahead biases, the extant empirical literature has focused on long-run returns and, therefore, only studied the empirical equivalents of $E_1\left[R_{l,1}\right]$ and $E_{n+1}\left[R_{l,n+1}\right]$. At least in the US context, empirical studies show that $E_1\left[R_{l,1}|\tilde{S} > 0\right] > 0$ while $E_{n+1}\left[R_{l,n+1}|\tilde{S} > 0\right]$ is either 0 or significantly close to 0 (compared to $E_1\left[R_{l,1}|\tilde{S} > 0\right]$).

Our empirical tests implement precisely this comparison by shifting attention from traditional, delayed-formation anomaly returns to disclosure-date returns. Because the original anomaly studies did not condition on disclosure dates, by construction, the disclosure-window returns are out-of-sample with respect to those studies. Accordingly, the rational (data-mining) model predicts no disclosure-window reaction both before and after publication, whereas the behavioral learning model predicts that disclosure-window reactions emerge or strengthen after publication—the mechanism underlying the publication effect and timeline compression documented in Table 2 and Figure 1 in the main text.

The returns immediately upon the release of the signals, captured by $E_1\left[R_{s,1}\right]$ and $E_{n+1}\left[R_{s,n+1}\right]$ in both models, have been overlooked in the extant literature. The initial omission of these returns ruled out look-ahead biases in return prediction. However, because they do not suffer from in-sample biases, this omission has rendered these returns useful from an econometrician's perspective. Moreover, the two models differ in their prediction regarding these returns. Under the rational model, these out-of-sample returns should be zero both before and after the academic publication. Any relation captured in the long-run returns is spurious under the rational model. As such, it should have no bearing on out-of-sample tests.

On the other hand, the behavioral model suggests that, initially, investors did not recognize that the signal predicts firm value and, therefore, $E_1\left[R_{s,1}|\tilde{S}_1 > 0\right] = 0$. However, upon learning the true nature of the signal's predictive ability of value, perfect Bayesian updating leads to pronounced price movements around the release of the information: $E_{n+1}\left[R_{s,n+1}|\tilde{S}_{n+1} > 0\right] > 0$.

Table A.1 summarizes these considerations.

Table A.1
Summary of Behavioral and Rational Model Predictions Before and After Publication

In Table A.1, "traditional anomaly portfolio formation timing" corresponds to the conventional long-horizon anomaly strategy implementation with a waiting period (to avoid look-ahead bias), whereas "accelerated anomaly portfolio formation timing" corresponds to the disclosure-date event-return design used in the paper (three-day returns around the information-relevant date).

	$E_p[R_l \mid \tilde{S} > 0]$ (traditional anomaly portfolio formation timing)		$E_p[R_s \mid \tilde{S} > 0]$ (accelerated anomaly portfolio formation timing)	
	Behavioral Model	Rational Model	Behavioral Model	Rational Model
Before Publication ($p = 1$)	+	+	**0**	**0**
After Publication ($p = n + 1$)	**0**	**0**	+	**0**