

Predicting Anomalies

Boone Bowles, Adam V. Reed,

Matthew C. Ringgenberg, and Jacob R. Thornock*

April 2026

ABSTRACT

We show that stock returns exhibit predictable patterns before the publication of anomaly trading signals. Moreover, anomaly trading signals derived from financial data are themselves predictable, making it possible to trade before financial statements are released. A trading strategy based on predicted anomaly signals earns an annualized return of 2.80% in the quarter before the signal is released. In recent periods, this return predictability is concentrated in signals that are harder to forecast, and returns are increasingly earned several quarters before signals are released. Our findings suggest anomalies are more anomalous than previously recognized.

Keywords: Anomalies, Asset Pricing, Information Economics, Return Predictability

JEL Classification Numbers: G12, G14

*Mays Business School, Texas A&M University (boone.bowles@tamu.edu), Kenan-Flagler Business School, University of North Carolina (adam.reed@unc.edu), David Eccles School of Business, University of Utah (matthew.ringgenberg@eccles.utah.edu), Marriott School of Business, Brigham Young University (thornock@byu.edu). The authors thank the Editor, an anonymous referee, Mike Cooper, Harald Hau, Jerry Hoberg, Roni Michaely, Andreas Neuhierl, and conference and seminar participants at the 2nd Annual HKU Summer Finance Workshop, Michigan State University, Jackson Hole Finance Group, and European Winter Finance Summit. All errors are our own. ©2024-2026 Boone Bowles, Adam V. Reed, Matthew C. Ringgenberg, and Jacob R. Thornock.

Asset pricing anomalies are trading signals that predict future abnormal returns.¹ For anomalies based on annual financial data, these returns are largest in the period immediately *after* the release of financial statements and decay quickly thereafter.² Yet, the literature has largely overlooked the period *prior* to the release of anomaly trading signals. In this paper, we show that anomaly returns exhibit predictable patterns in the weeks and months before anomaly signals are publicly released. We also show that anomaly signals are highly predictable, making it possible to trade before the signal is publicly revealed.

Figure 1 displays our first key result. While Bowles et al. (2024) show that anomaly returns are concentrated in the weeks immediately following the release of anomaly signals, we show that returns are predictable *before* information is released. In other words, prices respond to anomaly signals even before the signals are publicly known. But how can the market incorporate information before it has been publicly released? The answer is simple: many anomaly signals are highly predictable.

[Figure 1 about here.]

To study this predictability, we examine several anomaly prediction models, ranging from a sophisticated machine learning model to a simple martingale model. Surprisingly, the simple martingale model of the form $E[x_{t+1}|x_1...x_t] = x_t$ outperforms the more complex models. As an example, consider the asset growth anomaly (Cooper et al. (2008)), which forms long-short portfolios using each firm’s book value of assets; we show that asset growth measured using third-quarter financial statements is a reliable predictor of asset growth in the fourth quarter when annual financial statements are released. This suggests that it is possible to trade the asset growth anomaly (and other anomalies) early, a result that

¹As noted in Brennan and Xia (2001), “An asset pricing anomaly is a statistically significant difference between the realized average returns associated with certain characteristics of securities, or on portfolios of securities formed on the basis of those characteristics, and the returns that are predicted by a particular asset pricing model.”

²Bowles et al. (2024) show that the returns to many anomalies are increasingly concentrated in the period immediately after the public release of new information about anomaly trading signals. Ivkovic and Zekhnini (2024) show that this change in return patterns tends to start once academic research first documents a particular anomaly.

calls into question the very manner in which researchers define and measure asset pricing anomalies.

To analyze whether anomaly returns exhibit predictable patterns prior to the release of anomaly signals, we adopt a perfect-foresight model that assumes that an anomaly signal is perfectly predictable one quarter before its release. We then examine the returns to long-short anomaly portfolios around the release of the anomaly signals.³ However, instead of examining return patterns in the days, weeks, and months after the release of key financial data, we examine returns in the days, weeks, and months *before* the release of key financial data. We find that having perfect foresight and trading one quarter before the anomaly signal is released yields an additional return of 171 basis points for the average anomaly, or 7.10% on an annualized basis. Moreover, having perfect foresight and trading one month ahead (instead of one quarter ahead) earns 120 basis points (14.4% annualized), indicating that the results are strongest in the weeks immediately before the information is released.

Of course, since no trader has a crystal ball to perfectly predict future financial statements, it is not clear whether investors could take advantage of these return patterns. Accordingly, we analyze whether anomaly signals are predictable themselves. We show that they are. To show this, we test four different models to predict anomaly signals and the constituents of anomaly portfolios: a first-order autoregressive model that predicts the level of a stock’s anomaly signal (AR Signals), a first-order autoregressive model that predicts the relative ranking of a stock’s anomaly signal (AR Rankings), a machine learning model that uses a stock’s past anomaly signals and relative rankings to predict anomaly portfolio inclusion (ML), and a martingale model that assumes that the relative ranking of a stock’s anomaly signal next period is equal to its relative ranking this period.

³We use event time tests to examine return patterns around the release of anomaly signals. This requires anomalies that exclusively use information that arrives at infrequent, discrete, and observable points in time. Accordingly, we focus on the 28 anomalies used in Bowles et al. (2024) and which rely on information released in either earnings announcements or 10-K filings. See Section II for details.

The results from all four of these prediction models point to the same conclusion: most of the anomalies we study are highly predictable. For each anomaly and prediction method, we calculate the model’s F1 Score, which is analogous to the percentage of observations predicted correctly.⁴ In general, the two autoregressive models underperform the other two models, though they both still yield F1 Scores well above 60% for most anomalies. In other words, the worst-performing models are still able to accurately predict annual anomaly signals one period ahead more than 60% of the time. However, since autoregressive models impose a linear structure, they seem to perform poorly when predicting which stocks will be in the extreme deciles (i.e., the long and short legs of an anomaly portfolio).

In comparison, the machine learning model and the martingale model are consistently more accurate. Across all 28 anomalies and across our three-decade sample period, the F1 Score for the machine learning model averages 69% while the F1 Score for the martingale model averages 71%, meaning the models make correct predictions approximately 70% of the time. The success of the martingale model in particular is not only surprising but also indicates that anomaly signals are highly persistent. For the average anomaly, most of the stocks with high and low anomaly signals after the third quarter (Q3) still have high and low anomaly signals after the fourth quarter (Q4). Thus, a long-short anomaly portfolio constructed at Q3 will remain largely unchanged when annual information is released at Q4. Specifically, 66% of the annual portfolio’s stocks are holdovers from the portfolio constructed at Q3. In other words, by simply assuming persistence of anomaly signals from Q3 to Q4, a trader could make portfolio assignments one quarter early and be correct approximately two-thirds of the time.

Interestingly, we find that the level of persistence is heterogeneous across anomalies; some accounting numbers and anomaly signals are highly persistent, while others are not. For example, for the asset turnover and the profit margin anomalies, we find that over 85%

⁴The F1 Score is a measure of model accuracy that balances Type 1 and Type 2 errors. See Section A of the Internet Appendix for details on the F1 Score.

of the stocks in the annual portfolio are holdovers from Q3. As a result, our prediction models perform well for these two anomalies—the F1 Scores are around 90%. On the other hand, only 26% of the stocks in the annual earnings surprise anomaly portfolio were also in the portfolio at Q3 and our prediction models perform poorly for this anomaly (e.g., the martingale model has an F1 Score of just 29%). We find that the revenue surprise anomaly is similarly difficult to predict. Of course, it is not surprising that our prediction models have difficulty with these two anomalies since they are, by definition, constructed to be “surprises.”

Given the success of the machine learning model and the martingale model, we further examine their Type 1 (false positive) and Type 2 (false negative) error rates and study their forecast accuracy through time by charting F1 Scores for each anomaly and year in our sample. We find that both models are consistently accurate and have low error rates. Of course, the machine learning model could not have been implemented in the early years of our sample, so one advantage of the martingale model is that any trader could have easily implemented it throughout our entire sample period.

We then examine whether these prediction models can be used to earn abnormal returns even before anomaly signals are publicly released. In short, we find that they can. For the average anomaly, constructing anomaly portfolios three months *before* annual information releases using the simple martingale model results in additional abnormal returns of 70 basis points (or 2.80% on an annualized basis). These are *additional returns* since they are in addition to the returns earned from trading the anomaly after the anomaly signal information becomes public.

While the returns to trading early using the martingale model are economically significant, they are less than half of those from the perfect foresight model, which earns 171 basis points (7.10% annualized) over the same three-month period. Why is there such a large difference? Put differently, if the perfect foresight model earns large returns and our

prediction models are accurate, then why are the returns from trading using the prediction models so much less than the returns from the perfect foresight model?

We investigate this by examining the returns to correct predictions, false positives, and false negatives.⁵ With respect to correct predictions, they earn 127 basis points on average in the three months prior to the release of annual anomaly signals. In contrast, a portfolio of false positive predictions (Type 1 errors) *loses* 60 basis points in the same three months. Perhaps more interestingly, a portfolio of false negative predictions (Type 2 errors) earns an astonishing return of 297 basis points in the three months before annual information releases (12.42% annualized). This return is four times that of the martingale model and nearly double that of the perfect foresight model. This finding suggests that the stocks that were not expected to be in the annual long-short anomaly portfolio, but actually are ex-post, earn incredibly large returns.

This finding is even more striking when we zoom in on the one week before annual information releases and see that the martingale model earns 27 basis points, the perfect foresight model earns 69 basis points, yet the portfolio of false negatives earns 108 basis points (74.82% annualized). Thus, returning to the previously posed question, the reason the martingale model and the machine learning model underperform the perfect foresight model is that incorrect predictions may be rare, but they are incredibly costly.

We also study whether the returns to predicting anomalies and trading early have changed over time. The results are three-fold. First, returns to trading prior to annual information release have diminished in recent years. During the early part of our sample (1990–2006), the perfect foresight model earns 279 basis points for the average anomaly over the three months before annual information releases. In recent years (2007–2023), the perfect foresight model earns only 92 basis points.

⁵False positives are those stocks that were predicted to be in the annual long-short anomaly portfolio and were *not* in the annual anomaly portfolio after the signal was publicly released (i.e., Type 1 errors). False negatives are stocks that were not predicted to be in the annual anomaly portfolio but were in the annual portfolio when the signal was released (i.e., Type 2 errors).

Second, although the martingale and machine learning models are consistently accurate throughout our entire sample period, the returns from trading based on these models have been arbitrated away in recent years. In the early years of our sample, trading three months early using either model earns approximately 200 basis points. In the recent period, the martingale model loses 19 basis points, while the machine learning model loses 7 basis points. These findings suggest that investors are using prediction models to arbitrage away predictable returns in the quarter before annual information is released.

Third, correct predictions are less profitable in recent years while incorrect predictions are even more costly. Indeed, false negative predictions earn a much larger share of returns from trading prior to annual information releases. This finding suggests that investors have arbitrated away the predictable component of anomaly returns, but not the difficult to predict component.

Finally, we examine several mechanisms that could explain why anomaly prediction has not been fully arbitrated away. Theoretically, a number of limits to arbitrage have been proposed, including (i) rational inattention and costly information processing, (ii) career concerns and investment horizon constraints, and (iii) short-selling constraints.

To test the first explanation, we divide anomalies into “easy-to-predict” and “hard-to-predict” groups based on their F1 Score and find that the hard-to-predict anomalies consistently earn the highest returns, even in recent periods. We also show that this pattern is robust to alternate predictability cutoffs and to direct complexity proxies based on the number of accounting inputs and lags used to construct the anomaly. These results support Caplin and Dean (2015)’s rational-inattention story: investors arbitrage signals that are cheap to process and leave costly signals largely untouched.

To examine the career concerns and investment horizon channel, we study whether it is profitable to trade on predictable anomalies more than a quarter before information is released. Although returns to trading three months early have been arbitrated away in recent years, we find there are still predictable returns to trading *six* months early using

a two-quarter- ahead martingale model. In other words, predicting anomalies can still be profitable, but you need to act even earlier in event time. This is consistent with Greenwood and Sammon (2022), who show that index addition and deletion premiums have moved earlier and earlier in event time as a result of arbitrage competition (also see Arnott et al. (2023)). The fact that predictability remains at the six-month horizon supports the idea that traders are unwilling to open positions more than one quarter ahead, perhaps due to career concerns. In other words, because professional traders are often evaluated based on recent performance, they may be unwilling to open trades that will not yield a profit for several quarters (De Long et al. (1990)). As a result, mispricing remains at longer horizons.

Finally, to examine the impact of short-selling frictions, we use institutional ownership, a standard proxy for short-sale constraints (Nagel (2005)). We repeat the event-time analysis after separating stocks with low institutional ownership from those without low institutional ownership. We find that, once we exclude the stocks most likely to face short-sale constraints, the same pre-release return pattern remains. This suggests that short-sale frictions do not appear to be the primary driver of the pre-release drift.

Overall, our paper contributes to several strands of the literature. First, we contribute to the long-running debate on the nature and existence of anomalies. While several papers debate whether anomalies are the result of data mining or unmeasured costs (e.g., Hou et al. (2020) and Chen and Velikov (2023)) or whether anomalies are real and the result of mispricing (e.g., Chen and Zimmermann (2019), Chen and Zimmermann (2021), and Jensen et al. (2021)), our findings provide strong evidence that anomalies actually exist in the data and that they are related to information about anomaly signals. These results shed important light on the economic mechanism underlying many asset pricing anomalies.

Second, our paper relates to the large literature on the response of capital markets to the release of accounting data (starting with Ball and Brown (1968)). In this sense, our work contributes to the broader literature on the private versus social value of information releases (e.g., Hirshleifer (1971)). Our paper adds to this literature by generating a new question:

if traders are investing time and capital to predict anomalies, are they over-investing in information acquisition relative to the social optimum?

Finally, our paper also contributes to the extensive literature on market efficiency. The existence of publicly available information that predicts stock returns has long bothered financial economists because these so-called anomalies are apparent violations of market efficiency. We show that the violations are happening at a *deeper level* because the public information that predicts returns is itself predictable based on previously-released public information. In that sense, our results show that anomalies are more anomalous than previously realized.

Moreover, these findings have important economic implications. The fact that some anomalies are persistent suggests that they may have a larger impact on firm’s discount rates, possibly leading to the misallocation of capital in the real economy (van Binsbergen and Opp (2019), Keloharju et al. (2021), van Binsbergen et al. (2023), Baba-Yara et al. (2024)). Our results provide the first comprehensive evidence on the persistence of a large number of anomalies.

Our results also have a number of practical implications. First, our findings highlight that it is important to understand *when* returns are earned because timing impacts our measurement and economic interpretation of events. Second, we show that the returns to many anomalies have migrated over time and are now more concentrated around important information releases. As a result, the returns to many anomalies are larger than they initially appear when rebalancing using traditional methods.⁶ Third, our results have important implications for factor models. If returns are earned prior to information releases, then existing research may be mismeasuring factor returns. This mismeasurement could have implications for our models of risk, as well as the subsequent interpretations we assign to abnormal returns. Future research should continue to explore these important implications.

⁶See Bowles et al. (2024) and Ivkovic and Zekhnini (2024). See also Asness and Frazzini (2013) who show that the HML factor performs better if the calculation is updated to condition on recent price data.

The rest of the paper proceeds as follows. Section I provides background to motivate our empirical tests. Section II discusses our data and the construction of the anomaly variables. Section III presents our main findings that anomalies are predictable, and this information can be used to generate abnormal returns. Section IV concludes.

I. Background

Our study relates to several lines of research that examine anomalies based on accounting information. First, it links to the stream of papers that examine the existence, magnitude, and explanations of accounting-related anomalies (for surveys of this literature, see Kothari (2001), Lee (2001), Richardson et al. (2010), and Lee et al. (2015), among others). As noted in Lee et al. (2015), explanations for accounting-related anomalies are often tied to their connections with attributes of the firm, its performance, or its information environment. These attributes can include such factors as the cheapness of the stock (e.g., Fama and French (1992), Lakonishok et al. (1994)), the profitability of the stock (e.g., Piotroski (2000)), or the safety of the stock (as measured by leverage or financial distress (Dichev (1998), Campbell et al. (2008))). Indeed, Asness et al. (2019) pull these attributes together to create a quality score for stocks that are safe and profitable, demonstrating that accounting information signals are associated with predictable returns.

Perhaps most related to our paper are those that examine the quality of accounting information and its relation to future earnings or future stock returns. Most accounting numbers involve some amount of judgment, discretion, and estimation, all of which are susceptible to human error or human manipulation. For example, Sloan (1996) documents the differential persistence of accruals versus cash flows—accruals require more estimation and managerial discretion, which is likely related to accruals having lower persistence. Both Sloan (1996) and Xie (2001) show that this feature of accrual-based accounting numbers appears to be unanticipated and even mispriced by investors.

Further, Richardson et al. (2005) show that the less reliable components of accruals (i.e., those with greater propensity for manager discretion or manipulation) are less persistent and that investors fail to appreciate this differential reliability, leading to significant mispricing. Additionally, Fairfield et al. (2003) provide evidence that investors tend to misunderstand the growth embedded in accruals. Overall, the evidence in this line of research suggests that the persistence of certain accounting numbers helps explain mispricing of some accounting information. In a similar spirit, we show that investors seem to act as if they do not anticipate the persistence of a large number of accounting anomalies, leading to predictability in the anomaly portfolio assignment and significant mispricing.

Second, this study relates to a long literature on market reactions to the release of information. Samuelson (1965) famously proved that price changes should be unpredictable if they properly reflect all available information: “If one could be sure that a price will rise, it would have already risen.” A number of papers test this idea and the empirical evidence shows that since at least the 18th century, public equity markets have focused on the arrival of new information.⁷ Indeed, the arrival of new information creates a race to trade on it as quickly as possible. Going all the way back to Fama et al. (1969), researchers have documented that share prices move around the release of information.⁸

However, in the subsequent decades, researchers have documented that share prices continue to update and change in the weeks, months, and even years after information is released. The first accounting-based anomaly identified by researchers is post-earnings announcement drift (PEAD) (e.g., Ball and Brown (1968), Bernard and Thomas (1989), and Bernard and Thomas (1990)). The signature pattern of post-earnings announcement drift is the tendency

⁷Koudijs (2016) examines how shares in English companies traded on the Amsterdam exchange in the 1700s responded to new information via the arrival of ships from England.

⁸Most recently, Hartzmark and Solomon (2024) show that aggregate stock market returns move in a predictable manner due to price pressure from investors trading to capitalize on recently announced dividend payments.

of stock prices to continue moving in the direction of recently released earnings news.⁹ As noted in Richardson et al. (2010), post-earnings announcement drift can be summarized as investors underestimating the implications of current earnings for future earnings. A similar intuition applies in our setting as we show investors appear to underestimate the implications of current anomaly signals for future anomaly signals. Indeed, as discussed in Gabaix (2019), underreaction to news can occur when investors are inattentive to the true autocorrelation of a stochastic time series.

Third, our study links to the literature on the predictability of accounting information using prior information. Early work examines the time-series properties of earnings, considering whether earnings tend to follow a random-walk model and exhibit mean reversion (e.g., Kothari (2001)). Fama and French (2000) use a cross-sectional approach to earnings prediction and present evidence that profitability is mean-reverting. Beyond that evidence, multiple studies demonstrate that a broad set of accounting variables can predict future returns (e.g., Lev and Thiagarajan (1993), Abarbanell and Bushee (1997), Abarbanell and Bushee (1998)). Despite decades of research examining the basic predictability of future earnings and returns, our study offers a simple but important distinction: we rewind the clock to before anomaly signals are released and examine whether anomaly signals themselves can be predicted before they are in turn used to predict returns.

Building on Samuelson (1965), we examine whether accounting information used to form anomalies is “properly anticipated.”¹⁰ If it is, then stock returns should not exhibit predictability in the period prior to the release of accounting information about the anomaly signal.

Our empirical design tests this simple idea using a simple methodology: event studies. While there is a large literature debating the concept of market efficiency, there is also a large

⁹Researchers have since documented other forms of drift or momentum to dividend announcements and stock splits (e.g., Ikenberry and Ramnath (2002)) and related evidence that prices reflect information *before* earnings announcements (e.g., Sammon (2025)).

¹⁰Samuelson (1965) titled his article “Proof That Properly Anticipated Prices Fluctuate Randomly.”

literature critiquing the validity of many empirical designs used to test for market efficiency. But as noted in Fama (1991), “The cleanest evidence on market-efficiency comes from event studies, especially event studies on daily returns. When an information event can be dated precisely and the event has a large effect on prices, the way one abstracts from expected returns to measure abnormal daily returns is a second-order consideration. As a result, event studies can give a clear picture of the speed of adjustment of prices to information.” Our paper uses event studies to test whether prices react even before information is released.

II. Data, Anomaly Selection, & Anomaly Calculations

Our study relies on several key data sources. First, we use daily stock returns from the Center for Research in Security Prices (CRSP). From CRSP, we also collect prices, trading volume, and shares outstanding. Second, we use Compustat to collect firms’ quarterly and annual financial information to calculate anomaly variables. Third, we use the Compustat Snapshot database to identify point-in-time information releases, as discussed below.

A. *Compustat Snapshot*

Key to our study is identifying *when* anomaly-relevant information is first made available to investors. To do this, we rely on Compustat Snapshot (hereafter, “Snapshot”), a database that records when accounting information about a firm was made publicly available.¹¹ As noted in D’Souza et al. (2010), different financial statement line items are released at different times by different companies. Thus, for each line item on every financial statement, Snapshot helps identify when the line item was first publicly released. For example, Snapshot identifies whether the book value of assets was revealed on the earnings announcement date or at the 10-Q filing date for a given firm. Using Snapshot, we are able to determine—at the anomaly

¹¹Point in time data has also been used in several recent studies to evaluate departures of market price from fundamental value (Bartram and Grinblatt (2018), Bartram and Grinblatt (2021)) and to evaluate the timing of anomaly returns (Bowles et al. (2024), Ivkovic and Zekhnini (2024)).

level, the firm level, and the quarter level—the earliest date at which anomaly-relevant information is publicly known. We use Snapshot to identify anomaly signals for all quarterly information releases.

Our sample period begins in January 1990 and runs through 2023. Snapshot data identify information release dates beginning in the mid-1980s. We begin the sample in 1990 to allow for several anomalies that use multiple years of data in their construction.¹²

B. *Anomaly Selection*

Given our focus on anomaly returns before and after the release of financial information, we use a set of anomalies for which we can clearly observe the timing of the release of anomaly signals. Although our starting point is the 97 anomalies examined in McLean and Pontiff (2016), for many of these 97 anomalies the underlying data change constantly, making it difficult to establish a discrete information release date.¹³ Accordingly, as in Bowles et al. (2024), we focus on the subset of anomalies that depend entirely on information that has a precise revelation date.

We refine the list of anomalies in McLean and Pontiff (2016) by requiring each anomaly to be based entirely on information that is publicly revealed in quarterly financial statements (10-Ks or 10-Qs) or related releases (such as press releases around earnings announcement dates). The resulting sample contains 28 accounting-based anomalies. For each anomaly, we use Snapshot to identify the earliest date such that all information necessary to construct the anomaly signal is known. We call this date the *information release date*.¹⁴

¹²For more details on Snapshot see Section B of the Internet Appendix. Also see Bowles et al. (2024).

¹³For example, the earnings-to-price ratio (Basu (1977)) requires two data points for each stock: earnings and price. While annual earnings data are released on clearly identified dates and generally remain unchanged for the year, prices change constantly, making it difficult to define the precise information release date for this anomaly. Thus, we do not include earnings-to-price ratio in our study. In contrast, asset growth (Cooper et al. (2008)) is measured using only book assets, a value that is revealed at clear points in time and does not change frequently, so we include asset growth in our sample.

¹⁴For some composite measures based on multiple pieces of information, we use the latest information release date. For instance, to compute asset turnover, which requires sales, cash, long-term debt, short-term debt, common and preferred book equity, and minority interest, we use the latest information release date

C. Anomaly Calculations

Anomaly variables are generally constructed based on the paper originally citing the anomaly.¹⁵ However, we adjust the calculations for some of the anomalies since most of them are constructed using annual financial data while our study uses quarterly anomaly signals to predict annual anomaly portfolios. Specifically, the construction of anomalies that rely on revenue or earnings or other flow variables has been adjusted to use the year-to-date or end-of-quarter financial information. For example, the Asset Turnover anomaly is traditionally calculated by dividing annual sales by year-end net operating assets. In this study, we calculate Asset Turnover quarterly by dividing year-to-date sales by net operating assets as of the end of Q3. Thus, the Asset Turnover value for a firm after the third quarter financial statements are released is a measure of year-to-date Asset Turnover.

We also use the year-to-date construction for anomalies that are based on changes in stock variables, such as the Asset Growth anomaly. In these cases, prior research using an annual anomaly signal would compare last year’s annual financial statements to this year’s. However, to calculate these anomalies on a quarterly basis, we do not do a year-over-year or quarterly comparison, but instead maintain a year-to-date approach. For example, asset growth measured after the release of third quarter financial statements is calculated by comparing total assets at the end of the third quarter (e.g., Q32021) to total assets as of the end of last year (e.g., Q42020), capturing the year-to-date growth in assets. It is worth mentioning here that the focus of this study is to analyze how well past anomaly signals predict future annual anomaly signals. In our view, using the year-to-date approach to calculate quarterly anomaly signals will provide better predictors. This is because quarterly

for those variables as the information release date for the anomaly. If sales is revealed on Monday in an earnings announcement but the other variables are not revealed until Wednesday in the 10-K filing, the anomaly variable cannot be calculated until Wednesday. Thus, for this example, Wednesday would be the information release date. For further details on these anomalies and their construction in the original papers, see Section C of the Internet Appendix.

¹⁵We also verify the construction of these 28 anomalies with the code provided by Andrew Chen and Tom Zimmermann on www.OpenSourceAssetPricing.com and discussed in Chen and Zimmermann (2022).

year-to-date predictors will take into account all available information since the prior year end to forecast the coming year’s anomaly signal. That is, an investor wishing to predict a certain anomaly signal in next year’s annual report would likely use all the information available to them at the time of the prediction—hence the intuition for using the quarterly and year-to-date information in predicting future signals.¹⁶

Overall, we construct and analyze a panel of observations for each anomaly, firm, and year. For a given firm-year, we measure the values of the annual anomaly signal, the quarterly anomaly signal for Q3 and Q2, and the annual anomaly signal from the prior year, $Q4_{t-1}$. Each observation also records several important dates, especially the information release date. We extend the panel by incorporating daily returns and, for consistency, limiting the sample to firms with December fiscal year ends. In all of our tests, we follow the standard methodology of constructing anomaly portfolios to include the stocks in the top and bottom deciles for each anomaly. Summary statistics for our sample are detailed in Table I.

[Table 1 about here.]

III. Return Patterns and Signal Predictability

This section progresses in three stages as we detail our analyses and key findings. First, we describe abnormal return patterns to anomalies in the months before annual information releases. Next, we examine the predictability and persistence of anomaly signals and anomaly portfolios. Finally, we study whether models to predict anomaly signals can be used to generate abnormal anomaly returns in the months before the release of annual information.

¹⁶We have also performed much of the analysis in this study using an alternative year-over-year approach to calculating quarterly anomaly values. This approach resulted in lower predictive power; however, the main inferences of the paper remained unchanged.

A. *Anomaly Returns Before Information Release Dates*

A fundamental tenet of the efficient markets hypothesis (EMH) is that prices should react quickly to the release of new and relevant information and that prices should not continue moving in a predictable direction after the information is released.¹⁷ We propose a separate insight related to the EMH—prices should not move in a predictable direction *before* the release of new and relevant information. We test this idea by examining return patterns immediately before information release dates.

To do this, we use an event-time approach that tracks stock returns in the 60 trading days (three months) before annual information release dates (e.g., Bowles et al. (2024)). This technique mimics a hypothetical investor who is endowed with perfect foresight about future anomaly signals. The investor then forms anomaly portfolios based on annual financial information one quarter before that information is publicly released. Based on that intuition, we compute perfect-foresight returns for a portfolio of all 28 accounting-based anomalies in our sample, which we label as the Average anomaly portfolio. The results of trading anomalies with perfect foresight are shown in Figure 1 and Table II.

[Table 2 about here.]

Figure 1 shows the return profile for the Average anomaly from three months before to three months after the information release date. Most relevant for this study is the return profile *before* the information release date—the figure shows significant returns to trading early with perfect foresight. Table II provides returns to the event-time, perfect-foresight strategy over various periods. It shows that trading three months before information release dates yields a return of 171 basis points (7.10% on an annualized basis). This is slightly less than the return earned after information release dates (178 basis points or 7.70% annualized). In other words, almost 50% of the total six-month return is earned before the information release date.

¹⁷This latter idea is at the center of the post-earnings-announcement drift anomaly.

Not only are the returns to perfect foresight and trading early high, but they are also highly concentrated in the weeks nearest the information release date. Table II shows that having perfect foresight earns 2.8 basis points per day over the three months before the information release dates, 5.7 basis points per day over the one month before, and 11.5 basis points per day over the one week before. These results are striking as they challenge the core event-study intuition that has prevailed in the literature since at least Fama et al. (1969): that the time of the information release is the appropriate time at which stock prices should begin to change. Instead, these findings show returns move predictably prior to information releases, and this predictability strengthens as information releases near. That is, Figure 1 and Table II present evidence of a *pre*-information announcement drift, suggesting that abnormal returns to anomaly signals begin even before the anomaly signal is publicly known.

This raises two important questions. First, to what extent are anomaly signals predictable? That is, can market participants incorporate information about anomalies before the anomaly signal is publicly released? Second, is it possible to implement a prediction model and capture some of the returns earned before information release dates? We examine these questions in the next section.

B. Anomaly Signal Persistence

Table III shows the temporal correlation of anomaly signals for the 28 anomalies in our sample and for the Average anomaly. The table shows the pairwise correlations between the annual anomaly signal ($Q4_t$) and the anomaly signals calculated one quarter earlier ($Q3_t$), two quarters earlier ($Q2_t$), and one year earlier ($Q4_{t-1}$). In addition, we provide serial correlations for anomaly decile rankings. These analyses provide a simple examination of the extent to which anomaly signals and anomaly portfolio assignments are predictable.

[Table 3 about here.]

In general, Table III shows that anomaly signals exhibit positive serial correlation. The signals of the Average anomaly at ($Q4_t$) are correlated at 0.41 with the signal for the Average

anomaly calculated one quarter earlier ($Q3_t$). However, the correlations weaken as the time between anomaly signals increases. The correlation of annual anomaly signals with $Q2_t$ anomaly signals is 0.31 while the correlation with last year’s annual anomaly signal ($Q4_{t-1}$) is 0.05.

There is also significant variation across anomalies in anomaly signal serial correlations. The Asset Growth (Ag) anomaly, for example, has a correlation of 0.94 between the annual signal ($Q4_t$) and the signal at $Q3_t$. Other anomalies also have high $Q4_t - Q3_t$ correlations, such as Inventory Growth (Ig), Changes in Net Working Capital (Nwc), Profit Margin (Pm), and Sales Growth (Sg). On the other hand, some anomalies exhibit relatively low $Q4_t - Q3_t$ serial correlation, including Asset Turnover (At), Net Operating Assets (Noa), Non-current Operating Assets (Nca), Percent Operating Accruals (Poa), Return on Equity (Roe), and Taxes (Tx). Overall, while anomaly signals are serially correlated on average, there is variation in the strength of that correlation across anomalies.

The serial correlations improve for the *relative ranking* of anomaly signals, which is necessary for portfolio construction (i.e., the typical long-short portfolio is based on relative rankings into deciles). The rightmost columns of Table III show that the temporal correlations between anomaly signal rankings are much higher than the correlations between raw anomaly signals. The Average anomaly has a correlation coefficient of 0.81 between the ranking in $Q3_t$ and the ranking in $Q4_t$. Even the correlations between $Q2_t$ and $Q4_t$ are high at 0.71 for the rankings of the Average anomaly. And, with respect to the variation across anomalies, it is striking that 20 out of the 28 anomalies we study have correlation coefficients between Q3 and Q4 anomaly signal rankings that are at least 0.75. These high correlations support a simple, but important insight—the constituents of anomaly portfolios should be highly persistent.

We present evidence for this insight by examining the persistence of portfolio assignments. Table IV shows the percent of stocks in a given portfolio as of $Q4_t$ that were in the same portfolio in prior periods—one quarter earlier ($Q3_t$), two quarters earlier ($Q2_t$), and one year

earlier ($Q4_{t-1}$). We find that 66% of the stocks in the Average anomaly portfolio as of $Q4_t$ were also in the portfolio at $Q3_t$. In other words, two-thirds of the annual Average anomaly portfolio persists from $Q3_t$. That figure drops to 56% when considering stocks that remained in the anomaly portfolios from $Q2_t$ to $Q4_t$, and further to 32% for stocks remaining in the Average anomaly portfolio from last year ($Q4_{t-1}$). Table IV also shows that some anomaly rankings (e.g., Asset Turnover (At) and Profitability (Pro)) are highly persistent while others (e.g., Earnings Surprise (Es) and Revenue Surprise (Rs)) are not. Importantly, the fact that some anomaly portfolio assignments are persistent suggests they may have an impact on firm’s discount rates (van Binsbergen and Opp (2019), Keloharju et al. (2021), Baba-Yara et al. (2024)). As such, these findings add novel evidence on the potential relation between asset pricing anomalies and real economic outcomes.

[Table 4 about here.]

C. Four Models to Predict Anomalies

Given the high persistence of anomalies and anomaly portfolio assignments documented above, it follows that anomaly signals and the constituents of anomaly portfolios should be predictable. We explore this idea using four different prediction models and we compare the results to a simple benchmark that uses the historical persistence of each anomaly as a base-case. Below, we discuss each of these models, as well as the benchmark model, in detail.

Benchmark: We construct an anomaly-specific benchmark that uses the $Q3_t$ -to- $Q4_t$ persistence of each anomaly from Table IV. The benchmark predicts that a stock remains in its $Q3_t$ classification in $Q4_t$ with probability equal to the anomaly’s estimated persistence rate. For example, the Accruals anomaly has $Q3_t$ -to- $Q4_t$ persistence of 63%, so a stock in the long (short) leg at $Q3_t$ is predicted to remain in that leg at $Q4_t$ with probability 0.63.¹⁸ This

¹⁸Stocks not in a leg of the Accruals anomaly portfolio at Q3 are assigned to enter with probability $(1 - 0.63)/9$, ensuring that the expected size of the leg at Q4 remains 10%.

benchmark captures the mechanical predictability implied by anomaly portfolio persistence alone and provides a clear baseline for evaluating the incremental contribution of the four prediction models.

Autoregressive models: Two of the prediction models are autoregressive models, in which the current value of an anomaly signal is based on its preceding values, plus a stochastic term. These models represent the canonical random walk with drift. The first autoregressive model we test predicts the *level* of the anomaly signal given year-to-date information on that particular signal. This model is labeled the “AR Signals” model. The second autoregressive model predicts the *relative ranking* of anomaly signals using year-to-date information on the anomaly ranking (anomaly signals are ranked into deciles). We label this model “AR Rankings.” Both models are of the form:

$$x_{4,t} = \delta + \phi_1 x_{3,t} + \phi_2 x_{4,t-1} + w_{4,t}, \quad (1)$$

where $x_{4,t}$ is either the anomaly signal or the anomaly ranking at Q4_t; $x_{3,t}$ is the signal or ranking one quarter earlier at Q3_t; and $x_{4,t-1}$ is the signal or ranking four quarters earlier, Q4_{t-1}. The parameter of interest, ϕ_1 , measures the persistence in the autoregressive process, indicating the degree to which the anomaly signals or rankings at Q3_t influence the Q4_t value. The random error term is denoted as w .

It is important to note that we use a rolling-window strategy with all of our prediction models to prevent a look-ahead bias. Specifically, we use data up to year t to train the autoregressive models and then make a prediction about year $t + 1$. Then, using the predictions, we classify whether or not a stock will be in the long or short leg of the anomaly portfolio at Q4_t.

Machine Learning model: We also study a machine learning model that uses past anomaly signals, anomaly rankings, and other firm-specific characteristics to predict whether a stock will be in the long or short leg of a given annual anomaly portfolio. The features considered in the model include the following: anomaly rank at Q3_t, anomaly rank at Q2_t, the interaction

between the anomaly ranks at $Q2_t$ and $Q3_t$, the squares of the anomaly ranks at $Q3_t$ and $Q2_t$, the log of total assets, the interaction between the anomaly rank at $Q3_t$ and the log of total assets, the difference between the anomaly rank at $Q3_t$ and $Q2_t$, and the industry of the firm. We choose this feature set *ex ante* based on features that intuitively relate to anomaly predictability. Specifically, we include current and lagged anomaly ranks to capture anomaly persistence, interactions and squared terms to capture simple nonlinearities, firm size and interactions to capture scale effects, and industry fixed effects to capture industry heterogeneity. For comparability, we also keep the feature set constant across anomalies and years, rather than searching over alternate possible specifications.¹⁹

The specific machine learning algorithm we use is a high-performance gradient-boosting machine called the LightGBM Classifier. Similar to other gradient-boosting machines, it is a tree-based algorithm that iteratively builds models while minimizing a loss function. LightGBM is ideal for our tests as it is tailored for classification tasks (such as classifying a stock as in or out of an anomaly portfolio) and efficiently handles large datasets. To keep the model simple and comparable across anomalies, we do not use cross-validation to search over feature sets or to tune the model aggressively.²⁰ Aside from class weights used to address class imbalance, we leave the LightGBM hyperparameters at their default values so that they are consistent across each of the anomalies we consider.²¹

Martingale model: Motivated by the high persistence of anomaly portfolios from $Q3_t$ to $Q4_t$, we use a martingale model that simply assumes that the annual anomaly portfolio ranking will be the same as the ranking as of $Q3_t$. That is, the conditional expected value of the anomaly signal ranking at $Q4_t$ is best approximated with the $Q3_t$ ranking. It is effectively

¹⁹We also keep this feature set fixed across anomalies and years in a rolling expanding-window design.

²⁰We did test one other machine learning algorithm, XGBoost, which made nearly identical predictions but with less efficiency.

²¹We give a weight of 1 to stocks not inside the legs of anomaly portfolios (class=0) and a weight of 2 to stocks inside the legs of anomaly portfolios (class=1). In Internet Appendix D, we examine additional class weights for the LightGBM model. The results are qualitatively similar.

a “random walk without drift,” as the expectation of the error term is assumed to be zero. This model is of the form:

$$E[x_{4,t}|x_{1,t}\dots x_{3,t}] = x_{3,t}. \quad (2)$$

To assess the quality of each of these prediction models, we examine their F1 Scores. The F1 Score measures a model’s accuracy while considering both the precision and recall of the predictions, providing a single metric for prediction quality.²² The F1 Score ranges between 0 and 1, with higher values indicating a higher degree of precision and recall, and thus a better prediction model. For each anomaly and model, we compute separate F1 Scores for long-leg membership and short-leg membership (each a binary classification), and we report the average of the two.

Table V reports F1 Scores for the four prediction models and the persistence-based benchmark for each anomaly. The benchmark already achieves meaningful predictive performance (Average F1 Score of 0.50), reflecting the strong predictability implied by $Q3_t$ -to- $Q4_t$ persistence. All four prediction models improve upon this baseline on average, while the martingale model—the simplest model—performs best overall. For the Average anomaly its F1 Score is 0.71, an increase of 0.21 (or 42%) relative to the benchmark. The machine learning and AR Rankings models follow closely with average F1 Scores of 0.69 and 0.68, respectively. The differences across the top-performing models are modest for most anomalies, with the martingale model typically best or tied with the machine learning model. This suggests that even with a sophisticated algorithm and additional inputs, the machine learning model relies heavily on the anomaly rankings in $Q3_t$ to predict portfolio membership at $Q4_t$. In other words, the best model may be the simple martingale model.²³

²²Precision is the proportion of true positive predictions among all positive predictions. Recall is the proportion of true positive predictions among all actual positive cases. See Section A of the Internet Appendix for a formal discussion of prediction quality measures.

²³For visual clarity, Figure 2 is a heatmap of the F1 Scores for the four prediction models discussed in this section. Also, while the machine learning and martingale models are highly accurate on average, they do perform poorly for three anomalies: Earnings Surprise, Revenue Surprise, and the Tax anomaly. This is

[Table 5 about here.]

[Figure 2 about here.]

These models’ comparative performance can also be seen in Table VI and Figure 3 where the precision (Type 1 error rate) and recall (Type 2 error rate) are plotted. As was the case when comparing F1 Scores, the martingale and machine learning models dominate the autoregressive models. And again, the martingale and machine learning models are very similar.²⁴

[Figure 3 about here.]

[Table 6 about here.]

Finally, we explore whether the quality of these models has changed over time. As shown in Figure 4 the quality of the martingale model has been relatively constant over time. In other words, it consistently performs well at forecasting future anomaly signals and the constituents of anomaly portfolios throughout our sample period.²⁵

[Figure 4 about here.]

Overall, the results suggest that the returns to trading early with perfect foresight may be achievable by investors. We investigate this idea in the next section.

not surprising, however, since two of these three anomalies are, by definition, constructed to be “surprises” and the third, Taxes, is known to vary significantly from year to year while being less informative on a year-to-date basis.

²⁴See Section D of the Internet Appendix for robustness to alternative weighting choices in model evaluation and machine-learning training (including value-weighted F1 scores, value-weighted LightGBM training, and a grid over LightGBM class weights).

²⁵The results in Figure A2 of the Internet Appendix show a similar result for the machine learning model.

D. Trading on Predicted Anomaly Signals

Can investors profitably trade before annual anomaly signals are released? In short, we find that the answer is “yes.” Figure 5 shows returns to the Average anomaly portfolio before and after the annual information release date. The return to the perfect foresight model is shown again in blue (as in Figure 1), but now the return to the Average anomaly portfolio *predicted by the martingale model* is also shown in green.

Two facts are immediately clear. First, trading using the martingale model is profitable. As detailed in Table VII, trading three months early using the martingale model yields 70 basis points (2.80% on an annualized basis). Further, a strategy that forms portfolios three months early using the martingale model then rebalances after the annual information release date could earn 248 basis points (5.02% annualized) in the six-month window around the release date.²⁶ Thus, almost one-third of an achievable six-month around-release return is earned before the information release date.

[Figure 5 about here.]

[Table 7 about here.]

The second fact: though positive and significant, the return to trading early using the martingale model is less than half of the return possible with perfect foresight. This is perhaps surprising, especially given that anomalies are highly persistent and predictable.

D.1. Errors in Anomaly Prediction

To investigate this, we decompose the return earned by the martingale model based on whether the prediction was correct or incorrect. We further decompose incorrect predictions

²⁶The estimate of 248 basis points adds together the 70 basis points earned before the information release date and the 178 basis points earned afterward (see Table II). Also note that Figure 5 and Table VII are replicated using the machine learning model in addition to the martingale model. These results are shown in the Section E of the Internet Appendix. Given the similarities between the machine learning and martingale models with respect to performance, the main text typically addresses only the martingale model while placing results for the machine learning model in the Internet Appendix.

by the type of error. Specifically, we examine event-time return patterns for each of the following scenarios:

- correct predictions – cases in which a stock is correctly assigned to a long or short leg of an anomaly portfolio;
- false positive predictions – cases in which the model incorrectly predicts that a stock should be included in a particular portfolio, but in reality, it should not be included;
- false negative predictions – cases in which the model incorrectly predicts that a stock should not be included in a particular portfolio, but in reality, it should be included.

Figure 6 shows the return profiles of portfolios of correct, false positive, and false negative predictions from the martingale model. The figure makes it clear that incorrect predictions are very costly. This provides an explanation for why the return to trading early using the martingale model fails to achieve even half of the return earned using perfect foresight.

[Figure 6 about here.]

As shown in Table VIII (which adds numerical detail to Figure 6), false positive predictions result in losses of up to 60 basis points (2.42% annualized) in the three months before annual information releases. Nearer the information release date, one week before for instance, false positives can lose up to 33 basis points (18.7% annualized). And this trend continues after the information release date. Overall, making false positive predictions before annual information releases is very costly.

[Table 8 about here.]

The results with respect to false negative predictions are even more striking and showcase, again, the costs of incorrect predictions. False negatives earn very large returns, even up to 297 basis points (12.4% annualized), in the three months before annual information releases. Since they earn 297 basis points before the information release date and 198 basis points

after it, about 60% of the six-month around-release return arrives before the signal is publicly released. And again, the returns are highly concentrated around the annual information releases. In just the one week before annual information releases, these stocks earn 108 basis points (74.8% annualized). Given that these returns are “missed”, in a sense, the costs they impose on an investor using the martingale model are large.

Taken together, while correct predictions are profitable, even approaching the performance of the perfect foresight model as can be seen in Table VIII, earning negative returns due to false positives and missing out on the large returns from false negative predictions results in an achievable return that is much lower than the perfect foresight model. This is true even with the high-quality prediction models.

Importantly, it should also be noted that the false negatives earn exceptionally high returns in the week immediately before the information release date and have the largest returns on the information release date and in the days just after (Table VIII). This suggests that the high concentration of anomaly returns around information release dates (both before and after) is driven by *new* and *surprising* information.

D.2. Time Trends in Returns

Having shown predictable patterns in anomaly returns before annual information release dates, and that some of these returns can be earned using a simple prediction model, we next examine whether these predictable patterns have survived in recent years or whether they only existed in the earlier part of our sample. Our analysis is motivated by the findings in McLean and Pontiff (2016) indicating that anomaly returns are weaker after anomalies are discovered and published, and by Bowles et al. (2024) showing that in recent years anomaly returns are smaller and more concentrated in the weeks after annual information releases.²⁷ For this exercise, we divide our sample in half to compare the early period (1990–2006) with

²⁷See also Ivkovic and Zekhnini (2024).

the recent period (2007–2023) and perform the same analysis as in the prior section. Figure 7 shows our first result.

[Figure 7 about here.]

In the early part of our sample (1990–2006), returns to trading early using the martingale model were very high, even approaching the returns of the perfect foresight model. As can be seen in Figure 7, and also in Table IX, using the martingale model and trading three months early generated 198 basis points for the Average anomaly. Combining the returns over those three months with the return to holding the actual annual anomaly portfolios for three months after the information release date would result in a total return of approximately 400 basis points over those six months. In other words, in the 1990s and early 2000s the strategy of trading early on predictable anomalies earned a very large return.

[Table 9 about here.]

Two things have changed in the recent period. First, the perfect foresight model is almost one-third of the size as in the early period. Thus, even with perfect foresight, returns before information release dates are smaller. Second, there is no return to trading early using the martingale model.²⁸ This can be seen as the relatively flat line before trading-day zero in the second panel of Figure 7. One interpretation of this is that the returns to trading early using a simple prediction model (or even a more sophisticated machine learning model) have been arbitrated away. Additional evidence for this interpretation can be seen in Figure 8, which shows the returns to correct, false positive, and false negative predictions in the two different time periods (the results are also shown in Panel B of Table IX.)

[Figure 8 about here.]

In the early period, correct predictions earned very large returns in the weeks before information releases, up to 260 basis points over three months. False negative predictions

²⁸Section E of the Internet Appendix shows that the same holds true for the machine learning model.

did even better: 330 basis points over three months. Most interestingly, even false positive predictions yield returns of 58 basis points in this early period. All together, in the early period correct predictions were profitable and false positive predictions were not costly. It was still the case, however, that missing returns from the false negative predictions was costly.

In the recent period, the returns are far different as correct predictions earned low returns and, in fact, almost no returns until the week before the annual information release date. This pattern suggests that these predictable returns being arbitrated away in recent years and that making prediction errors became increasingly costly. False positives from the martingale model lost 138 basis points in the three months before information dates. Together, earning much less for correct predictions and losing much more for false positive predictions results in the zero returns to trading early in recent years.

Perhaps the most striking aspect of these trends relates to the false negatives. In recent years, the false negative predictions still earn high returns: these errors from the martingale model earned 268 basis points in the three months before information dates. Focusing on just the one week before the information dates, the false negatives earn 124 basis points. (On an annualized basis, this is 90%!) It is also interesting that the return after information release dates is highest for these stocks. The evidence in Figure 8 suggests that in the two weeks around information release dates, these stocks earn approximately 200 basis points (67% on an annualized basis). What does this pattern suggest?

For one, it suggests that the market is more efficient in recent years. While there was a time when investors could use the simple martingale model to predict annual anomaly portfolios and earn high returns in the three months before annual information release dates, that time has passed. Those returns have been arbitrated away as signal processing costs have diminished and the publicity of anomalies has increased.²⁹

²⁹See McLean and Pontiff (2016) and Bowles et al. (2024).

Second, and perhaps more importantly, it suggests that anomaly returns still exist in recent years, but they are increasingly related to the release of new information. Returns to anomalies are not only highly concentrated around information release dates, but they are also highly concentrated in the stocks where the information is actually *new*. Again, Figure 8 shows that the false negative predictions earn approximately 200 basis points in the two weeks around information release dates. Correct predictions (e.g., stocks that are unsurprisingly in the annual anomaly portfolios) earn almost nothing in these same two weeks. This is perhaps the strongest evidence in our paper that anomaly returns are driven by the release of new information.³⁰

Given that returns to using the prediction models to trade three months (one quarter) early have been arbitrated away in recent years, we next test whether predictions could have been made, and whether returns could have been earned, by trading six months (two quarters) early. Figure 9 shows the returns to a strategy that uses the martingale model as of the second quarter ($Q2_t$) and trades six months before annual information release dates. The results suggest that even in recent years, investors could have earned a positive return with this strategy. Indeed, the returns to correct and false positive predictions are both positive in the period from six months before to three months before information releases. This suggests that while much of the returns to making predictions and trading three months early have been arbitrated away, investors can potentially move earlier in event time and earn abnormal returns by making predictions and trading six months ahead of time. These findings are consistent with Greenwood and Sammon (2022), who show that index addition and deletion premiums are moving earlier and earlier in event time as a result of competition.

[Figure 9 about here.]

³⁰In related analysis, we investigate abnormal trading volume for these anomalies both before and after annual information release dates and during both the early and recent periods in our sample. The main result from this analysis is that there is much more trading volume on information release dates in recent years. A larger discussion of this analysis is included in Section F of the Internet Appendix.

D.3. Mechanisms: Predictability, Horizon, and Short-Selling Constraints

The previous results motivate us to look at the economic mechanisms that explain why anomaly prediction, though diminished, has not been fully arbitrated away. Several mechanisms could be responsible: (i) rational inattention and costly information processing, (ii) career concerns and investment horizon constraints, and (iii) short-selling constraints. In this section, we examine three tests to investigate these mechanisms.

Rational Inattention and Costly Information Processing: First, we test whether surviving pre-release returns concentrate in signals that are intrinsically harder to predict. To do this, we compute each anomaly’s F1 Score from the martingale model and split the sample of anomalies into an “easy-to-predict” group (F1 Scores ≥ 0.70) and a “hard-to-predict” group (F1 Scores < 0.70). Figure 10 shows the distribution of F1 Scores across the 28 anomalies in our sample—the mean F1 Score is 0.70—and the resulting split of 10 hard-to-predict and 18 easy-to-predict anomalies. This design targets the first potential mechanism, rational inattention and costly information acquisition: if investors arbitrage the “cheap-to-process” signals first, then any pre-release drift that remains should be concentrated in the hard-to-predict anomalies (Caplin and Dean (2015)).

[Figure 10 about here.]

We then redo the event-time tests by predictor quality (Hard-to-predict versus Easy-to-predict) and by sample period (Early versus Recent years), focusing on the three-month horizon before information releases. The results are shown for both the perfect foresight and martingale models in Figure 11, with corresponding results reported in Table X. In the early period, both the hard-to-predict and easy-to-predict groups earn positive and economically meaningful pre-release returns under both models, with the hard-to-predict group drifting significantly more than the easy-to-predict group. For example, Panel A of Table X shows perfect foresight returns of 368 basis points for the hard-to-predict anomalies compared to 233 basis points for the easy-to-predict group during the early period. In the recent period,

however, pre-release returns have been arbitrated away for the martingale model in both groups, while under perfect foresight the pre-release drift survives only for the hard-to-predict group and is essentially negligible for the easy anomalies. In other words, what was broad-based in the early period is now concentrated in the subset of anomalies that are intrinsically harder to forecast.

[Table 10 about here.]

[Figure 11 about here.]

Taken together, these results are consistent with a rational inattention and costly information processing channel: investors appear to have competed away the pre-release returns related to “cheap” signals first. This interpretation also reconciles the time-trend evidence—why perfect foresight drift persists even as martingale returns vanish—and motivates the next set of tests on when in event time the remaining gains now accrue.

Our interpretation is also robust to alternative ways of classifying anomalies as easy or hard to predict. Section G of the Internet Appendix shows that the same patterns obtain when we define easy- and hard-to-predict anomalies using alternative cutoffs of the F1 distribution, including median splits and tercile-based splits. That section also presents anomaly-level scatter plots relating predictability to pre-release returns, as well as similar tests that sort anomalies by the complexity of the underlying accounting signal based on the number of required inputs and lags. Across these alternative definitions and exercises, the same broad conclusion emerges: pre-release returns are strongest where anomaly signals are harder to anticipate, which further supports the rational inattention and costly information processing interpretation.

Career Concerns and Investment Horizon: Next, we test the career concerns and investment horizon channel with a straightforward idea: if one-quarter-ahead opportunities have been competed away, returns may have migrated earlier in event time. If investors have career concerns (De Long et al. (1990)) and are evaluated on relatively short time horizons,

then they may be reluctant to initiate trades that will not be fully profitable until several quarters ahead. To test this possibility, we extend the event window back one additional quarter and re-estimate event-time returns. This design allows us to examine whether return drift is simply earlier (consistent with horizon constraints) or instead reflects something else about the anomaly signals themselves. Figure 12 illustrates the results: consistent with horizon constraints, we see that return predictability appears earlier in event-time.

[Figure 12 about here.]

In the recent period, perfect foresight earns over 5% for the hard-to-predict anomalies and about 2% for the easy-to-predict group over the 6 months pre-release. Strikingly, the martingale model also earns economically meaningful pre-release returns when looking earlier in event time: approximately 2% for both the easy- and hard-to-predict groups. Notably, this return drift for the martingale model exists entirely in the -120 to -60 trading day window, not in the last 60 trading days before the information release date. In other words, implementable returns exist when the trade is opened two quarters ahead, but they largely do not exist in the final quarter where most investors would prefer to concentrate activity.

These patterns are consistent with the career concerns and investment horizon channel: as competition has intensified in the weeks leading up to the information release date, returns have moved earlier in event time, to the -120 to -60 window (mirroring the “disappearing index effect” in Greenwood and Sammon (2022)).³¹

Short-selling Constraints: Finally, we examine whether short-selling constraints drive the pre-release return pattern. To do this, we use institutional ownership (IO) as a standard proxy for short-sale constraints (Nagel (2005)), and repeat the event-time analysis after separating stocks into those with low IO (more likely to be short-sale constrained) and

³¹Section G of the Internet Appendix shows that this result is robust to alternate definitions of easy- and hard-to-predict.

those that do not have low IO.³² Figure 13 compares long-short hedge returns for stocks in the lowest-IO decile with returns for the remaining stocks that do not have low IO. Once we exclude the stocks most likely to face short-sale constraints, the main return pattern remains. That is, stocks with medium or high institutional ownership exhibit the same pre-release drift that appears throughout the paper in both the perfect foresight and martingale specifications. If short-sale constraints were the main driver of our findings, this pattern should largely disappear once the most constrained stocks are removed. It does not.

[Figure 13 about here.]

These results indicate that short-sale constraints are not the primary explanation for our central findings. At most, short-sale frictions may amplify return magnitudes for the most constrained stocks, but they do not generate the pre-release drift in the first place. Combined with the predictability and horizon tests above, the evidence points to costly information processing and investor horizon constraints as the more plausible mechanisms behind the timing of anomaly returns.

IV. Conclusion

There is a large literature on asset pricing anomalies, yet by tradition, many papers form anomaly portfolios annually, on June 30th, to ensure that trading signals have been publicly available long enough to prevent a look-ahead bias (e.g., Fama and French (1992)). More recent studies have shown that this leads to the use of stale information, and they instead advocate for forming anomaly portfolios on information release dates (Bowles et al. (2024)) or around news events (e.g., Engelberg et al. (2018)). The economic motivation is intuitive: the returns to anomaly signals should begin to accrue as soon as anomaly signals are publicly revealed.

³²Institutional ownership is commonly used as a proxy for lending supply because institutional investors are the primary suppliers of lendable shares, and stocks with low institutional ownership tend to exhibit lower lending supply and higher borrowing costs (Nagel (2005); Beneish et al. (2015)).

We examine a simple but important follow-up question: Can an investor earn anomaly returns *before* anomaly signals are even released? We first show that many anomaly signals are highly predictable. By examining various models, including a sophisticated machine learning algorithm and simple autoregressive models, we find that a straightforward martingale model often provides the best predictive ability for portfolio assignment. In other words, a significant portion of the stocks in the long and short legs of anomaly portfolios remain in the portfolios from quarter to quarter. Our results show that investors (and researchers) do not need to wait until information is released to form portfolios based on anomaly signals.

If anomaly signals are highly predictable *ex ante*, an investor who moves faster than others should be able to outperform the competition. Our evidence supports this idea. Specifically, we show robust evidence that trading on predictable anomalies is highly profitable if done before the information release date. The evidence suggests that returns begin to react to anomaly signals even before they are publicly released. On average, an investor who trades on predictable anomalies using the previous quarter’s information generates an additional return of approximately 3% on an annualized basis.

Overall, our study brings a simple intuition—that anomaly signals are predictable—to a mature and well-developed asset pricing literature. An investor with an understanding of the predictability of anomaly signals can significantly outperform others by simply trading on predictable signals. Moreover, the investor does not need a highly complex forecasting model, as even fairly simple models are highly predictive of future anomaly signals.

The predictability of anomaly signals poses an additional challenge to standard notions of market efficiency. It implies that information may be incorporated into prices more slowly than previously thought. Our results indicate that investors underreact to the predictable component of accounting data used in anomaly portfolio formation. As a result, returns move predictably in the direction of information that has yet to be released. Moreover, the persistence of anomaly predictability suggests that anomalies may meaningfully affect firms’ discount rates, with potential implications for capital misallocation in the economy.

In sum, our findings challenge the traditional methods for defining and measuring asset pricing anomalies. The results are inconsistent with anomaly explanations related to data mining or risk, and instead suggest that anomalies are related to information about anomaly signals. In other words, many anomalies are more anomalous than previously recognized.

References

- Abarbanell, Jeffery S, and Brian J Bushee, 1998, Abnormal returns to a fundamental analysis strategy, *Accounting Review* 19–45.
- Abarbanell, Jeffrey S, and Brian J Bushee, 1997, Fundamental analysis, future earnings, and stock prices, *Journal of Accounting Research* 35, 1–24.
- Alwathainani, Abdulaziz M, 2009, Consistency of firms’ past financial performance measures and future returns, *The British Accounting Review* 41, 184–196.
- Arnott, Robert D, Christopher Brightman, Vitali Kalesnik, and Lillian Wu, 2023, Earning alpha by avoiding the index rebalancing crowd, *Financial Analysts Journal* 79, 76–97.
- Asness, Clifford, and Andrea Frazzini, 2013, The Devil in HML’s Details, *Journal of Portfolio Management* 39, 49–68.
- Asness, Clifford S, Andrea Frazzini, and Lasse Heje Pedersen, 2019, Quality minus junk, *Review of Accounting Studies* 24, 34–112.
- Baba-Yara, Fahiz, Martijn Boons, and Andrea Tamoni, 2024, Persistent and transitory components of firm characteristics: Implications for asset pricing, *Journal of Financial Economics* 154, 103808.
- Balakrishnan, Karthik, Eli Bartov, and Lucile Faurel, 2010, Post loss/profit announcement drift, *Journal of Accounting and Economics* 50, 20–41.
- Ball, Ray, and Philip Brown, 1968, An empirical evaluation of accounting income numbers, *Journal of Accounting Research* 6, 159–178.
- Bartram, Söhnke M., and Mark Grinblatt, 2018, Agnostic fundamental analysis works, *Journal of Financial Economics* 128, 125–147.

- Bartram, Söhnke M., and Mark Grinblatt, 2021, Global market inefficiencies, *Journal of Financial Economics* 139, 234–259.
- Basu, Sanjoy, 1977, Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis, *The Journal of Finance* 32, 663–682.
- Beneish, M.D., C.M. Lee, and D.C. Nichols, 2015, In short supply: Short-sellers and stock returns, *Journal of Financial Economics* 60, 33–57.
- Bernard, Victor L, and Jacob K Thomas, 1989, Post-earnings-announcement drift: delayed price response or risk premium?, *Journal of Accounting Research* 27, 1–36.
- Bernard, Victor L, and Jacob K Thomas, 1990, Evidence that stock prices do not fully reflect the implications of current earnings for future earnings, *Journal of Accounting and Economics* 13, 305–340.
- Bowles, Boone, Adam V Reed, Matthew C Ringgenberg, and Jacob R Thornock, 2024, Anomaly time, *The Journal of Finance* 79, 3543–3579.
- Bradshaw, Mark T, Scott A Richardson, and Richard G Sloan, 2006, The relation between corporate financing activities, analysts’ forecasts and stock returns, *Journal of Accounting and Economics* 42, 53–85.
- Brennan, Michael J., and Yihong Xia, 2001, Assessing asset pricing anomalies, *The Review of Financial Studies* 14, 905–942.
- Campbell, John Y, Jens Hilscher, and Jan Szilagyi, 2008, In search of distress risk, *The Journal of Finance* 63, 2899–2939.
- Caplin, Andrew, and Mark Dean, 2015, Revealed preference, rational inattention, and costly information acquisition, *American Economic Review* 105, 2183–2203.

- 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, Andrew Y, and Tom Zimmermann, 2019, Publication bias and the cross-section of stock returns, *Review of Asset Pricing Studies* , Forthcoming.
- Chen, Andrew Y, and Tom Zimmermann, 2021, Open source cross-sectional asset pricing, *Critical Finance Review*, Forthcoming , Forthcoming.
- Chen, Andrew Y., and Tom Zimmermann, 2022, Open source cross-sectional asset pricing, *Critical Finance Review* 27, 207–264.
- Cooper, Michael J, Huseyin Gulen, and Michael J Schill, 2008, Asset growth and the cross-section of stock returns, *The Journal of Finance* 63, 1609–1651.
- 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.
- Dichev, Ilia D, 1998, Is the risk of bankruptcy a systematic risk?, *The Journal of Finance* 53, 1131–1147.
- D’Souza, Julia, K Ramesh, and Min Shen, 2010, Disclosure of gaap line items in earnings announcements, *Review of Accounting Studies* 15, 179–219.
- Engelberg, Joseph, R David McLean, and Jeffrey Pontiff, 2018, Anomalies and news, *The Journal of Finance* 73, 1971–2001.
- Fairfield, Patricia M, J Scott Whisenant, and Teri Lombardi Yohn, 2003, Accrued earnings and growth: Implications for future profitability and market mispricing, *The Accounting Review* 78, 353–371.
- Fama, Eugene F., 1991, Efficient Capital Markets: II, *The Journal of Finance* 46, 1575–1617.

- Fama, Eugene F, Lawrence Fisher, Michael C Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information, *International Economic Review* 10, 1–21.
- Fama, Eugene F, and Kenneth R French, 1992, The cross-section of expected stock returns, *The Journal of Finance* 47, 427–465.
- Fama, Eugene F, and Kenneth R French, 2000, Forecasting profitability and earnings, *The Journal of Business* 73, 161–175.
- Foster, George, Chris Olsen, and Terry Shevlin, 1984, Earnings releases, anomalies, and the behavior of security returns, *The Accounting Review* 574–603.
- Gabaix, Xavier, 2019, Behavioral inattention, in *Handbook of behavioral economics: Applications and foundations 1*, volume 2, 261–343 (Elsevier).
- Greenwood, Robin M., and Marco Sammon, 2022, The disappearing index effect, *National Bureau of Economic Research Working Paper* .
- Hafzalla, Nader, Russell Lundholm, and E Matthew Van Winkle, 2011, Percent accruals, *The Accounting Review* 86, 209–236.
- Hartzmark, Samuel M., and David H. Solomon, 2024, Marketwide predictable price pressure, *Working Paper* .
- Haugen, Robert A., and Nardin L. Baker, 1996, Commonality in the determinants of expected stock returns, *The Journal of Financial Economics* 41, 401–439.
- Hirshleifer, David, Kewei Hou, Siew Hong Teoh, and Yinglei Zhang, 2004, Do investors overvalue firms with bloated balance sheets?, *Journal of Accounting and Economics* 38, 297–331.
- Hirshleifer, Jack, 1971, The private and social value of information and the reward to inventive activity, *The American Economic Review* 61, 561–574.

- Hou, Kewei, Chen Xue, and Lu Zhang, 2020, Replicating anomalies, *The Review of Financial Studies* 33, 2019–2133.
- Ikenberry, David L, and Sundaresh Ramnath, 2002, Underreaction to self-selected news events: The case of stock splits, *The Review of Financial Studies* 15, 489–526.
- Ivkvic, Zoran, and Morad Zekhnini, 2024, Dissecting the timing of anomalies, *Working Paper* .
- Jegadeesh, Narasimhan, and Joshua Livnat, 2006, Revenue surprises and stock returns, *Journal of Accounting and Economics* 41, 147–171.
- Jensen, Theis Ingerslev, Bryan T Kelly, and Lasse Heje Pedersen, 2021, Is there a replication crisis in finance?, *National Bureau of Economic Research Working Paper* .
- Keloharju, Matti, Juhani T. Linnainmaa, and Peter Nyberg, 2021, Long-term discount rates do not vary across firms, *Journal of Financial Economics* 141, 946–967.
- Kothari, Sri Prakash, 2001, Capital markets research in accounting, *Journal of Accounting and Economics* 31, 105–231.
- Koudijs, Peter, 2016, The boats that did not sail: Asset price volatility in a natural experiment, *The Journal of Finance* 71, 1185–1226.
- Lakonishok, Josef, Andrei Shleifer, and Robert W Vishny, 1994, Contrarian investment, extrapolation, and risk, *The Journal of Finance* 49, 1541–1578.
- Lee, Charles MC, 2001, Market efficiency and Accounting Research: a discussion of ”Capital Market Research in Accounting” by SP Kothari, *Journal of Accounting and Economics* 31, 233–253.
- Lee, Charles MC, Eric C So, et al., 2015, Alphanomics: The informational underpinnings of market efficiency, *Foundations and Trends® in Accounting* 9, 59–258.

- Lev, Baruch, and Doron Nissim, 2004, Taxable income, future earnings, and equity values, *The Accounting Review* 79, 1039–1074.
- Lev, Baruch, and S Ramu Thiagarajan, 1993, Fundamental information analysis, *Journal of Accounting Research* 31, 190–215.
- Lockwood, Larry, and Wikrom Prombutr, 2010, Sustainable growth and stock returns, *The Journal of Financial Research* 33, 519–538.
- McLean, R David, and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability?, *The Journal of Finance* 71, 5–32.
- Nagel, Stefan, 2005, Short sales, institutional investors and the cross-section of stock returns, *Journal of Financial Economics* 78, 277–309.
- Novy-Marx, Robert, 2010, Operating leverage, *Review of Finance* 15, 103–134.
- Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of Financial Economics* 108, 1–28.
- Piotroski, Joseph D., 2000, Value investing: The use of historical financial statement information to separate winners from losers, *Journal of Accounting Research* 38, 1–41.
- Richardson, Scott, Irem Tuna, and Peter Wysocki, 2010, Accounting anomalies and fundamental analysis: A review of recent research advances, *Journal of Accounting and Economics* 50, 410–454.
- Richardson, Scott A, Richard G Sloan, Mark T Soliman, and Irem Tuna, 2005, Accrual reliability, earnings persistence and stock prices, *Journal of Accounting and Economics* 39, 437–485.
- Sammon, Marco, 2025, Passive ownership and price informativeness, *Management Science* 71, 4582–4598.

- Samuelson, Paul A., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* 6, 41–49.
- Sloan, Richard G, 1996, Do stock prices fully reflect information in accruals and cash flows about future earnings?, *The Accounting Review* 289–315.
- Soliman, Mark T, 2008, The use of dupont analysis by market participants, *The Accounting Review* 83, 823–853.
- Thomas, Jacob K, and Huai Zhang, 2002, Inventory changes and future returns, *Review of Accounting Studies* 7, 163–187.
- Titman, Sheridan, K. C. John Wei, and Feixue Xie, 2004, Capital Investments and Stock Returns, *Journal of Financial and Quantitative Analysis* 39, 677–700.
- van Binsbergen, Jules H, Martijn Boons, Christian C Opp, and Andrea Tamoni, 2023, Dynamic asset (mis) pricing: Build-up versus resolution anomalies, *Journal of Financial Economics* 147, 406–431.
- van Binsbergen, Jules H., and Christian C. Opp, 2019, Real anomalies, *The Journal of Finance* 74, 1659–1706.
- Xie, Hong, 2001, The mispricing of abnormal accruals, *The Accounting Review* 76, 357–373.




Figure 1. Anomaly Returns in Event Time

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after annual information release dates. The returns have been scaled to be zero on the information release date. The compound returns, relative to the information release date, have been highlighted on the 20th trading day before and the 20th trading day after the information date.




Figure 2. Heatmap of Model Quality - F1 Scores

The heatmap shows F1 Scores by anomaly and using four different prediction models: AR Signals, AR Rankings, Machine Learning, and the Martingale model. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives).




Figure 3. Scatter plot of Model Quality - Precision and Recall

The scatter plot shows the precision and recall scores for the 28 anomalies in our study and using three different prediction models: AR Rankings, Machine Learning, and the Martingale model. The point for the Average anomaly is shown with the model's abbreviation. Precision measures the proportion of true positive predictions among all positive predictions. Recall measures the proportion of true positive predictions among all actual positive cases.




Figure 4. Heatmap of Martingale Model Quality by Year - F1 Scores

The heatmap shows F1 Scores by anomaly and by year using the Martingale model. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly and year has a balance of high precision (few false positives) and high recall (few false negatives).




Figure 5. Anomaly Returns with the Martingale Predictions

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after the information release dates. The blue line shows the return for the Perfect Foresight model. The green line shows the return for the Martingale model. The returns have been scaled to be zero on the information release date.






Figure 6. Anomaly Returns for Prediction Outcomes

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after the information release dates and according to the prediction outcomes from the Martingale model. The blue line shows the return for the correct predictions. The purple line shows the return for false positive predictions. The yellow line shows the return for the false negative predictions. The returns have been scaled to be zero on the information release date.




(a) Early Sample Period (1990-2006)




(b) Recent Sample Period (2007-2023)

Figure 7. Anomaly Returns with the Martingale Predictions - Time Trends

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after the information release dates and according to the time period. The top panel shows returns over the early part of the sample period (1990-2006) while the bottom panel shows returns over the recent part of our sample (2007-2023). The blue line shows the return for the Perfect Foresight model. The green line shows the return for the Martingale model. The returns have been scaled to be zero on the information release date.




(a) Early Sample Period (1990-2006)




(b) Recent Sample Period (2007-2023)

Figure 8. Anomaly Returns for Prediction Outcomes - Time Trends

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after the annual information release dates and according to the prediction outcomes from the Martingale model and according to the time period. The top panel shows returns over the early part of the sample period (1990-2006) while the bottom panel shows returns over the recent part of our sample (2007-2023). The blue line shows the return for the correct predictions. The purple line shows the return for false positive predictions. The yellow line shows the return for the false negative predictions. The returns have been scaled to be zero on the information release date.



(a) Early Sample Period (1990-2006)



(b) Recent Sample Period (2007-2023)

Figure 9. Anomaly Returns for Prediction Outcomes at Q2 - Time Trends

This figure shows the compound anomaly returns to the Average anomaly in the 120 trading days (six months) before and 60 trading days after annual information release dates and according to the prediction outcomes for the Martingale model at Q2 and according to the time period. The top panel shows returns over the early part of the sample period (1990-2006) while the bottom panel shows returns over the recent part of our sample (2007-2023). The blue line shows the return for the correct predictions. The purple line shows the return for false positive predictions. The yellow line shows the return for the false negative predictions. The returns have been scaled to be zero on the information release date.






Figure 10. Boxplot of F1 Scores by Anomaly Group


This figure shows the F1 Scores of all anomalies and the anomalies in both the hard-to-predict and easy-to-predict groups.




(a) Perfect Foresight (Early)



(b) Perfect Foresight (Recent)




(c) Martingale (Early)




(d) Martingale (Recent)

Figure 11. Anomaly Returns by Predictor Quality and Time Trends

This figure shows the pre-release returns for the Average anomaly among easy-to-predict anomalies and hard-to-predict anomalies. The figure also shows comparisons between the early years in our sample (1990–2006) and the recent period (2007–2023). Panels (a) and (b) show the Perfect Foresight returns across the two predictor-quality groups and for the two periods. Panels (c) and (d) show the returns to the Martingale model. All returns are scaled to zero at the information release date.




(a) Perfect Foresight (Recent) Q2




(b) Martingale (Recent) Q2

Figure 12. Anomaly Returns by Predictor Quality and Time Trends at Q2

This figure compares the pre-release returns for the Average anomaly among easy-to-predict anomalies and hard-to-predict anomalies, including 120 days before the information release. The figure compares the Perfect Foresight Model (Panel a) and the Martingale Model (Panel b) for the recent period (2007–2023). All returns are scaled to zero at the information release date.



(a) Perfect Foresight



(b) Q3 Model

Figure 13. Anomaly Returns and Institutional Ownership

This figure shows full-sample anomaly returns using both the Perfect Foresight specification (Panel a) and the Q3 Martingale model (Panel b) after splitting stocks by institutional ownership using the low-decile (D1) tail. All returns are scaled to zero at the information release date.

Table I.
Summary Statistics

The table provides summary statistics for our sample which runs from 1990 through 2023. Panel A provides summary statistics for daily returns and total assets (in millions of USD) for all stocks in our sample. Panel B provides summary statistics for each of the anomaly variables.

(1)	(2)	(3)	(4)	(5)
Panel A. Daily Returns and Total Assets				
	Mean	Std. Dev.	Median	N
Daily Raw Returns	8 bps	508 bps	0 bps	112,254,604
Assets	\$7,915	\$63,698	\$477	2,091,805
Panel B. Anomaly Characteristics				
Anomaly (abbreviation)	Mean	Std. Dev.	Median	N
Accruals (Acc)	-0.04	0.09	-0.04	70,672
Asset Growth (Ag)	0.14	0.42	0.06	89,681
Asset Turnover (At)	2.17	3.10	1.34	85,117
Change In Asset Turnover (Cat)	0.00	1.21	0.00	73,479
Change In Profit Margin (Cpm)	0.01	1.03	0.00	87,341
Earnings Consistency (Ec)	0.04	0.77	0.07	43,667
Earnings Surprise (Es)	-0.39	3.18	-0.05	64,078
Gross Profitability (Gp)	0.21	0.29	0.18	103,622
Inventory Growth (Ig)	0.01	0.04	0.00	88,261
Investments (Inv)	1.03	0.80	0.89	54,644
Growth In LT Net Operat. Assets (Ltg)	0.05	0.16	0.04	74,790
Non-Current Operating Assets (Nca)	0.00	0.12	0.00	89,697
Net Operating Assets (Noa)	0.52	0.39	0.54	83,702
Net Working Capital (Nwc)	0.00	0.08	0.00	70,666
Operating Leverage (Ol)	0.76	0.73	0.58	103,626
O-Score (Osc)	-0.52	3.36	-1.01	68,007
Profit Margin (Pm)	-0.15	3.36	0.35	100,963
Percent Operating Accruals (Poa)	-1.97	6.13	-0.66	94,524
Profitability (Pro)	-0.05	0.29	0.02	89,666
Percent Operating Accruals (Pta)	0.75	6.21	0.27	85,690
Return On Equity (Roe)	-0.06	0.85	0.07	103,606
Revenue Surprise (Rs)	-0.36	3.78	0.15	60,110
Sales Growth (Sag)	1,124	483	1,083	47,753
Sustainable Growth (Sg)	0.17	0.62	0.07	84,947
Sales Growth Less Invest. Growth (Sli)	-0.06	1.03	0.03	53,740
Sales Growth Less Exp. Growth (Slx)	0.02	0.47	0.00	60,104
Taxes (Tx)	0.95	1.27	0.99	72,233
Total External Finance (Txf)	0.07	0.24	0.00	83,374

Table II. Event Time Returns and Perfect Foresight

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates. The first row shows the abnormal return over the first three months (60 trading days) after the information release date. The next rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points, basis points per day, and annualized return in percent.

(1)	(2)	(3)	(4)
	Compound Return (bps)	Basis Points Per Day	Annualized Return (%)
Return from trading <i>after</i> the information release date and holding for 3 months.	178	3.1	7.7
Additional return for trading <i>before</i> the information release date:	Compound Return (bps)	Basis Points Per Day	Annualized Return (%)
1 trading day before	39	19.6	49.4
3 trading days before	53	13.2	33.2
1 week before	69	11.5	29.0
2 weeks before	89	8.1	20.4
1 month before	120	5.7	14.4
2 months before	161	3.9	9.9
3 months before	171	2.8	7.1

Table III. Serial Correlation of Anomaly Signals

The table shows pairwise correlations between year end information ($Q4_t$) with the information from previous quarters (e.g., $Q3_t$, $Q2_t$, and $Q4_{t-1}$). Correlations are reported separately for anomaly signals and anomaly rankings (into ten deciles). The row for the *Average* measures the average correlation across all 28 anomalies.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Anomaly	Correlations with $Q4_t$					
	Signals			Rankings (1-10)		
	$Q3_t$	$Q2_t$	$Q4_{t-1}$	$Q3_t$	$Q2_t$	$Q4_{t-1}$
Average	.41	.31	.05	.81	.71	.37
Accruals	.59	.60	-.48	.75	.62	.22
Asset Growth	.94	.99	.01	.84	.69	.26
Asset Turnover	.11	.02	.07	.97	.96	.92
Chg in Asset Turnover	.18	.01	-.62	.77	.67	.12
Chg in Profit Margin	.53	.48	-.64	.82	.70	-.03
Earnings Consistency	.38	.23	.81	.83	.73	.63
Earnings Surprise	.16	.10	-.07	.46	.30	-.37
Gross Profitability	.45	.19	.04	.97	.95	.88
Inventory Growth	.80	.14	.07	.70	.56	.10
Investments	.15	.18	.01	.88	.77	.31
Growth in LT Net Operating Assets	.77	.48	.28	.74	.60	.06
Non-current Operating Assets	.06	.11	.29	.70	.54	-.08
Net Operating Assets	.08	.11	.00	.94	.90	.82
Net Working Capital	.70	.71	-.50	.65	.51	-.16
Operating Leverage	.34	.22	.12	.97	.96	.91
O-Score	.29	.18	.22	.93	.87	.80
Profit Margin	.67	.60	.29	.97	.95	.88
Percent Operating Accruals	-.00	-.00	-.00	.77	.63	.40
Profitability	.56	.38	.17	.93	.88	.72
Percent Total Accruals	.00	.00	-.00	.71	.57	.08
Return on Equity	.00	.00	.00	.87	.79	.59
Revenue Surprise	.23	.03	-.00	.67	.43	-.23
Sales Growth	.90	.87	.85	.84	.79	.76
Sustainable Growth	.78	.21	-.00	.84	.71	.24
Sales Growth Less Inventory Growth	.29	.24	-.00	.75	.60	.23
Sales Growth Less Expenses Growth	.99	.96	-.07	.91	.83	.33
Tax	.00	.01	-.00	.56	.53	.57
Total External Financing	.65	.59	.58	.86	.73	.41

Table IV. Persistence of Anomaly Portfolios

The table shows the percent of stocks in the anomaly portfolio as of the release of year end information ($Q4_t$) that remained the same as of previous information releases. For example, 63% of the Accruals anomaly portfolio at $Q4_t$ remained unchanged from the $Q3_t$ information release. The row for *Average* measures the average percentages across all 28 anomalies.

(1)	(2)	(3)	(4)
	Pct. of Q4 Portfolio from Prior Quarter		
	$Q3_t$	$Q2_t$	$Q4_{t-1}$
Anomaly	Remained	Remained	Remained
Average	66	56	32
Accruals	63	50	22
Asset Growth	72	59	23
Asset Turnover	85	82	60
Chg in Asset Turnover	67	59	26
Chg in Profit Margin	68	57	19
Earnings Consistency	57	43	42
Earnings Surprise	26	13	2
Gross Profitability	85	78	66
Inventory Growth	68	56	22
Investments	66	54	31
Growth in LT Net Operating Assets	65	52	12
Non-current Operating Assets	59	45	12
Net Operating Assets	77	74	59
Net Working Capital	57	45	11
Operating Leverage	84	78	63
O-Score	81	72	57
Profit Margin	87	81	64
Percent Operating Accruals	55	42	24
Profitability	78	68	52
Percent Total Accruals	51	38	14
Return on Equity	75	65	44
Revenue Surprise	49	24	4
Sales Growth	57	50	45
Sustainable Growth	69	56	21
Sales Growth Less Inventory Growth	67	54	26
Sales Growth Less Expenses Growth	86	77	31
Tax	35	30	25
Total External Financing	68	55	31

Table V. F1 Scores for Prediction Models

The table shows the F1 Scores by anomaly and using four different prediction models plus the benchmark: AR Signals, AR Rankings, Machine Learning, and the Martingale model. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. For each anomaly, the reported F1 Score is the average of the F1 Scores for the long-leg and short-leg binary classifications. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives). The row for *Avg* measures the average F1 Score across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)
Anomaly	Benchmark	AR Signals	AR Rankings	Machine Learning	Martingale
Avg	.50	.62	.68	.69	.71
Acc	.41	.58	.60	.63	.63
Ag	.53	.72	.72	.75	.74
At	.77	.87	.88	.90	.91
Cat	.48	.44	.70	.70	.71
Cpm	.52	.55	.69	.77	.77
Ec	.38	.64	.65	.70	.71
Es	.13	.35	.30	.19	.29
Gp	.76	.85	.87	.89	.89
Ig	.49	.71	.70	.71	.71
Inv	.51	.52	.75	.76	.77
Ltg	.44	.62	.64	.65	.66
Nca	.40	.26	.63	.64	.64
Noa	.62	.79	.79	.80	.80
Nwc	.35	.55	.59	.58	.59
Ol	.76	.86	.89	.90	.90
Osc	.66	.79	.81	.82	.82
Pm	.79	.88	.89	.91	.91
Poa	.35	.05	.55	.61	.61
Pro	.66	.84	.83	.84	.85
Pta	.31	.57	.57	.57	.59
Roe	.59	.68	.77	.78	.78
Rs	.28	.50	.44	.48	.50
Sag	.40	.70	.71	.68	.70
Sg	.51	.72	.72	.74	.74
Sli	.45	.65	.64	.66	.67
Slx	.73	.82	.80	.85	.85
Tx	.14	.25	.23	.18	.30
Txf	.51	.70	.70	.77	.75

Table VI. Precision and Recall for Prediction Models

The table shows the precision and recall scores for the Avg anomaly using three different prediction models: AR Rankings, Machine Learning, and the Martingale model. Precision measures the proportion of true positive predictions among all positive predictions. Recall measures the proportion of true positive predictions among all actual positive cases. The table also reports the number of predictions that are True Positives, False Positives, False Negatives, and True Negatives.

(1)	(2)	(3)	(4)
Metric	AR Rankings	Machine Learning	Martingale
Precision	67%	74%	71%
Recall	72%	70%	73%
True Positives	225,119	219,600	226,936
False Positives	110,578	75,669	90,727
False Negatives	87,684	93,203	85,867
True Negatives	2,958,761	2,993,670	2,978,612

Table VII. Event Time Returns and Prediction Models

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and using both the Perfect Foresight Model and the Martingale Model. The rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points and basis points per day.

(1)	(2)	(3)	(4)	(5)
	Perfect Foresight		Martingale	
Additional return for trading <i>before</i> the info release date:	Compound Return (bps)	Basis Points per day	Compound Return (bps)	Basis Points per day
1 trading day before	39	19.6	11	5.5
3 trading days before	53	13.2	19	4.9
1 week before	69	11.5	27	4.5
2 weeks before	89	8.1	38	3.4
1 month before	120	5.7	57	2.7
2 months before	161	3.9	77	1.9
3 months before	171	2.8	70	1.2

Table VIII. Event Time Returns and Prediction Outcomes

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and relative to the outcomes using the Perfect Foresight Model and the Martingale Model. The prediction outcomes are Correct, False Positives, and False Negatives. The column Correct records the return to anomaly that were correctly predicted to be in the anomaly portfolio at the annual information release date. The False Positives column records the return to anomaly that were predicted to be in the annual anomaly portfolio, but were not (Type 1 errors). The False Negatives column records the return to stocks that were not predicted to be in the annual portfolio but were (Type 2 errors). The rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points.

(1)	(2)	(3)	(4)	(5)
	Compound Return (bps)			
	Perfect Foresight	Martingale		
Additional return for trading <i>before</i> the info release date:		Correct	False Positives	False Negatives
1 trading day before	39	29	-25	69
3 trading days before	53	42	-30	82
1 week before	69	55	-33	108
2 weeks before	89	73	-40	134
1 month before	120	103	-46	167
2 months before	161	128	-30	256
3 months before	171	127	-60	297
Return 3 months <i>after</i>	178	171	-43	198
N	147,943	108,726	42,601	39,217

Table IX. Event Time Returns and Prediction Models - Time Trends

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and relative to the outcomes from various prediction models: the Perfect Foresight Model (PF), the Martingale Model (Q3), and a Machine Learning Model (ML). The sample is also divided into an early period (1990-2006) and a recent period (2007-2023). The prediction outcomes are Correct, False Positives, and False Negatives. The column Correct records the return to stocks that were correctly predicted to be in the anomaly portfolio at the annual information release date. The False Positives column records the return to anomaly that were predicted to be in the annual anomaly portfolio, but were not (Type 1 errors). The False Negatives column records the return to stocks that were not predicted to be in the annual portfolio but were (Type 2 errors). The rows show the additional return from trading before the info date. The returns are presented as compound returns in basis points.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Basis Points Before Info Dates by Prediction Model						
	Early Period (1990-2006)			Recent Period (2007-2023)		
	PF	Q3	ML	PF	Q3	ML
1 trading day before	30	25	29	46	2	4
1 week before	83	56	60	61	8	14
1 month before	176	131	137	83	7	12
3 months before	279	198	201	92	-19	-7
Panel B. Basis Points Before Info Dates by Prediction Outcomes (Martingale)						
	Early Period (1990-2006)			Recent Period (2007-2023)		
	Correct	False Positives	False Negatives	Correct	False Positives	False Negatives
1 trading day before	35	5	14	25	-46	105
1 week before	83	1	82	38	-56	124
1 month before	191	-3	135	46	-74	187
3 months before	260	58	330	30	-138	268

Table X. Returns, Predictor Quality, and Time Trends

The table shows the event time anomaly returns to the Average anomaly from both the hard-to-predict (Hard) and easy-to-predict (Easy) anomaly groups and between the early (1990–2006) and recent (2007–2023) periods. The first row shows the abnormal return over the first three months (60 trading days) after the information release date. The next rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Perfect Foresight Model						
Predictor Quality	Compound Returns (bps)					
	Full Sample		Early Period		Recent Period	
	Hard	Easy	Hard	Easy	Hard	Easy
Return from trading <i>after</i> the information release date, holding 3 months.	163	186	153	224	167	160
Additional return for trading <i>before</i> the information release date:						
1 trading day before	43	37	18	36	59	39
3 trading days before	54	52	37	67	65	43
1 week before	79	64	68	90	87	48
2 weeks before	111	78	101	118	117	53
1 month before	137	111	157	185	124	61
2 months before	213	135	244	222	191	71
3 months before	307	102	368	233	263	4
Panel B. Martingale Model						
Predictor Quality	Compound Returns (bps)					
	Full Sample		Early Period		Recent Period	
	Hard	Easy	Hard	Easy	Hard	Easy
Return from trading <i>after</i> the information release date, holding 3 months.	163	186	153	224	167	160
Additional return for trading <i>before</i> the information release date:						
1 trading day before	-0	17	19	29	-11	10
3 trading days before	2	28	20	49	-8	15
1 week before	8	37	33	68	-8	17
2 weeks before	22	46	63	88	-3	19
1 month before	25	73	104	145	-25	24
2 months before	62	84	167	154	-5	34
3 months before	114	47	289	154	-1	-30

Internet Appendix for “Predicting Anomalies”

This internet appendix provides additional empirical evidence to supplement the main text.¹

¹Citation format: Bowles, Boone, Adam V. Reed, Matthew C. Ringgenberg, and Jacob R. Thornock, Internet Appendix to “Predicting Anomalies.” Working Paper, 2026.

A. Prediction Quality Measures

We study the quality of the four different prediction models by analyzing each model’s precision, recall, and F1 Score. The calculations for the three variables are detailed just below. As a description, the precision of a model measures the proportion of true positive predictions among all positive predictions. A high precision indicates few Type 1 errors. Recall measures the proportion of true positive predictions among all actual positive predictions. A high recall indicates few Type 2 errors. The F1 Score is the harmonic mean of precision and recall and provides a single measure that balances both Type 1 and Type 2 errors.

$$\text{Precision} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}}$$

$$\text{Recall} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}}$$

$$\text{F1 score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

B. *Snapshot Details*

Snapshot is a database within Compustat that provides fundamental information for firms and details when fundamental information was updated in the public record. From the Snapshot manual:²

Compustat Snapshot North America (“Snapshot”) is the premier historical analysis and backtesting database that provides fundamental financial data on publicly traded companies and shows all values obtained in the data collection process. The Snapshot database creates a historical investment environment by showing the information that was available at that time in history.

Snapshot contains fundamental financial data from the reported annual and quarterly sources, retaining original values and all succeeding changes from the income statement, balance sheet, and cash flow statements. In addition, Snapshot provides both the preliminary data and final data with annual, quarterly and year-to-date periodicities known at user-specified observation dates. This information enables you to identify what item values have changed and when they changed.

Within Snapshot, perhaps the most salient variable is the *Effective Date*, which captures the date when a company’s “fundamental data is captured from productions of the S&P Capital IQ’s Compustat database.” In other words, the effective date records that date that Compustat noticed that a new book assets or net income amount had been released to the public. Prior to December 2008, Compustat used weekly productions to record data. Beginning in 2008, Compustat uses intraday productions to record data changes.

Based on the Snapshot manual, the earliest effective date is December 14, 1986. This date applies to financial statements prior to 1986 as this was the date Compustat recorded the fundamental information. Obviously, information released in financial statements in 1976 or

²See Compustat Snapshot User Guide, July 24, 2023.

1984, for example, would have been public knowledge prior to December of 1986. However, the effective date is capturing *when* Compustat recorded the data, which, after 1986 took place weekly and after 2008 took place at least daily. As a result of this, we limit our sample to information releases beginning in 1990.

Finally, we utilize SAS code provided in the Snapshot manual to use the effective dates for fundamental information. In essence, we follow Snapshot's guidance to identify the date when information was updated via the release of new financial information.

C. Construction of Anomaly Variables

This appendix provides details regarding the anomalies included in the paper. Each anomaly is listed, along with the paper originally citing the anomaly, a description of how frequently the original paper rebalanced the anomaly portfolios, and details on anomaly construction.

Anomaly	Paper	Original Rebalancing	Our Calculation
Accruals (<i>Acc</i>)	Sloan (1996)	Firms are ranked into deciles based on accruals. Hedge returns for one year ahead are calculated beginning four months after the end of the fiscal year.	$ACC = \frac{WC_t - WC_{t-1}}{\frac{1}{2}(A_t + A_{t-1})}$ WC = Working Capital A = Total Assets WC = current assets - cash and equivalents - current liabilities + debt in current liabilities + taxes payable - depreciation and amortization
Asset Growth (<i>Ag</i>)	Cooper et al. (2008)	Ranked into deciles at the end of June.	$AG = \frac{A_t - A_{t-1}}{A_{t-1}}$ A = Total Assets
Asset Turnover (<i>At</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$AT = \frac{Sales_t}{\frac{1}{2}(NOA_t + NOA_{t-1})}$ NOA = Net Operating Assets
Change in Asset Turnover (<i>Cat</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$CAT = AT_t - AT_{t-1}$ AT = Asset Turnover (defined previously)
Change in Profit Margin (<i>Cpm</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$CPM = PM_t - PM_{t-1}$ PM = Profit Margin (defined below)
Earnings Consistency (<i>Ec</i>)	Alwathainani (2009)	Consistency is based on the number of years in the preceding five years that the firm has had high earnings (low earnings), defined as falling within the top (bottom) 30th percentile. Portfolio returns are calculated beginning on either January or April first.	$EC = \frac{1}{5}(EG_t + EG_{t-1} + EG_{t-2} + EG_{t-3} + EG_{t-4})$ $EG = \frac{EPS_t - EPS_{t-1}}{(\frac{1}{2})EPS_{t-1} + EPS_{t-2}} //$ If EPS_t is opposite sign of EPS_{t-1} then don't include.

Anomaly	Paper	Original Rebalancing	Our Calculation
Earnings Surprise (<i>Es</i>)	Foster et al. (1984)	Earnings surprises are ranked into deciles quarterly.	$ES = \frac{EPS_t - EPS_{t-1} - Drift}{SD}$ $Drift = \text{mean quarterly } EPS \text{ over the preceding seven quarters.}$ SD is the standard deviation of the difference between the preceding seven <i>EPS</i> values and the drift.
Gross Profitability (<i>GP</i>)	Novy-Marx (2013)	Ranked into quintiles at the end of June.	$GP = \frac{Sales_t - COGS_t}{A_t}$
Inventory Growth (<i>Ig</i>)	Thomas and Zhang (2002)	Ranked into deciles annually. Return calculations begin four months following fiscal year-end.	$INV = \frac{Inv_t - Inv_{t-1}}{(A_t + A_{t-1})/2}$ $Inv = \text{Inventory}$
Investments (<i>Inv</i>)	Titman et al. (2004)	Ranked into deciles annually. Return calculations begin four months following fiscal year-end.	$INV = \frac{CE_t}{(\frac{1}{3})(CE_{t-1} + CE_{t-2} + CE_{t-3})}$ $CE = \frac{CAPX}{Sales}$
Growth in Long Term Net Operating Assets (<i>Ltg</i>)	Fairfield et al. (2003)	Stocks are sorted into deciles based on growth in long-term net operating assets. Returns are calculated beginning in April after fiscal year-end.	$LTG = \frac{NOA_t - NOA_{t-1} - ACC_t}{NOA}$ $NOA = \text{receivables} + \text{inventory} + \text{other current assets} + \text{PP\&E} + \text{intangible assets} + \text{other assets} - \text{accounts payable} - \text{other current liabilities} - \text{other liabilities}$ all scaled by total assets. ACC is defined previously.
Non-current Operating Assets (<i>Nca</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$NCA = \frac{chgOA}{\frac{1}{2}(AT_t + AT_{t-1})}$ $OA = AT - ACT - IVAO - LT + DLC + DLTT$
Net Operating Assets (<i>Noa</i>)	Hirshleifer et al. (2004)	Stocks are sorted into deciles. Returns are calculated beginning 4 months after fiscal year-end	$NOA = \frac{OA_t - OL_t}{A_{t-1}}$ $OA_t = AT_t + CHE_t$ $OL_t = AT_t - DLTT_t - MIB_t - PSTK_t - CEQ_t$

Anomaly	Paper	Original Rebalancing	Our Calculation
Net Working Capital (<i>Nwc</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$NWC = \frac{NWC_t - NWC_{t-1}}{\frac{1}{2}(A_t + A_{t-1})}$ $NWC_t = ACT_t - CHE_t - LCT_t + DLC_t$
Operating Leverage (<i>Ol</i>)	Novy-Marx (2010)	Ranked into quintiles at the end of June.	$OL = \frac{COGS_t + SG\&A_t}{A_t}$
O-Score (<i>Osc</i>)	Dichev (1998)	Ranked into deciles. Returns are calculated beginning six months after fiscal year-end.	$OSC = -1.32 - 0.407(\ln(A)) + 6.03(\frac{L}{A}) - 1.43(\frac{CA-CL}{A}) + 0.076(\frac{CL}{CA}) - 1.72I(L > A) - 2.37(\frac{NI}{A}) - 1.83(\frac{IO}{L}) + 0.285I(NI_t + NI_{t-1} < 0) - 0.521(\frac{NI_t - NI_{t-1}}{ NI_t + NI_{t-1} })$ A = total assets L = total liabilities CA = current assets CL = current liabilities NI = net income IO = income from operations $I()$ is the indicator operator taking the value of one if true and zero otherwise.
Profit Margin (<i>Pm</i>)	Soliman (2008)	Measures control variables from last fiscal year-end. Starts calculating monthly returns during the first month of the fiscal year.	$PM = \frac{Sales - COGS}{Sales}$
Percent Operating Accruals (<i>Poa</i>)	Hafzalla et al. (2011)	Returns are calculated beginning four months after the end of the fiscal year.	$POA = \frac{IB_t - OANCF_t}{ IB_t }$ IB = Income before extraordinary items $OANCF$ = Net cash flow
Profitability (<i>Pro</i>)	Balakrishnan et al. (2010)	Measures profitability at date of earnings announcement and measures returns from earnings announcement date.	$PRO = \frac{Earnings_t}{Assets_{t-1}}$

Anomaly	Paper	Original Rebalancing	Our Calculation
Percent Total Accruals (<i>Pta</i>)	Hafzalla et al. (2011)	Returns are calculated beginning four months after the end of the fiscal year.	$PTA =$ $NI + SSTKY - PRSTKCY -$ $DVY - OANCFY -$ $FINCFY - IVNCFY$ All scaled by absolute value of net income. NI = Income before extraordinary items $OANCFY$ = Net cash flow $SSTKY$ = Sale of common and preferred stock $PRSTKCY$ = Purchase of common and preferred stock DVY = Cash Dividends $FINCFY$ = Net cash from financing activities $IVNCFY$ = Net cash from investment activities
Return on Equity (<i>Roe</i>)	Haugen and Baker (1996)	“We assume a reporting lag of 3 months.” We take this to mean they start 3 months after the fiscal year-end.	$ROE = \frac{NI_t}{BE_t}$ BE = Common Equity+Deferred Taxes NI = Net Income
Revenue Surprise (<i>Rs</i>)	Jegadeesh and Livnat (2006)	Revenue surprises are ranked into quintiles quarterly. Abnormal returns are measured from the earnings announcement date.	$RS = \frac{REV_t - REV_{t-1} - Drift}{SD}$ $Drift$ = mean quarterly REV over the preceding seven quarters. SD is the standard deviation of the difference between the preceding seven REV values and the drift.
Sales Growth (<i>Sag</i>)	Lakonishok et al. (1994)	Ranked into quintiles at the end of April.	$SAG =$ $(5 \times R_t) + (4 \times R_{t-1}) + (3 \times$ $R_{t-2}) + (2 \times R_{t-3}) + (1 \times R_{t-4})$ All scaled by 15. R = rank of sales growth as of earnings announcement.

Anomaly	Paper	Original Rebalancing	Our Calculation
Sustainable Growth (<i>Sg</i>)	Lockwood and Prombutr (2010)	Ranked into deciles or quintiles at the end of June.	$SG = \frac{BE_t}{BE_{t-1}}$
Sales Growth Less Inventory Growth (<i>Sli</i>)	Abarbanell and Bushee (1998)	Anomaly variable measured from annual financial statements for December 31 st fiscal-year-end firms. Returns are measured from April 1 st to March 31 st of following year.	$SLI = SAG - IVG$ $SAG = \frac{REV_t - \frac{1}{2}(REV_{t-1} + REV_{t-2})}{\frac{1}{2}(REV_{t-1} + REV_{t-2})}$ $IVG = \frac{INV_t - \frac{1}{2}(INV_{t-1} + INV_{t-2})}{\frac{1}{2}(INV_{t-1} + INV_{t-2})}$ INV = inventory.
Sales Growth Less Expenses Growth (<i>Slx</i>)	Abarbanell and Bushee (1998)	Anomaly variable measured from annual financial statements for December 31 st fiscal-year-end firms. Returns are measured from April 1 st to March 31 st of following year.	$SLX = SAG - XG$ $SAG = \frac{REV_t - \frac{1}{2}(REV_{t-1} + REV_{t-2})}{\frac{1}{2}(REV_{t-1} + REV_{t-2})}$ $XG = \frac{XSGA_t - \frac{1}{2}(XSGA_{t-1} + XSGA_{t-2})}{\frac{1}{2}(XSGA_{t-1} + XSGA_{t-2})}$ $XSGA$ = selling and administrative expenses.
Tax (<i>Tx</i>)	Lev and Nissim (2004)	Anomaly variable is updated annually at the beginning of May.	$TX_t = \frac{TXFO_t + TXFED_t}{0.35 \times IB_t}$
Total External Financing (<i>Txf</i>)	Bradshaw et al. (2006)	Returns are calculated beginning four months after the end of the fiscal year.	$TXF =$ $SSTKY - PRSTKY - DVY + DLTISY - DLTRY$ Scaled by average of the preceding two years of total assets. $SSTKY$ = Sale of common stock. $PRSTKY$ = Purchase of common stock. DVY = cash dividends. $DLTISY$ = Long term debt issuance. $DLTRY$ = Long term debt reduction.

D. Robustness to Weighting Choices in Model Evaluation and Training

This appendix examines whether our conclusions about relative prediction performance depend on how we weight observations when (i) computing model quality and (ii) training the machine learning model. In the main text, Table V reports equally-weighted F1 Scores and the machine learning model is trained with equally-weighted observations (with class weights to address class imbalance). However, value-weighting may be economically relevant. We therefore replicate the main F1-score comparison under several alternative weighting schemes that separately vary the weighting used to compute F1 Scores and the weighting used in the LightGBM training procedure.

Specifically, Tables AI, AII, and AIII replicate the main comparison while adjusting (a) whether F1 Scores are computed equally-weighted or value-weighted, and (b) whether the machine learning model is trained using equally-weighted or value-weighted observations (these training weights apply only to the machine learning model). Figure A1 summarizes these variants visually. Across all variants, the central takeaway is unchanged: the martingale model remains best overall. The machine learning model is never better than the martingale and, if anything, its relative performance weakens under value-weighting.

[Table AI about here.]

[Table AII about here.]

[Table AIII about here.]

[Figure A1 about here.]

We also evaluate sensitivity of the machine learning model to the class-weight parameter used to upweight the minority class (portfolio-leg membership). Table AIV reports results for class weights on class= 1 observations of 1, 1.5, 2, 2.5, and 3 (the baseline specification in the paper is a class-weight of 2). The results show that a class-weight of 2.0 is better or at

least as good as the other options, and none of these machine learning variants outperform the martingale model.

[Table AIV about here.]

Finally, we also evaluate a machine learning model that predicts percentile ranks (1–100) instead of portfolio inclusion. We do this by adjusting the LightGBM classifier to a LightGBMRegressor algorithm. We compare the prediction quality of this continuous machine learning model with the martingale model in Table AV. The results show that the machine learning model is slightly better than the martingale model overall, but, not at placing stocks in the extreme deciles. For example, the martingale model, abbreviated in Table AV as Q3, has a “hit rate” of 70% (meaning it accurately predicts 70% of the long- and short-leg stocks) while the machine learning model (LightGBMRegressor) has a hit rate of 66%.

[Table AV about here.]

E. Results from Machine Learning Model

Several of the results described in the paper have been replicated while including results with respect to the Machine Learning model.

[Figure A2 about here.]

[Figure A3 about here.]

[Table AVI about here.]

[Figure A4 about here.]

[Table AVII about here.]

[Figure A5 about here.]

[Table AVIII about here.]

[Figure A6 about here.]

F. Time Trends in Trading Volume

Given the findings with respect to returns, we also examine patterns in abnormal trading volume around annual information release dates. To do this we calculate daily abnormal trading volume as a stock’s trading volume on a given day divided by the average daily volume over the 121-trading day period around information release dates. We then compute the average daily abnormal trading volume for the Average anomaly in event time and assuming perfect foresight. For instance, an abnormal volume measure of 0.20 indicates that volume on that day was 20% higher than the six-month average. As in the prior analysis, we divide our sample into early and recent periods to investigate potential shifts over time. Our findings are presented in Figure A7.

[Figure A7 about here.]

Figure A7 shows that trading activity has become more concentrated on information release dates in recent years. In this period, abnormal trading volume peaks at 0.80 on annual information release dates, indicating that volume is 80% higher than the average daily volume. By comparison, abnormal volume in the early period peaked at approximately 0.10. Together, these results indicate that information release dates have gained importance for both anomaly stocks and investors.³

Furthermore, in the early part of our sample abnormal volume remained relatively elevated (around 0.10) for the entire week leading up to annual information release dates. This is in contrast to the singular peak seen in recent years. Although this might initially appear to suggest front-running in the early years, it more likely reflects changes in information release timing. As demonstrated in Bowles et al. (2024), during the 1990s and early 2000s, earnings announcements often preceded 10-K filings by several weeks and typically lacked detailed financial statements. This sequence likely explains the elevated trading volume leading up to annual information releases. By contrast, in the late 2000s and 2010s, firms frequently issued

³This finding corroborates the results in Bowles et al. (2024).

detailed financial statements alongside earnings announcements, which were quickly followed by 10-K reports. Thus, in recent years, the earnings announcement date often functions as the primary information release date.

G. Robustness of Easy-versus-Hard To Predict Results

Figure 11 and Table X showed that returns to easy-to-predict anomalies are lower than returns to hard-to-predict anomalies. These two exhibits classified easy/hard anomalies using the average F1 score among them. In this appendix, we provide evidence that the main result—investors appear to have competed away the pre-release returns related to “cheap” signals first, leaving drift where anomaly signals are harder to predict—survives alternative easy/hard classifications. First, Table AIX and Figure A8 show that the main results are robust to using median or tercile cutoffs to classify easy/hard anomalies using F1 scores. They also show that the results are robust to using an alternative predictability measure, the portfolio persistence measured in Column 2 of Table IV.

Second, Figure A9 and Figure A10 show that these alternate cutoffs for easy/hard also produce similar results as in Figure 11 and Figure 12. That is, the return patterns are consistent with the career concerns and investment horizon channel.

Third, Figure A11 illustrates the relation between anomaly-predictability and pre-release returns using anomaly-level scatter plots. The fitted line is negative in all three plots shown, suggesting that anomalies that are easier to predict and more persistent tend to have lower pre-release returns, while anomalies that are harder to predict and less persistent tend to have higher pre-release returns. This evidence complements our previous results because it shows the same relation in the cross section of anomalies without relying on any cutoff.

Finally, we construct two direct proxies for information-processing difficulty from the anomaly definitions themselves. Specifically, we classify anomalies using two complexity measures. The first, a *simple* complexity score, is based on the number of underlying accounting inputs and the lag depth required for construction. The second, a *hybrid* complexity score, also accounts for uncommon inputs and formula features such as ranking, scaling, differencing, log transformations, and conditional rules.⁴ Figure A12 and Table AX show that

⁴An uncommon input is defined as an accounting variable used in the construction of two or fewer of the 28 anomalies in our sample.

anomalies with higher measured complexity generally exhibit larger pre-release returns. We interpret these results as corroborating evidence that information-processing difficulty helps shape the timing of anomaly returns.

[Table AIX about here.]

[Figure A8 about here.]


[Figure A9 about here.]

[Figure A10 about here.]


[Figure A11 about here.]

[Figure A12 about here.]

[Table AX about here.]



(a) F1 Scores for All Models



(b) F1 Scores for Machine Learning Model

Figure A1. F1 Scores with different weighting schemes

This figure shows the F1 scores from different weighting schemes. F1 scores are calculated using either equal-weighting or value-weighting. Also, the Machine Learning model uses either equal-weighting or value-weighting in the loss function of the LightGBM algorithm.




Figure A2. Heatmap of Machine Learning Model Quality by Year - F1 Scores
The heatmap shows F1 Scores by anomaly and by year using the Machine Learning model. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly and year has a balance of high precision (few false positives) and high recall (few false negatives).






Figure A3. Anomaly Returns with Prediction Models

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after annual information release dates. The blue line shows the return for the Perfect Foresight model. The green line shows the return for the Martingale model. The purple line shows the return for the Machine Learning model. The returns have been scaled to be zero on the information release date.




(a) Martingale Model




(b) Machine Learning Model

Figure A4. Anomaly Returns for Prediction Outcomes

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after the information release dates and according to the prediction quality of the model. The blue line shows the return for the correct predictions. The purple line shows the return for false positive predictions. The yellow line shows the return for the false negative predictions. The top panel shows the various returns for the Martingale model while the bottom panel shows returns for the Machine Learning model. The returns have been scaled to be zero on the information release date.




(a) Early Sample Period (1990-2006)




(b) Recent Sample Period (2007-2023)

Figure A5. Anomaly Returns with Prediction Models - Time Trends

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after annual information release dates and according to the time period. The top panel shows returns over the early part of the sample period (1990-2006) while the bottom panel shows returns over the recent part of our sample (2007-2023). The blue line shows the return for the Perfect Foresight model. The green line shows the return for the Martingale model. The purple line shows the return for the Machine Learning model. The returns have been scaled to be zero on the information release date.



(a) Early Sample Period (1990-2006)



(b) Recent Sample Period (2007-2023)

Figure A6. Anomaly Returns for Machine Learning Prediction Outcomes - Time Trends

This figure shows the compound anomaly returns to the Average anomaly in the 60 trading days (three months) before and after annual information release dates and according to the prediction outcomes from the Machine Learning model and according to the time period. The top panel shows returns over the early part of the sample period (1990-2006) while the bottom panel shows returns over the recent part of our sample (2007-2023). The blue line shows the return for the correct predictions. The purple line shows the return for false positive predictions. The yellow line shows the return for the false negative predictions. The returns have been scaled to be zero on the information release date.






Figure A7. Abnormal Trading Volume - Time Trends


This figure shows the abnormal trading volume for the Average anomaly in the 60 trading days (three months) before and after annual information release dates and according to the time period. The figure shows abnormal volume for the perfect foresight model over the early part of the sample period (1990-2006) and the recent part of our sample (2007-2023).



(a) F1 Median Split



(b) F1 Tercile Cuts (Low vs. High)



(c) Persistence Tercile Cuts (Low vs. High)

Figure A8. Anomaly Returns by Predictor Quality Robustness

This figure shows pre-release returns for the Average anomaly among easy-to-predict anomalies and hard-to-predict anomalies in the full sample under the Perfect Foresight specification. The first panel splits anomalies using median F1 scores. The second panel splits by F1 terciles, reporting only the low and high terciles. The third panel splits by portfolio persistence terciles. All returns are scaled to zero at the information release date.




Figure A9. Anomaly Returns by Predictor Quality - Recent Period

This figure shows pre-release returns for the Average anomaly among easy-to-predict anomalies and hard-to-predict anomalies in the recent sample period under both the Perfect Foresight specification and the Q3 Martingale model. The rows correspond to easy vs. hard classifications using F1 median split, F1 tercile split (low/high only), and persistence tercile split (low/high only). The left column (Panels a, c, and e) reports Perfect Foresight and the right column (Panels b, d, and f) reports Q3 Martingale. All returns are scaled to zero at the information release date.






Figure A10. Anomaly Returns by Predictor Quality, Recent Period, Q2


This figure shows pre-release returns for the Average anomaly among easy-to-predict anomalies and hard-to-predict anomalies in the recent sample period under both the Perfect Foresight specification and the Q2 Martingale model. The rows correspond to easy vs. hard classifications using F1 median split and F1 tercile cuts (low/high only). The left column (Panels a and c) reports Perfect Foresight and the right column (Panels b and d) reports Q2 Martingale. All returns are scaled to zero at the information release date



(a) F1 Score



(b) Portfolio Persistence



(c) Rank Persistence

Figure A11. Anomaly-Level Scatterplots, Full Sample, Perfect Foresight

This figure shows anomaly-level scatterplots relating pre-release returns over the three months before information release dates to the F1 score (Panel a), portfolio persistence (Panel b), and rank persistence (Panel c) in the full sample for the Perfect Foresight specification. The fitted line is the univariate OLS line.




Figure A12. Anomaly Returns and Complexity Scores

This figure shows pre-release returns for the Average anomaly less complex anomalies (easy-to-predict) and more complex anomalies (hard-to-predict) in the full sample period under both the Perfect Foresight specification and the Q3 Martingale model. The top row divides anomalies using the simple complexity score (at the median) and bottom row shows results cutting anomalies at the hybrid-complexity median. All returns are scaled to zero at the information release date.

Table AI. F1 Scores for Prediction Models w/ EW F1 Scores and VW ML Model

The table shows the F1 Scores (equally-weighted) by anomaly and using five different prediction models: Benchmark, AR Signals, AR Rankings, Machine Learning, and the Martingale model. (The Machine Learning prediction model is value-weighted in this table.) The F1 Score measures a model’s accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test’s accuracy. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives). The row for *Avg* measures the average F1 Score across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)
Anomaly	Benchmark	AR Values	AR Rankings	Machine Learning	Martingale
Avg	.50	.62	.68	.68	.71
Acc	.41	.58	.60	.62	.63
Ag	.53	.72	.72	.73	.74
At	.77	.87	.88	.89	.91
Cat	.48	.44	.70	.69	.71
Cpm	.52	.55	.69	.76	.77
Ec	.38	.64	.65	.68	.71
Es	.13	.35	.30	.19	.29
Gp	.76	.85	.87	.88	.89
Ig	.49	.71	.70	.69	.71
Inv	.51	.52	.75	.75	.77
Ltg	.44	.62	.64	.64	.66
Nca	.40	.26	.63	.62	.64
Noa	.62	.79	.79	.79	.80
Nwc	.35	.55	.59	.57	.59
Ol	.76	.86	.89	.89	.90
Osc	.66	.79	.81	.80	.82
Pm	.79	.88	.89	.90	.91
Poa	.35	.05	.55	.60	.61
Pro	.66	.84	.83	.83	.85
Pta	.31	.57	.57	.56	.59
Roe	.59	.68	.77	.77	.78
Rs	.28	.50	.44	.46	.50
Sag	.40	.70	.71	.66	.70
Sg	.51	.72	.72	.73	.74
Sli	.45	.65	.64	.65	.67
Slx	.73	.82	.80	.83	.85
Tx	.14	.25	.23	.20	.30
Txf	.51	.70	.70	.76	.75

Table AII. F1 Scores for Prediction Models w/ VW F1 Scores

The table shows the F1 Scores (value-weighted by market cap) by anomaly and using five different prediction models: Benchmark, AR Signals, AR Rankings, Machine Learning, and the Martingale model. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives). The row for *Avg* measures the average F1 Score across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)
Anomaly	Benchmark	AR Values	AR Rankings	Machine Learning	Martingale
Avg	.41	.59	.65	.67	.68
Acc	.22	.51	.53	.57	.54
Ag	.36	.67	.66	.70	.68
At	.70	.81	.82	.85	.87
Cat	.46	.32	.68	.65	.69
Cpm	.41	.43	.65	.66	.69
Ec	.29	.61	.55	.65	.70
Es	.12	.32	.25	.15	.24
Gp	.62	.77	.82	.83	.84
Ig	.42	.66	.65	.67	.66
Inv	.31	.56	.74	.70	.76
Ltg	.38	.54	.60	.60	.62
Nca	.30	.16	.58	.61	.59
Noa	.52	.74	.75	.77	.77
Nwc	.22	.47	.53	.52	.52
Ol	.73	.90	.91	.93	.94
Osc	.53	.81	.82	.84	.83
Pm	.66	.82	.83	.86	.86
Poa	.29	.02	.56	.63	.62
Pro	.45	.79	.78	.83	.80
Pta	.25	.55	.53	.57	.58
Roe	.52	.65	.79	.78	.79
Rs	.27	.53	.48	.51	.52
Sag	.47	.76	.75	.77	.78
Sg	.37	.65	.64	.67	.66
Sli	.38	.66	.64	.67	.68
Slx	.69	.77	.75	.81	.82
Tx	.13	.27	.25	.23	.28
Txf	.39	.63	.66	.75	.71

Table AIII. F1 Scores for Prediction Models w/ VW F1 Scores and VW ML Model

The table shows the F1 Scores (value-weighted by market cap) by anomaly and using five different prediction models: Benchmark, AR Signals, AR Rankings, Machine Learning, and the Martingale model. (The Machine Learning prediction model is value-weighted in this table.) The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives). The row for *Avg* measures the average F1 Score across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)
Anomaly	Benchmark	AR Values	AR Rankings	Machine Learning	Martingale
Avg	.41	.59	.65	.64	.68
Acc	.22	.51	.53	.53	.54
Ag	.36	.67	.66	.66	.68
At	.70	.81	.82	.82	.87
Cat	.46	.32	.68	.59	.69
Cpm	.41	.43	.65	.62	.69
Ec	.29	.61	.55	.58	.70
Es	.12	.32	.25	.11	.24
Gp	.62	.77	.82	.82	.84
Ig	.42	.66	.65	.63	.66
Inv	.31	.56	.74	.68	.76
Ltg	.38	.54	.60	.58	.62
Nca	.30	.16	.58	.55	.59
Noa	.52	.74	.75	.76	.77
Nwc	.22	.47	.53	.48	.52
Ol	.73	.90	.91	.91	.94
Osc	.53	.81	.82	.83	.83
Pm	.66	.82	.83	.84	.86
Poa	.29	.02	.56	.61	.62
Pro	.45	.79	.78	.80	.80
Pta	.25	.55	.53	.52	.58
Roe	.52	.65	.79	.76	.79
Rs	.27	.53	.48	.45	.52
Sag	.47	.76	.75	.74	.78
Sg	.37	.65	.64	.65	.66
Sli	.38	.66	.64	.66	.68
Slx	.69	.77	.75	.81	.82
Tx	.13	.27	.25	.28	.28
Txf	.39	.63	.66	.73	.71

Table AIV. F1 Scores for Machine Learning Model with Class-Weights Grid

The table shows the F1 Scores by anomaly and the Machine Learning with different class-weights. Specifically, a weight of 1 is applied to all class=0 observations while a different weight, noted below, is applied to class=1 observations. The F1 Score measures a model's accuracy while considering both the precision (the proportion of true positive predictions among all positive predictions) and recall (the proportion of true positive predictions among all actual positive cases) of the predictions, providing a single metric for the test's accuracy. A higher F1 Score indicates that a model for a given anomaly has a balance of high precision (few false positives) and high recall (few false negatives). The row for *Avg* measures the average F1 Score across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Equally-Weighted					Value-Weighted				
	Weight Applied to Class = 1									
Anomaly	1.0	1.5	2.0	2.5	3.0	1.0	1.5	2.0	2.5	3.0
Avg	.67	.69	.69	.70	.70	.64	.67	.67	.67	.67
Acc	.59	.62	.63	.63	.63	.47	.56	.57	.58	.57
Ag	.74	.74	.75	.74	.74	.69	.68	.70	.69	.68
At	.90	.90	.90	.90	.90	.85	.85	.85	.85	.84
Cat	.68	.70	.70	.70	.70	.59	.65	.65	.67	.67
Cpm	.74	.76	.77	.77	.77	.63	.66	.66	.68	.68
Ec	.68	.69	.70	.70	.70	.65	.65	.65	.65	.67
Es	.08	.15	.19	.23	.26	.07	.10	.15	.18	.20
Gp	.89	.89	.89	.89	.89	.85	.85	.83	.84	.83
Ig	.69	.70	.71	.71	.70	.64	.67	.67	.68	.67
Inv	.74	.75	.76	.76	.75	.66	.70	.70	.69	.70
Ltg	.63	.65	.65	.65	.64	.58	.59	.60	.62	.63
Nca	.61	.63	.64	.64	.64	.56	.59	.61	.59	.60
Noa	.80	.80	.80	.80	.79	.74	.77	.77	.77	.76
Nwc	.53	.57	.58	.59	.59	.44	.49	.52	.53	.52
Ol	.90	.90	.90	.89	.89	.93	.93	.93	.92	.92
Osc	.81	.82	.82	.82	.82	.85	.85	.84	.84	.84
Pm	.91	.91	.91	.91	.91	.86	.86	.86	.86	.83
Poa	.56	.60	.61	.62	.61	.60	.63	.63	.64	.63
Pro	.84	.84	.84	.84	.84	.82	.83	.83	.83	.83
Pta	.52	.56	.57	.57	.58	.52	.56	.57	.57	.56
Roe	.77	.78	.78	.78	.77	.77	.78	.78	.78	.77
Rs	.43	.46	.48	.48	.48	.45	.49	.51	.51	.52
Sag	.66	.67	.68	.68	.67	.75	.78	.77	.76	.76
Sg	.73	.74	.74	.74	.74	.67	.67	.67	.67	.67
Sli	.64	.66	.66	.66	.66	.65	.67	.67	.67	.67
Slx	.84	.84	.85	.85	.84	.80	.82	.81	.82	.81
Tx	.08	.14	.18	.22	.24	.12	.19	.23	.26	.29
Txf	.77	.77	.77	.77	.77	.74	.75	.75	.75	.74

Table AV. Metrics Comparing Martingale and Machine Learning Regressions

The table shows metrics comparing the quality of the continuous machine learning model (ML uses LightGBMRegressor) with the quality of the Martingale model (abbreviated below as Q3) using the percentile ranking of anomaly variables from 1 to 100. The results compare the root mean square error, the Spearman ρ , the R-squared, and the hit rate for when the continuous predictions correctly hit in the extreme deciles. The row for *Avg* measures the average metrics across all 28 anomalies. Abbreviations are used for the anomaly names. The full names can be found in the Internet Appendix and in a previous table.

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Anomaly	RMSE		Spearman ρ		R-squared		Hit Deciles	
	Q3	ML	Q3	ML	Q3	ML	Q3	ML
Avg	15.90	14.65	0.82	0.82	0.68	0.69	0.70	0.66
Acc	19.98	18.28	0.76	0.76	0.57	0.59	0.63	0.62
Ag	15.75	14.80	0.84	0.85	0.71	0.73	0.74	0.70
At	5.11	5.01	0.98	0.98	0.97	0.97	0.90	0.87
Cat	19.12	17.96	0.77	0.77	0.59	0.59	0.70	0.64
Cpm	15.71	14.87	0.85	0.85	0.72	0.72	0.75	0.71
Ec	13.82	13.24	0.86	0.86	0.75	0.74	0.67	0.54
Es	30.07	25.63	0.46	0.46	0.21	0.21	0.30	0.27
Gp	5.91	5.85	0.98	0.98	0.96	0.96	0.89	0.89
Ig	21.82	19.70	0.71	0.72	0.51	0.53	0.71	0.69
Inv	13.10	12.73	0.89	0.89	0.79	0.79	0.75	0.71
Ltg	20.21	18.78	0.75	0.75	0.56	0.56	0.66	0.62
Nca	21.71	19.84	0.71	0.71	0.50	0.51	0.64	0.61
Noa	9.02	8.92	0.95	0.95	0.90	0.90	0.80	0.78
Nwc	23.58	21.24	0.66	0.66	0.43	0.45	0.59	0.57
Ol	5.41	5.24	0.98	0.98	0.97	0.97	0.91	0.91
Osc	9.48	9.19	0.94	0.95	0.89	0.90	0.82	0.81
Pm	5.59	5.57	0.98	0.98	0.96	0.96	0.90	0.89
Poa	18.82	16.96	0.78	0.80	0.61	0.64	0.61	0.54
Pro	9.62	9.35	0.94	0.94	0.89	0.89	0.85	0.84
Pta	21.36	19.55	0.71	0.71	0.50	0.51	0.58	0.48
Roe	14.59	13.50	0.87	0.88	0.76	0.78	0.79	0.74
Rs	23.12	21.40	0.68	0.67	0.46	0.45	0.50	0.48
Sag	15.25	14.43	0.86	0.86	0.73	0.74	0.70	0.67
Sg	15.23	14.36	0.85	0.86	0.73	0.74	0.74	0.68
Sli	19.33	17.75	0.75	0.76	0.57	0.59	0.66	0.59
Slx	10.78	10.39	0.93	0.93	0.86	0.86	0.85	0.78
Tx	26.97	21.82	0.55	0.62	0.30	0.41	0.29	0.18
Txf	14.75	13.87	0.86	0.86	0.74	0.75	0.74	0.68

Table AVI. Event Time Returns and Prediction Models

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and using various prediction models: the Perfect Foresight Model, the Martingale Model, and a Machine Learning Model. The rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points and basis points per day. This table is analogous to Table VII.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Perfect Foresight		Martingale		Machine Learning	
Additional return for trading <i>before</i> the info release date:	Compound Return (bps)	Basis Points per day	Compound Return (bps)	Basis Points per day	Compound Return (bps)	Basis Points per day
1 trading day before	39	19.6	11	5.5	13	6.7
3 trading days before	53	13.2	19	4.9	23	5.7
1 week before	69	11.5	27	4.5	31	5.2
2 weeks before	89	8.1	38	3.4	42	3.8
1 month before	120	5.7	57	2.7	61	2.9
2 months before	161	3.9	77	1.9	84	2.0
3 months before	171	2.8	70	1.2	79	1.3

Table AVII. Event Time Returns and Prediction Outcomes

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and relative to the outcomes from various prediction models: the Perfect Foresight Model, the Martingale Model, and a Machine Learning Model. The prediction outcomes are Correct, False Positives, and False Negatives. The column Correct records the return to stocks that were correctly predicted to be in the anomaly portfolio at the annual information release date. The False Positives column records the return to anomaly that were predicted to be in the annual anomaly portfolio, but were not (Type 1 errors). The False Negatives column records the return to stocks that were not predicted to be in the annual portfolio but were (Type 2 errors). The rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points. This table is analogous to Table VIII.

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Compound Return (bps)						
	Perfect Foresight	Martingale		Machine Learning			
Additional return for trading <i>before</i> the info release date:	Correct	Correct	False Positives	False Negatives	Correct	False Positives	False Negatives
1 trading day before	39	29	-25	69	28	-22	68
3 trading days before	53	42	-30	82	40	-23	83
1 week before	69	55	-33	108	53	-26	109
2 weeks before	89	73	-40	134	71	-32	135
1 month before	120	103	-46	167	98	-32	175
2 months before	161	128	-30	256	122	-14	257
3 months before	171	127	-60	297	116	-23	306
Return 3 months <i>after</i>	178	171	-43	198	167	-46	208
N	147,943	108,726	42,601	39,217	105,249	35,338	42,695

Table AVIII. Event Time Returns and Prediction Models - Time Trends

The table shows the event time anomaly returns to the Average anomaly over various time horizons relative to annual information release dates and relative to the outcomes from various prediction models: the Perfect Foresight Model (PF), the Martingale Model (Q3), and a Machine Learning Model (ML). The sample is also divided into an early period (1990-2006) and a recent period (2007-2023). The prediction outcomes are Correct, False Positives, and False Negatives. The column Correct records the return to stocks that were correctly predicted to be in the anomaly portfolio at the annual information release date. The False Positives column records the return to anomaly that were predicted to be in the annual anomaly portfolio, but were not (Type 1 errors). The False Negatives column records the return to stocks that were not predicted to be in the annual portfolio but were (Type 2 errors). The rows show the additional return from trading before the info date. The returns are presented as compound returns in basis points. This table is analogous to Table IX with the addition of Panel C.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Basis Points Before Info Dates by Prediction Model						
	Early Period (1990-2006)			Recent Period (2007-2023)		
	PF	Q3	ML	PF	Q3	ML
1 trading day before	30	25	29	46	2	4
1 week before	83	56	60	61	8	14
1 month before	176	131	137	83	7	12
3 months before	279	198	201	92	-19	-7
Panel B. Basis Points Before Info Dates by Prediction Outcomes (Martingale)						
	Early Period (1990-2006)			Recent Period (2007-2023)		
	Correct	False Positives	False Negatives	Correct	False Positives	False Negatives
1 trading day before	35	5	14	25	-46	105
1 week before	83	1	82	38	-56	124
1 month before	191	-3	135	46	-74	187
3 months before	260	58	330	30	-138	268
Panel C. Basis Points Before Info Dates by Prediction Outcomes (Machine Learning)						
	Early Period (1990-2006)			Recent Period (2007-2023)		
	Correct	False Positives	False Negatives	Correct	False Positives	False Negatives
1 trading day before	37	11	13	22	-44	107
1 week before	81	7	88	37	-48	124
1 month before	186	12	155	42	-62	189
3 months before	248	80	344	22	-95	276

Table AIX. Returns and Predictor Quality Robustness

The table shows the event time anomaly returns to the Average anomaly for the Perfect Foresight specification in the full sample period. The columns report returns for anomalies grouped by the F1 median split, F1 tercile cuts (low/high only), and portfolio-persistence tercile cuts (low/high only). The first row shows the abnormal return over the first three months (60 trading days) after the information release date. The next rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points.

(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Compound Returns (bps)					
	F1 Score Median		F1 Score Terciles		Persistence Terciles	
Predictability Group	Hard	Easy	Hard	Easy	Hard	Easy
Return from trading <i>after</i> the information release date and holding for 3 months.	188	170	163	71	215	71
Additional return for trading <i>before</i> the information release date:						
1 trading day before	31	46	43	61	43	61
3 trading days before	41	62	54	71	54	71
1 week before	64	73	79	76	82	76
2 weeks before	90	89	111	84	115	84
1 month before	116	123	137	99	144	99
2 months before	196	133	213	70	223	70
3 months before	303	66	307	-89	338	-89

Table AX. Event Time Returns and Anomaly Complexity

The table shows event-time anomaly returns to the Average anomaly over the full sample period for complexity-sorted groups. Panel A uses the *simple* complexity score, defined as the number of underlying variables plus lag depth. Panel B uses the *hybrid* complexity score, which adds hard-variable counts (unique variables) and counts for variable transformations to the simple score. The first row of each panel shows the compound return over the first three months (60 trading days) after the information release date. The next rows show the additional return from trading before the information release date. The returns are presented as compound returns in basis points.

(1) Complexity Group	(2)	(3)	(4)	(5)
	Compound Returns (bps)			
	Perfect Foresight	Q3 Model		
	Less	More	Less	More
Panel A. Simple Median (Vars + Lags)				
Return from trading after the information release date and holding for 3 months.	214	131	147	57
Additional return for trading before the information release date:				
1 trading day before	53	21	9	14
3 trading days before	70	30	15	26
1 week before	91	41	22	33
2 weeks before	112	60	27	51
1 month before	137	98	39	78
2 months before	171	147	40	124
3 months before	171	171	11	146
Panel B. Hybrid Median (Adds Unique Variables and Transformations)				
Return from trading after the information release date and holding for 3 months.	196	157	146	61
Additional return for trading before the information release date:				
1 trading day before	35	44	10	12
3 trading days before	49	57	19	19
1 week before	65	74	29	25
2 weeks before	82	98	35	41
1 month before	107	138	53	61
2 months before	137	196	70	89
3 months before	127	233	64	81