A Work Project, presented as part of the requirements for the Award of a Master's degree in Business Analytics from the Nova School of Business and Economics.

SEASONAL MOMENTUM IN EQUITY MARKETS: THE ROLE OF CALENDAR, EARNINGS AND MONETARY POLICY EFFECTS

DAVID WALCH

Work project carried out under the supervision of:

Afonso Januário

01/2026

ABSTRACT

This study examines whether combining distinct seasonality drivers across macroeconomic factors such as monetary policy, microeconomic factors such as earnings announcement intensity, and historical equity return patterns can improve systematic trading performance. While the combined signal historically delivered stronger risk-adjusted returns than its individual components and a buy-and-hold benchmark, its effectiveness declines sharply in recent decades. The results suggest that those market anomalies have weakened over time, consistent with rising market efficiency. These results are in line with research showing that that return anomalies tend to diminish as they attract attention and are arbitraged away by market participants.

Keywords: Calendar Effects; Seasonal Momentum; Systematic Trading Strategies; Market Efficiency; Post-Earnings Announcement Drift (PEAD); Pre-FOMC Drift

This work was funded by Fundação para a Ciência e a Tecnologia (UID/00124/2025, UID/PRR/124/2025, Nova School of Business and Economics) and LISBOA2030 (DataLab2030 - LISBOA2030-FEDER-01314200).

1. Introduction

The search for systematic patterns and elements of predictability in financial markets has long been a central focus for both academics and industry professionals. A central debate within the literature concerns the validity of the Efficient Market Hypothesis (EMH) originally proposed by Fama (1970), which argues that asset prices fully and instantaneously reflect all available information. Supporters of the EMH argue that stock prices follow a stochastic process absent of predictable structure, thereby leaving forecasting efforts ineffective (Enow 2023). Nevertheless, a substantial amount of empirical research has documented persistent deviations from market efficiency, often referred to as anomalies (Lo 2004; Schwert 2002). Among these, seasonality effects, which refer to systematic and recurring patterns in asset returns associated with specific times of the year, remain some of the most extensively studied phenomena in financial research.

This paper explores whether different seasonality effects in equity markets, including calendar-based return patterns, earnings-announcement intensity, and monetary-policy timing, can be jointly used within a single predictive framework. While each of these effects has been documented in prior research, they have typically been analysed in isolation, leaving open the question of how they interact and whether their combined information can improve return predictability. Using a systematic long-short trading framework, the analysis evaluates the individual contribution of each seasonal component and examines whether integrating them leads to superior risk-adjusted performance. The results show that the combined signal historically delivers stronger performance than its individual components, but also reveal pronounced time variation. In particular, the predictive power of the composite signal is largely concentrated in earlier decades and weakens substantially in more recent periods, suggesting that the economic relevance of seasonality-based trading strategies has declined over time.

2. Literature Review: Stock Market Seasonality

It is wieldy documented by academic literature that stock market returns often follow recognizable seasonal patterns. One of the most studied is the month effect, where some months, particularly January, tend to show higher average returns than others. Wachtel (1942) first formally documented that stock returns in January were unusually strong compared to the rest of the year, marking one of the earliest observations of systematic seasonality in equity markets. Twenty-five years later, Rozeff and Kinney (1976) confirmed this pattern with more extensive data, and the so-called January effect has since become one of the most frequently studied and debated return anomalies in financial economics. In addition to calendar-based anomalies, markets exhibit systematic patterns around major information events. One well-documented case is the "Pre-FOMC Announcement Drift", where stock prices tend to rise in the days preceding Federal Open Market Committee announcements, reflecting investor positioning ahead of monetary policy decisions (Lucca and Moench 2015). Similar dynamics occur around corporate earnings releases, as trading activity and volatility often increase not only on the announcement date but also in the days leading up to it, driven by expectations and speculation (Ball and Brown 1968, 170-171). These findings indicate that seasonality in returns extends beyond calendar effects and is also shaped by the anticipation of policy actions and corporate disclosures.

Nevertheless, recent evidence suggests that the effect has largely disappeared in recent years. Martineau (2021) finds that the post-earnings-announcement drift (PEAD) has largely disappeared for large firms since the mid-2000s. Kettell, McInnis, and Zhao (2022) attributes this decline to reduced persistence in earnings news, suggesting that earnings surprises themselves have become less informative about future performance. Fink (2021) similarly document that PEAD effects have become much weaker in recent years. Richardson and Veenstra (2022) even argues that previous research may have misattributed the drift to market inefficiency because it failed to consider new economic information that becomes available after the earnings announcement. A parallel trend is observed for the pre-FOMC announcement drift. Kurov, Wolfe, and Gilbert (2020) show that it has faded in the post-2015 period, and Hu et al. (2022) offer risk-based explanations for its variability. Lee and Wang (2025) further find that pre-FOMC gains are often offset by post-announcement reversals. Overall, both earnings- and policy-related announcement effects appear to have diminished, reflecting a greater market- and information efficiency.

3. Research Question and Hypothesis

The first research question evaluates whether a constructed return-seasonality signal captures well-established calendar effects documented in the literature, ensuring that the measure reflects meaningful intra-year patterns rather than statistical noise.

Research Question 1: Does the constructed return-seasonality signal successfully capture well-documented calendar effects, such as the January effect, turn-of-the-month patterns, holiday effects, and other recurring intra-year return anomalies?

- H_{01}: The return-seasonality signal **does not** capture established calendar effects and **does not** reflect meaningful intra-year return patterns.

- H_{11}: The return-seasonality signal **does** capture established calendar effects and **does** reflects meaningful intra-year return patterns.

Second, it investigates whether these seasonal return patterns can be explained by the timing of micro- and macroeconomic announcements, such as earnings-reporting intensity or monetary-policy communications, to assess whether seasonality reflects its own underlying pattern or is mostly a result of when information is released throughout the year.

Research Question 2: Can return seasonality be explained by micro- / macroeconomic announcements such as monetary policy timing or earnings announcement intensity?

- H_{02}: Return seasonality **is not** fully explained by micro- / macroeconomic announcements such as monetary policy timing or earnings announcement intensity
- H_{12}: Return seasonality **is** fully explained by micro- / macroeconomic announcements such as monetary policy timing or earnings announcement intensity

Third, the thesis explores whether the combined information from return seasonality, earnings intensity, and monetary-policy timing can be used to construct a simple real-time trading rule that indicates improved risk-adjusted performance. This final step serves as a proof of concept and abstracts from transaction costs, slippage, and other practical implementation frictions.

Research Question 3: Can the combined signals be used to create a real-time trading rule that indicates improved risk-adjusted performance?

- H_{03}: The combined signal **does not** indicate improved risk-adjusted performance.
- H_{13}: The combined signal **does** indicate improved risk-adjusted performance.

4. Methodology

4.1 Data

This study combines several data sources to capture equity market dynamics from both micro- and macroeconomic perspectives. Equity market returns are obtained from Yahoo Finance using adjusted daily closing prices, which account for dividends, stock splits, and other corporate actions. The sample spans January 1950 through December 2024, covering more than seven decades of daily observations across multiple market cycles and policy regimes. The daily frequency allows the analysis to capture short-term market reactions to events, while the long sample period facilitates the assessment of persistence and time variation in seasonal patterns. This return series serves as the foundation of the analysis and is augmented with additional datasets on index membership, earnings announcements, and monetary policy events.

Firm-level earnings announcement data are drawn from the Institutional Brokers' Estimate System (IBES), which provides widely used information on reported earnings, prices, and shares outstanding. These data are used to identify announcement dates and to construct firm-level market capitalization measures. By combining this information with historical S&P 500 index membership, the analysis constructs a daily series of the share of index market capitalization represented by firms announcing earnings. Historical index membership is reconstructed from multiple publicly available sources. Constituents from 1996 to 2019 are taken from a dataset compiled by Andreas Clenow, while subsequent years are supplemented using an open-source, web-scraped dataset published on GitHub under the MIT License. This approach ensures transparency and reproducibility while avoiding reliance on proprietary data.

Monetary policy events are obtained from official Federal Open Market Committee (FOMC) statements, which provides a reliable record of policy decisions. The archive extends back only to 1994, which sets the maximum available history for this dataset and restricts the analysis to the modern era of U.S. monetary policy. Prior to 1994, the Federal Reserve did not publicly disclose its policy decisions, which makes consistent data on announcement dates unavailable. The announcement calendar is obtained using the FedTools Python package and transformed into a binary indicator that equals one on statement days and zero otherwise. This structure allows policy announcements to be combined with stock market and earnings data, making it possible to analyse how markets behave around central bank meetings.

4.1.1 Exploratory Data Analysis (EDA)

Figure 1 illustrates the daily returns of the S&P 500 from the 1950s to the early 2020s and highlights several well-known characteristics of financial time series. The distribution of returns is centered around zero, yet distinct periods of heightened volatility are observable during major crises, including the 1987 market crash, the 2008 global financial crisis, and the 2020 COVID-19 shock. The data demonstrates clear volatility clustering, with extended periods of relative stability followed by intervals of extreme fluctuation. Several tail events are visible, most notably the Black Monday collapse of October 19, 1987, when the index declined by more than 20% in a single trading day.

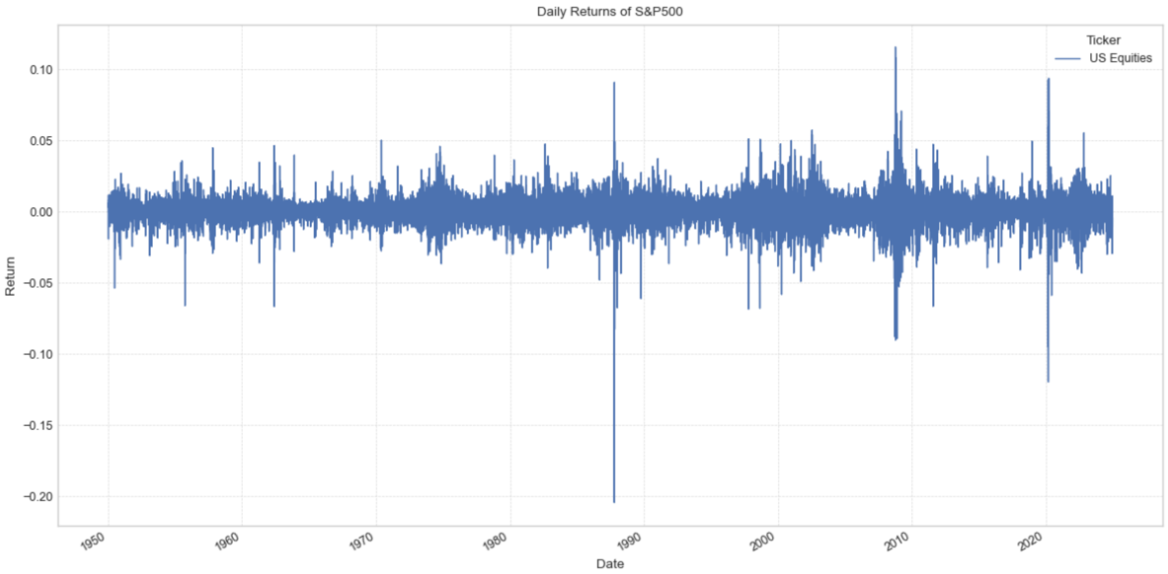

Figure 1: Daily log returns of the S&P 500 (1950 to 2025). Returns are computed from Yahoo Finance daily price data using log differences of the index level. Data source: Yahoo Finance.

As displayed in Figure 2, monetary policy announcements follow clear schedules but the frequency of FOMC statements has not always been as regular as it is today. In earlier decades, the Committee met less frequently, producing fewer statements each year. In recent years, however, the pattern has stabilized at about eight meetings scheduled per year ('Federal Open Market Committee' 2025). March 2020, however, provides a clear example of how this pattern can shift in times of crisis. As the COVID-19 pandemic escalated, the Committee released four statements in a single month. These included two emergency rate cuts that brought the federal funds target range back to near zero, along with measures to support liquidity and financial stability, such as expanded repurchase operations and large-scale asset purchases (Board of Governors of the Federal Reserve System 2020a; Board of Governors of the Federal Reserve System 2020b).

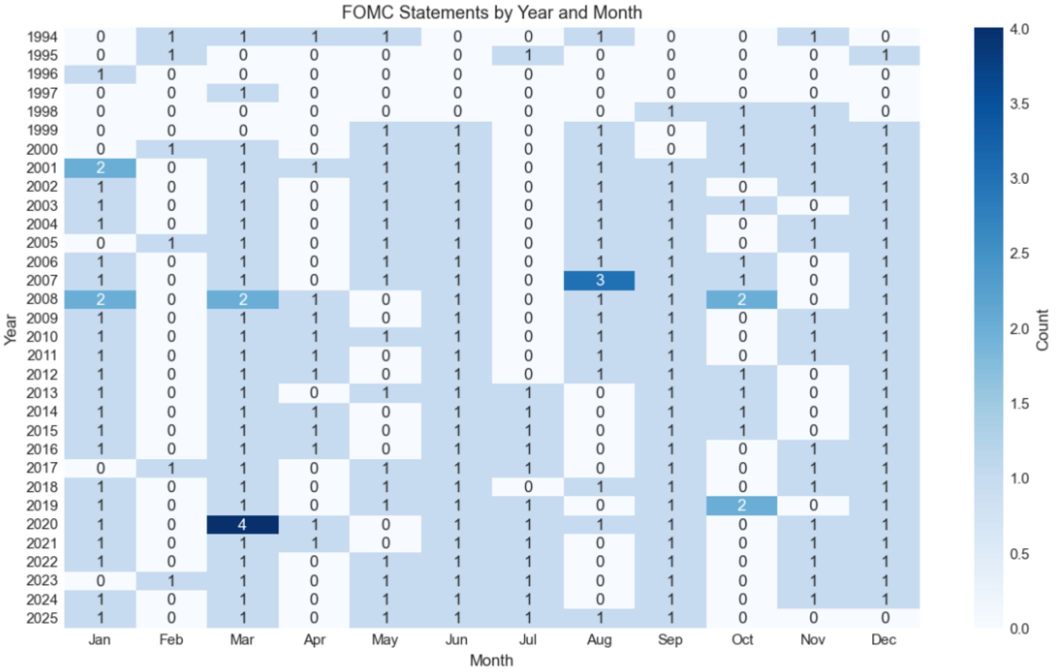

Year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1994	0	1	1	1	1	0	0	1	0	0	1	0
1995	0	1	0	0	0	0	1	0	0	0	0	1
1996	1	0	0	0	0	0	0	0	0	0	0	0
1997	0	0	1	0	0	0	0	0	0	0	0	0
1998	0	0	0	0	0	0	0	0	1	1	1	0
1999	0	0	0	0	1	1	0	1	0	1	1	1
2000	0	1	1	0	1	1	0	1	0	1	1	1
2001	2	0	1	1	1	1	0	1	1	1	1	1
2002	1	0	1	0	1	1	0	1	1	0	1	1
2003	1	0	1	0	1	1	0	1	1	1	0	1
2004	1	0	1	0	1	1	0	1	1	0	1	1
2005	0	1	1	0	1	1	0	1	1	0	1	1
2006	1	0	1	0	1	1	0	1	1	1	0	1
2007	1	0	1	0	1	1	0	3	1	1	0	1
2008	2	0	2	1	0	1	0	1	1	2	0	1
2009	1	0	1	1	0	1	0	1	1	0	1	1
2010	1	0	1	1	1	1	0	1	1	0	1	1
2011	1	0	1	1	0	1	0	1	1	0	1	1
2012	1	0	1	1	0	1	0	1	1	1	0	1
2013	1	0	1	0	1	1	1	0	1	1	0	1
2014	1	0	1	1	0	1	1	0	1	1	0	1
2015	1	0	1	1	0	1	1	0	1	1	0	1
2016	1	0	1	1	0	1	1	0	1	0	1	1
2017	0	1	1	0	1	1	1	0	1	0	1	1
2018	1	0	1	0	1	1	0	1	1	0	1	1
2019	1	0	1	0	1	1	1	0	1	2	0	1
2020	1	0	4	1	0	1	1	1	1	0	1	1
2021	1	0	1	1	0	1	1	0	1	0	1	1
2022	1	0	1	0	1	1	1	0	1	0	1	1
2023	0	1	1	0	1	1	1	0	1	0	1	1
2024	1	0	1	0	1	1	1	0	1	0	1	1
2025	1	0	1	0	1	1	1	1	1	0	0	0

Figure 2: Frequency of FOMC Statements by Year and Month (1994–2025). The figure displays the number of FOMC communication events for each year and calendar month based on official release dates published by the Federal Reserve. Data source: Federal Reserve.

Earnings announcements also display pronounced seasonality, with corporate disclosures clustered in specific months rather than evenly distributed throughout the year. As shown in Figure 3, the fraction of index market capitalization reporting earnings spikes in January-February, April, July, and October, which correspond to the traditional earnings seasons following the close of each fiscal quarter. In contrast, months such as March, June, September, and December reveal minimal reporting activity, functioning as relatively quiet periods in the corporate information cycle.

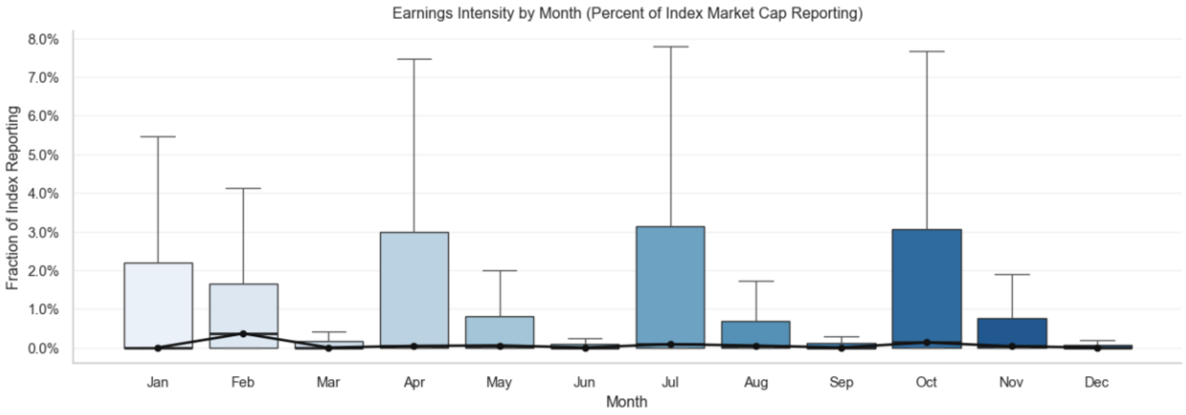

Figure 3: Seasonal Patterns in Earnings Reporting Intensity (1994 to 2025). Daily earnings-report dates are obtained from IBES and linked with market-capitalisation data to compute the percentage of index market cap reporting on each trading day. Boxplots show the monthly distribution of these values. Data sources: IBES..

4.1.2 Market Returns Around Key Economic and Corporate Events

As discussed in the literature review, markets often display both "pre-announcement" and "post-announcement" effects around key economic and corporate events. Although this phenomenon has been well documented, recent studies indicate that its impact has weakened in recent years. To examine whether this pattern also holds in this dataset and whether it could potentially serve as a source of excess return, a brief empirical analysis was conducted.

Figure 4 illustrates the average cumulative returns of the S&P 500 around FOMC announcement dates and heavy earnings days. Heavy-earnings days are defined as trading days on which a large share of the S&P 500's market capitalization reports earnings. These days are identified by applying a threshold set at the 97.5th percentile of the daily percentage of index market capitalization reporting earnings. Days at or above this threshold are treated as events. Around FOMC announcements, returns are largely stable prior to the event but exhibit a distinct upward jump immediately at day 0, indicating a short-term positive market reaction to policy announcements. In contrast, returns around heavy earnings days show a steady upward drift both before and after the event, reflecting sustained positive market sentiment during major earnings periods. The magnitude of the cumulative return is notably larger in the earnings window, suggesting that corporate earnings activity has a stronger overall impact on market performance compared to FOMC announcements.

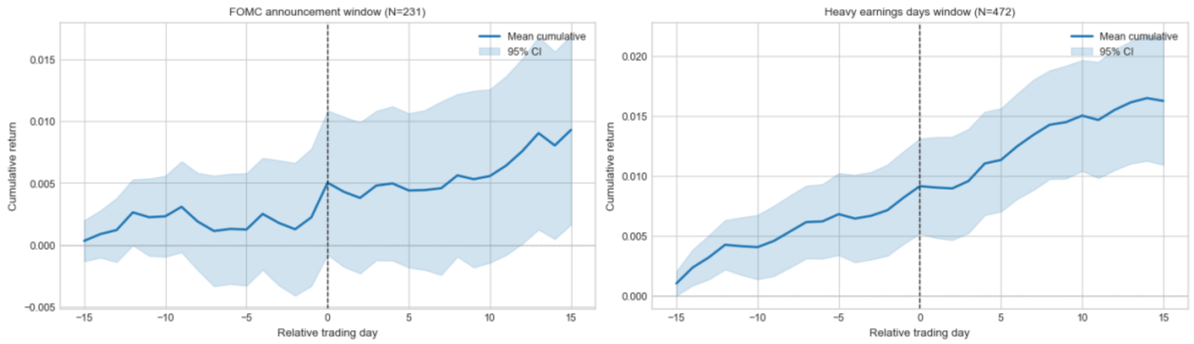

Figure 4: Market Behavior Before and After Key Economic and Earnings Events (1994 to 2025). Returns are aligned relative to event day zero and cumulated over a symmetric window from −15 to +15 trading days. Heavy earnings days are defined as trading days in the top 2.5 percent of S&P 500 market-capitalisation reporting activity. Data sources: Yahoo Finance, IBES, and Federal Reserve.

Furthermore, it is noteworthy that this FOMC jump effect in the dataset has weakened over the past few years, aligning with findings reported in the academic literature. As illustrated in Figure 5, only the pre-2010 period shows a distinct increase between day -1 and day 0. In the subsequent periods (2010-2019 and 2020-present), this effect appears to have largely dissipated, likely reflecting greater market efficiency, as discussed in the literature review. The higher overall magnitude of both pre- and post-announcement drifts observed in the 2020-present period may be explained by the elevated significance of FOMC announcements during this time, particularly in the context of the COVID-19 shock and the 2022 inflation cycle. On the contrary, the impact of earnings announcements seems to have grown in recent years, possibly because the index has become increasingly concentrated in a small number of individual stocks. As a result, index performance is increasingly tied to the earnings outcomes of a few large-cap names, leading to sharper market reactions and greater volatility around reporting periods.

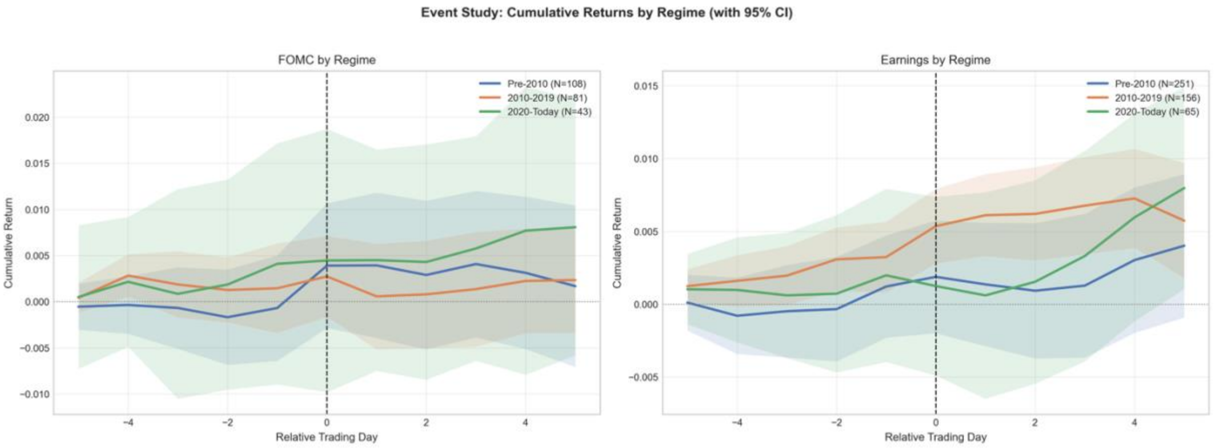

Figure 5: FOMC/Earnings Announcement Drift by Period (1994 to 2025). Returns are aligned relative to event day zero and cumulated over a symmetric window from −5 to +5 trading days. Regimes are defined as Pre-2010, 2010–2019, and 2020–2025. Heavy earnings days correspond to trading days in the top 2.5 percent of S&P 500 market-capitalisation reporting activity. Data sources: Yahoo Finance, IBES, and Federal Reserve.

4.1 Signal Construction

The trading signal is constructed by combining three information sources, namely historical return seasonality, earnings reporting intensity, and the timing of FOMC statements into a single composite predictor. The objective is to create one unified daily signal that captures recurring calendar effects, periods of concentrated earnings activity, and scheduled monetary policy announcements. Their combination is intended to provide a more consistent and informative predictor than any individual input and improve model performance.

4.1.1 Seasonal Return Component

The seasonal component identifies recurring intra-year patterns in daily returns. Daily returns are first reorganized into a year-by-day panel in which rows correspond to calendar years and columns represent the day of the year. This alignment places the same calendar date across multiple years into a common column, enabling the detection of systematic seasonal structure. To ensure a genuine one-step-ahead forecasting framework, each calendar row is shifted down by one year, so that the estimate for year t is based exclusively on data from years $\leq t - 1$, thereby eliminating look-ahead bias.

Across the year dimension, the panel is smoothed using an exponentially weighted moving average (EWMA) with a half-life of λ_signal = 3 years, assigning greater weight to more recent observations while still incorporating longer historical patterns. To further reduce high-frequency noise along the day-of-year dimension, a short, symmetric 5-day kernel with weights (0.05, 0.20, 0.50, 0.20, 0.05) is applied via convolution. The center day receives 50% of the weight, adjacent days receive 20% each, and days two positions away receive 5% each, attenuating random day-to-day fluctuations while preserving the broader seasonal profile. After this panel-based transformation, the smoothed values are mapped back into a date-indexed time series and subjected to an additional 3-day rolling mean to remove residual noise.

4.1.2 FOMC Component

The FOMC component captures the documented tendency for equity returns to behave differently around Federal Reserve policy announcements. Statement dates, sourced from Federal Reserve public records spanning 1994 to the present, are encoded in a binary indicator that takes the value 1 on announcement days and 0 otherwise. Daily market returns are multiplied by this indicator to isolate returns realized specifically on FOMC statement days, with all non-event days set to zero. To maintain a strict one-step-ahead forecasting structure, the conditioned series is shifted forward by one trading day so that information available at time t is not used until time t + 1.

4.1.3 Earnings Intensity Component

The earnings-intensity component captures periods of concentrated corporate disclosure, when a substantial fraction of index constituents release quarterly results simultaneously. For each trading day, the share of total S&P 500 market capitalization represented by firms announcing earnings is computed, ensuring that disclosures by larger firms exert proportionally greater influence. Daily returns are then scaled by this earnings-intensity measure, following the same conditioning procedure as in the FOMC component. The resulting series is likewise shifted forward by one trading day to eliminate look-ahead bias and preserve the forecasting structure of the model.

4.1.4 Standardization via EWMA Z-Score

All three components (the seasonal-return series, the FOMC-conditioned return, and the earnings-conditioned return) are standardized using a rolling z-score computed from an exponentially weighted mean and standard deviation. This EWMA-based standardization, governed by a prespecified half-life, places greater emphasis on recent observations while down-weighting older data and ensures that all predictors operate on a comparable scale prior to aggregation.

4.1.5 Composite Signal

The composite predictor is calculated as the equal-weighted average of the three standardized components: the seasonal-return series, the earnings-intensity measure, and the FOMC-conditioned return. Equal weighting ensures that each source of information contributes proportionally to the final signal and prevents any single component from dominating the outcome. This aggregation functions similarly to an ensemble or majority-vote mechanism, in which the composite signal reflects the consensus direction implied by the underlying indicators. When multiple independent components point in the same direction, the resulting signal is amplified, whereas conflicting inputs yield more muted values. This structure reduces the influence of noise in any one series and helps produce a more stable and robust predictor.

4.2 Position Construction and Trade Sizing

Positions are derived directly from the composite signal through a transparent, rule-based implementation that scales portfolio exposure in direct proportion to the signal’s magnitude. This proportional mapping aligns position sizes with the model’s assessed confidence: stronger signals generate larger allocations, while weaker signals produce more moderate exposure.

Position scaling follows a linear transformation relative to a predefined threshold, after which values are clipped to the range [−1, 1], allowing both long and short positions within symmetric bounds. A position of +1 represents a fully invested long stance, −1 denotes a fully invested short stance, and 0 corresponds to no exposure. For example, a signal of +0.5 results in a 50 percent long position, while a signal of +2.0 is capped at a fully invested long position.

4.3 Look-Ahead Bias Prevention

To ensure a realistic representation of execution dynamics and to prevent any forward-looking information from entering the decision process, all trades are implemented with a one-day lag so that positions at time t + 1 are based solely on information observable at time t. All rolling calculations, including the EWMA mean and standard deviation, rely strictly on trailing windows, and all event-based variables such as FOMC and earnings indicators are shifted forward by one trading day prior to signal construction. These adjustments eliminate look-ahead bias and ensure that the forecasting structure reflects feasible real-time implementation. An initial warm-up period of 252 trading days (approximately one year), corresponding to the length of the rolling estimation windows, is excluded from performance evaluation to avoid distortions arising from incomplete data histories.

4.4 Baseline Parameter Choices

Baseline parameter choices are fixed ex-ante to provide a consistent reference for analysis. A summary of these baseline settings is provided in the parameter table in Appendix Table 6. To assess the robustness and generalizability of the framework, these parameters are systematically varied in a structured sensitivity analysis presented in Section 6. Specifically, λ_signal is tested over the range {0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0} years, λ_zscore is tested over the same range, and θ is tested over {0, 0.5, 1.0, 1.5, 2.0}. This yields a grid of 500 parameter combinations, and results demonstrate the stability of performance across alternative calibration choices.

5. Empirical Results

5.1 Signal Evaluation

Figure 6 presents the smoothed daily z-score of the return-seasonality signal aggregated by month and illustrates the typical intra-year pattern through its median profile, while the shaded region reflects the 10–90 percent range of historical variation. The figure shows that the constructed signal captures several well-documented calendar effects. Early January displays a pronounced rise that aligns with the January effect, indicating that equity returns tend to be stronger at the beginning of the year. The pattern also shows recurring oscillations across months, with local peaks frequently appearing around turn-of-the-month periods, consistent with short-term timing effects widely discussed in the literature. Mid-year and late-year fluctuations are visible as well. The signal tends to weaken in June and July, which corresponds to the seasonal slowdown often described in the “Sell in May” literature, and it strengthens again through October and November, reflecting historically robust autumn and year-end performance.

While the seasonality signal is smooth and continuous, the earnings-intensity and FOMC components behave very differently, as illustrated in the diagnostic plots provided in the Appendix (Figure 10). These two signals remain near zero on most days but generate sharp spikes during concentrated announcement windows, which interrupt the smoother seasonal pattern and contribute to the composite signal’s occasional extreme values. Averaging across components mitigates some of these outliers, yet the composite still inherits noticeable tail behaviour on event days while preserving the underlying seasonal shape during non-event periods. Summary statistics for all four signals, reported in Table 5 in the Appendix, confirm these differences by documenting the extreme kurtosis and wide value ranges of the announcement-driven indicators relative to the well-behaved seasonality component.

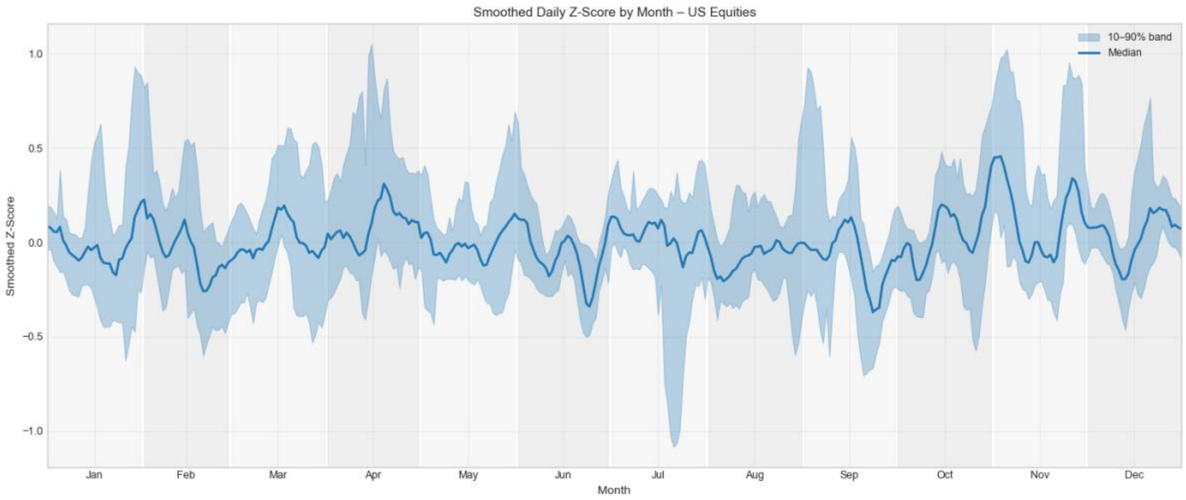

Figure 6: Monthly Pattern of Smoothed Daily Z-Scores (1950 to 2025). Z-scores are obtained by applying an exponential-weighted moving average to daily returns and standardising them using rolling variance estimates. Shaded areas represent the 10–90 percent range. Data source: Yahoo Finance.

5.2 Is Seasonality Explained by Announcement Timing?

The empirical evidence provides strong support for the conclusion that the return-seasonality signal cannot be explained by the timing of macro- or microeconomic announcements. As shown in the correlation matrix in the Appendix (Table 4), the association between seasonality and the announcement-based components is essentially zero: the correlation with the FOMC signal is 0.003 and with the earnings-intensity signal only 0.016. These values are economically negligible and stand in clear contrast to the much stronger correlations between each component and the composite signal, which suggests that the seasonality, earnings, and FOMC indicators each capture distinct information. This interpretation is reinforced by the regression results reported in the Appendix (Table 11 & Table 12), which formally test whether earnings activity or monetary-policy timing can account for variations in seasonal returns. The model yields an R-squared of 0.0004, indicating that the two announcement variables together explain virtually none of the variation in the seasonality series. Both coefficients are economically small, statistically weak, and lack robustness even under heteroscedasticity- and autocorrelation-consistent standard errors. In practical terms, these findings imply that the seasonal return pattern persists after controlling for clustered information releases and is not simply a byproduct of earnings seasons or FOMC cycles. Instead, seasonality appears to reflect its own independent return structure that operates alongside, but separate from, predictable macro- and microeconomic news flows. This provides clear support for rejecting H_{02} and confirms that the seasonality signal captures behaviour not attributable to scheduled information events.

5.3 Composite Signal vs. Single-Component Performance

To check whether the combined signal improves performance over any single component, the strategy is run individually on each component as displayed in Table 1. The results indicate that while both the Return and Earnings components generate positive annualized returns (1.12% and 4.83%, respectively), the FOMC component alone exhibits a slightly negative performance of –0.18%. However, the composite signal achieves the highest annualized return of 7.50% with a Sharpe ratio of 0.74, substantially outperforming all individual components on a risk-adjusted basis. This suggests that the combination of signals provides diversification benefits and captures complementary sources of return that are not fully exploited when the components are traded in isolation.

Metric	Return Component	FOMC Component	Earnings Component	Composite Signal
Annualized Return	0.0112	-0.0018	0.0483	0.0750
Annualized Volatility	0.0995	0.0193	0.1114	0.1015
Sharpe Ratio	0.1127	-0.0914	0.4335	0.7393

Table 1: Performance Metrics by Component. Metrics are based on daily Yahoo Finance return data from 1950 to 2025, annualised using standard conventions. Data source: Yahoo Finance.

5.4 Strategy Performance

The trading strategy is implemented as a dynamic long-short allocation that scales exposure symmetrically around zero based on the composite signal. Positions may therefore take both positive and negative values, allowing the strategy to adjust between long and short market exposure in response to prevailing seasonal conditions. By construction, the strategy is designed to remain approximately market-neutral over time, with returns driven primarily by timing rather than sustained directional exposure. While a long-only implementation is conceptually feasible, all empirical results reported in this paper refer to the long-short specification. Given the long-short and approximately market-neutral nature of the strategy, the buy-and-hold benchmark is not intended as a direct point of comparison. Instead, it provides intuitive context for assessing the magnitude, volatility, and drawdown behaviour of the strategy. Accordingly, risk-adjusted measures such as the Sharpe ratio, drawdowns, and alpha estimates offer a more informative basis for evaluating economic performance.

Figure 7 shows that the strategy outperforms the benchmark early in the sample, with a steeper equity curve and a widening gap through the mid-1990s. After this period, the slope flattens and several prolonged drawdowns appear, while the benchmark continues to compound at a steadier pace. By the end of the sample, the S&P 500 closes most of the historical gap and slightly overtakes the strategy, indicating that the excess returns observed before 2000 do not persist in later decades.

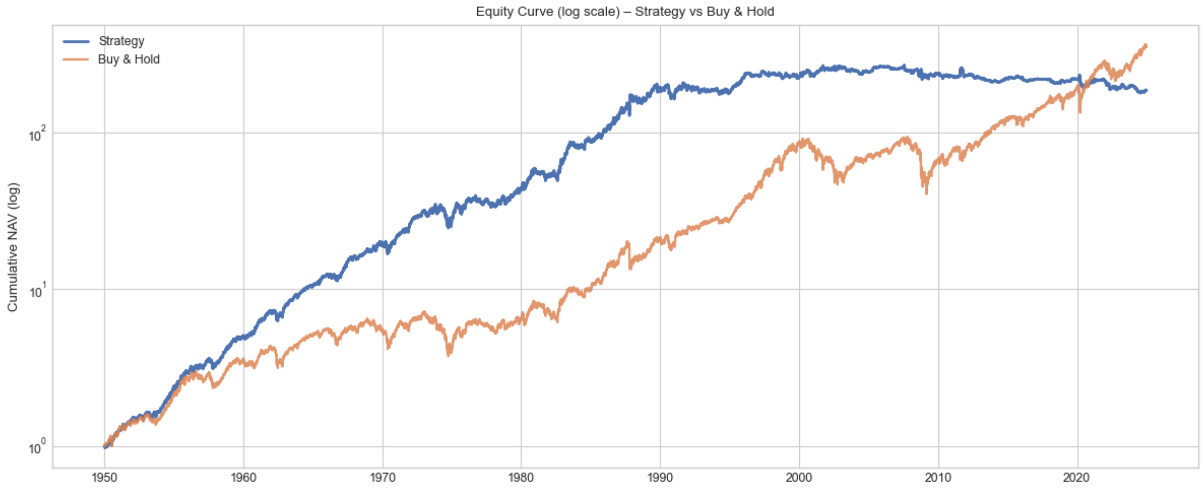

Figure 7: Log-Scaled Equity Curve for Strategy and Benchmark (1950 to 2025). The strategy applies the composite signal with symmetric position scaling between −1 and +1. Data source: Yahoo Finance.

Table 2 summarizes the key performance statistics. Although the strategy delivers a lower average return than the market (7.5 percent versus 9.1 percent annualized), its substantially lower volatility (10.2 percent versus 15.7 percent) results in a higher Sharpe ratio (0.74 versus 0.58). The return distribution also exhibits more favorable tail characteristics, with mildly positive skewness, lower excess kurtosis, and a maximum drawdown roughly half that of the benchmark. Hit rates are similar, indicating comparable frequency of gains but smoother return dynamics for the trading strategy.

Metric	Strategy	Buy & Hold
Annualized Return	0.0750	0.0908
Annualized Volatility	0.1015	0.1572
Sharpe Ratio	0.7393	0.5774
Positive Days	0.5133	0.5308
Positive Months	0.6000	0.6111
Skew	0.3166	-0.6275
Kurtosis	12.0566	20.0754
Max Drawdown	-0.3306	-0.5678
Number of Days	18870	18870

Table 2: Performance and Risk Metrics: Trading-Strategy vs. US Equities (1950 to 20205). Metrics are based on daily Yahoo Finance log returns and include annualised return, annualised volatility, Sharpe ratio, distributional measures, maximum drawdown, and sample size.

Additional diagnostics, including rolling volatility and rolling Sharpe ratios, are provided in the Appendix (Figure 11 and Figure 12). Together these measures show that while the strategy historically offered lower variance and stronger risk-adjusted performance, its advantage has weakened substantially in recent decades. A likely explanation is the diminishing influence of the announcement-related effects examined earlier. As markets have become more informationally efficient and event-driven return patterns have been arbitraged away, the underlying drivers of the strategy appear to have lost much of their historical strength. As a result, what was once a robust and persistent premium has evolved into a markedly weaker signal, yielding inferior risk-adjusted performance relative to passive equity exposure in recent years.

5.5 Alpha and Beta Exposure

Table 7, Table 8, and Figure 8 present the annual return decomposition of the strategy into beta and alpha components. The results reveal a clear structural break around 1990, with markedly different return characteristics before and after this period. In the first three decades of the sample (1960s–1980s), the strategy generated substantial alpha, averaging between 7.29 percent and 12.98 percent annually, with positive alpha in 22 of 27 years. Beta contributions during this period were comparatively modest, ranging from 1.56 to 3.75 percent, indicating that the strong historical performance was not primarily driven by market exposure but by genuine return predictability consistent with a long-short design. The strategy's low and stable market beta is not incidental but arises directly from its long-short construction and symmetric position scaling. By centering the signal around zero and constraining exposure within fixed bounds, the framework minimizes systematic market exposure, resulting in an approximately market-neutral profile where performance is driven by timing rather than directional market risk. In this context, achieving Sharpe ratios comparable to or exceeding those of the equity market is particularly remarkable.

The pattern reverses sharply from the 1990s onward. Alpha turns negative in four consecutive decades, averaging -2.04 percent in the 1990s and deteriorating further to -6.72 percent in the 2020s. Notably, only 9 out of 35 years since 1990 exhibit positive alpha. The beta contribution remains relatively stable at 3 to 4 percent, indicating that the strategy continues to provide some market exposure, but the value-added component has largely disappeared. Across the full sample, the strategy delivers an average annual return of 5.75 percent, composed of 2.55 percent from beta exposure and 3.09 percent from alpha, with exactly half of all years exhibiting positive alpha. However, this aggregate figure masks the pronounced time variation evident in the data. The historical alpha appears concentrated in the pre-1990 period, raising questions about whether the seasonal patterns that drove early outperformance have since been arbitraged away diminishing their effectiveness.

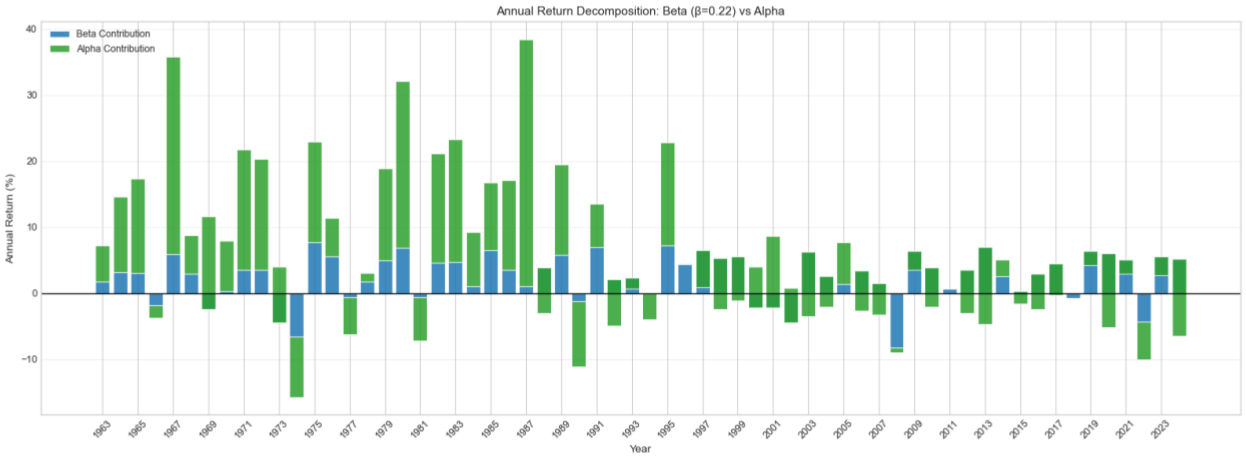

Figure 8: Alpha and Beta Contribution to Annual Returns (1950 to 2025). Beta returns are derived using the estimated market beta multiplied by annual S&P 500 excess returns; alpha represents the difference between total strategy return and beta return. Daily returns are aggregated to annual frequency. Data source: Yahoo Finance.

6. Fama French Factor Analysis

To assess the source of the strategy’s apparent alpha, a Fama–French factor regression is estimated. This allows the returns to be decomposed into exposures to established risk premia and a residual intercept that captures any remaining abnormal performance. A key objective of the factor regression analysis is to determine whether the strategy’s returns reflect compensation for systematic risk or arise from timing-based effects embedded in the trading rule. Given the long-short and signal-scaled construction of the strategy, exposure to the equity market is not intended to be a primary driver of performance. Consistent with this design, the estimated market beta is close to zero and statistically insignificant, indicating that the strategy is approximately market-neutral on average.

Table 3 presents the results of the six-factor model. Once market, size, value, profitability, investment, and momentum factors are included, the strategy’s alpha becomes statistically insignificant. The estimated market beta is effectively zero (0.0019) and insignificant, consistent with the design of the signal, which is centered around zero and clipped symmetrically, leading to a largely market-neutral exposure. This implies that returns are not compensation for bearing systematic risk but instead reflect timing-based effects associated with seasonal information. As shown in the Appendix (Table 9) the model explains only a modest share of return variation ($R^2 = 0.165$). The alpha estimate is slightly negative and insignificant, and the momentum loading is also insignificant. Only the value factor approaches significance, suggesting a mild growth tilt. Overall, the factor model should be interpreted with caution, as the generally weak factor loadings indicate substantial noise in the data.

Variable	Coef.	Std. Err.	z	P > z	[0.025]	[0.975]
const	-0.0020	0.003	-0.767	0.443	-0.007	0.003
Mkt_RF	0.0019	0.056	0.035	0.972	-0.109	0.112
SMB	-0.0055	0.085	-0.065	0.948	-0.172	0.161
HML	-0.1288	0.067	-1.911	0.056	-0.261	0.003
RMW	-0.1046	0.072	-1.453	0.146	-0.246	0.037
CMA	-0.0294	0.113	-0.259	0.795	-0.252	0.193
Mom	-0.0649	0.070	-0.934	0.350	-0.201	0.071

Table 3: Regression of Excess Returns on Fama–French Six Factors. Estimates are based on HAC-robust standard errors. Data sources: Yahoo Finance and Kenneth French Data Library.

7. Sensitivity and Robustness

To evaluate whether the strategy genuinely possesses predictive power and to reduce the risk of overfitting to a particular parameter configuration, the model was tested across a broad grid of alternative specifications. The resulting Sharpe ratio surface, shown in Figure 9, plots performance across combinations of signal half-lives, z-score half-lives, and trading thresholds. The figure reveals a smooth and well-behaved performance landscape, with Sharpe ratios tightly clustered between roughly 0.70 and 0.77 for the vast majority of parameter settings. This consistency indicates that neither the responsiveness of the signals nor the speed of standardization materially alters the qualitative outcome of the strategy.

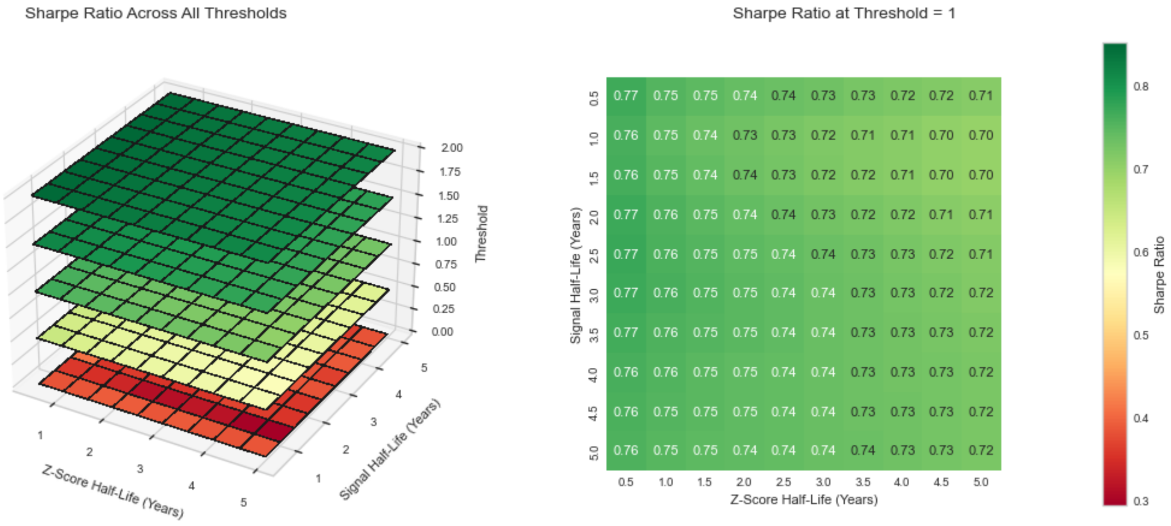
Sharpe Ratio Across All Thresholds

Sharpe Ratio at Threshold = 1

Signal Half-Life (Years) \ Z-Score Half-Life (Years)	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
0.5	0.77	0.75	0.75	0.74	0.74	0.73	0.73	0.72	0.72	0.71
1.0	0.76	0.75	0.74	0.73	0.73	0.72	0.71	0.71	0.70	0.70
1.5	0.76	0.75	0.74	0.74	0.73	0.72	0.72	0.71	0.70	0.70
2.0	0.77	0.76	0.75	0.74	0.74	0.73	0.72	0.72	0.71	0.71
2.5	0.77	0.76	0.75	0.75	0.74	0.74	0.73	0.73	0.72	0.71
3.0	0.77	0.76	0.75	0.75	0.74	0.74	0.73	0.73	0.72	0.72
3.5	0.77	0.76	0.75	0.75	0.74	0.74	0.73	0.73	0.73	0.72
4.0	0.76	0.76	0.75	0.75	0.74	0.74	0.73	0.73	0.73	0.72
4.5	0.76	0.75	0.75	0.75	0.74	0.74	0.73	0.73	0.73	0.72
5.0	0.76	0.75	0.75	0.74	0.74	0.74	0.74	0.73	0.73	0.72

Figure 9: Sharpe Ratio Sensitivity to Signal and Z-Score Half-Lives. The left panel shows the Sharpe ratio surface evaluated over all thresholds from 0 to 2. The right panel reports the Sharpe ratio for a fixed threshold of 1 across the grid of exponential-weighting half-lives. Performance metrics are based on daily Yahoo Finance returns over the full sample period. Data source: Yahoo Finance.

The figure also highlights the role of the trading threshold. As the threshold increases, the Sharpe ratio generally improves, reflecting the benefit of executing trades only when the composite signal reaches stronger conviction levels. This filters out noise and concentrates exposure in periods where the underlying indicators provide clearer directional guidance. However, the performance gains flatten once the threshold approaches about 1.5, suggesting that stricter entry criteria produce diminishing marginal benefits. Extremely high thresholds eventually reduce trade frequency and limit portfolio exposure, which decreases total return without meaningfully improving risk-adjusted performance.

Across the entire parameter grid, the strategy performs consistently across a wide range of smoothing speeds, standardization windows, and conviction thresholds. This stability indicates that the composite signal captures a persistent underlying structure rather than exploiting a narrow or fragile parameter choice, and it strengthens confidence that the observed predictive power is not the result of overfitting.

8. Limitations

A key limitation is that the study does not incorporate transaction costs or slippage. The analysis abstracts from real-world frictions such as trading fees, execution costs, financing rates, and short-sale constraints. The results are also based solely on U.S. equities and show clear regime dependence, which limits their general applicability across markets and time periods. Although best practices were followed, including avoiding look-ahead bias and conducting extensive parameter-sensitivity checks, the research design remains ex post. This raises the possibility of data-snooping bias, where part of the observed performance may reflect historical chance rather than persistent predictive power. Only a live implementation could fully address this concern. Since the strategy's performance advantage has weakened in recent years, such validation is no longer feasible, making it difficult to distinguish structural alpha from potential data-mining effects.

9. Discussion and Conclusion

A central finding of this paper is the pronounced time variation in the effectiveness of seasonality-based trading strategies. While the combined signal delivers strong risk-adjusted performance in earlier decades, particularly prior to the 1990s, this outperformance weakens substantially and largely disappears after 2000. Importantly, much of the historical performance is achieved with minimal systematic market exposure, as reflected in the strategy's near-zero beta and comparatively high Sharpe ratios during the early sample period. Individual components contribute unevenly to performance, with some signals exhibiting more persistent predictive power than others. The results further show that seasonal effects are measurable and cannot be fully explained by the timing of major information releases. Earnings-announcement intensity and monetary-policy timing instead capture distinct sources of return predictability that operate alongside broader calendar-based patterns. Although integrating these dimensions produces a composite signal with historical predictive power, its declining effectiveness in recent decades suggests that greater market efficiency and the widespread documentation of seasonal anomalies have led to their gradual arbitrage and erosion. This trend is consistent with the broader decline in the profitability of seasonal and announcement-related strategies as outlined in the literature review section.

Regarding Research Question 1, the analysis demonstrates that the constructed return-seasonality signal captures several well-established calendar effects. The monthly and daily profiles reveal patterns that correspond to the January effect, turn-of-the-month behavior, mid-year slowdowns, and stronger performance toward the end of the year. These features appear consistently in both the median series and the historical range of variation, which indicates that the measure captures genuine intra-year structure rather than noise, supporting H_{11}.

Research Question 2 asked whether return seasonality can be explained by concentrated information events, such as earnings-reporting intensity or FOMC announcement timing. The correlation and regression results show that these announcement-related variables account for only a small portion of the variation in the seasonality signal. While earnings seasons and monetary-policy announcements display distinct return patterns, they do not account for the broader seasonal structure observed in the data, leading to rejection of H_{02}.

Research Question 3 explored whether the three signals can be combined into a simple real-time trading rule. The composite signal outperforms the individual components on a risk-adjusted basis and exhibits smoother drawdowns, but its performance is strongly time-varying. Most of the outperformance occurs prior to 2000, with substantially weaker results thereafter. Because the analysis does not include transaction costs or slippage, the trading results should be viewed strictly as a proof of concept. Within this context, the evidence is sufficient to reject H_{03}.

10. References

Ball, Ray, and Philip Brown. 1968. 'An Empirical Evaluation of Accounting Income Numbers'. *Journal of Accounting Research* 6 (2): 159–78. https://doi.org/10.2307/2490232.

Board of Governors of the Federal Reserve System. 2020a. 'FOMC (a)'. https://www.federalreserve.gov/newsevents/pressreleases/monetary20200303a.htm.

Board of Governors of the Federal Reserve System. 2020b. 'FOMC (b)'. https://www.federalreserve.gov/newsevents/pressreleases/monetary20200315a.htm.

Enow, Samuel Tabot. 2023. 'Random Walk and Modelling Stock Return: Evidence from International Stock Markets'. *International Journal of Research in Business and Social Science (2147- 4478)* 12 (3): 353–60. https://doi.org/10.20525/ijrbs.v12i3.2443.

Fama, Eugene F. 1970. 'Efficient Capital Markets: A Review of Theory and Empirical Work'. *The Journal of Finance* 25 (2): 383–417. https://doi.org/10.2307/2325486.

'Federal Open Market Committee'. 2025. https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm.

Hu, Grace Xing, Jun Pan, Jiang Wang, and Haoxiang Zhu. 2022. 'Premium for Heightened Uncertainty: Explaining Pre-Announcement Market Returns'. *Journal of Financial Economics* 145 (3): 909–36. https://doi.org/10.1016/j.jfineco.2021.09.015.

Kettell, Laura, John McInnis, and Wuyang Zhao. 2022. *Why Has PEAD Declined Over Time? The Role of Earnings News Persistence*. 2022.

Kurov, Alexander, Marketa Wolfe, and Thomas Gilbert. 2020. 'The Disappearing Pre-FOMC Announcement Drift'. SSRN Scholarly Paper No. 3134546. Rochester, NY: Social Science Research Network, September 26. https://doi.org/10.2139/ssrn.3134546.

Lee, Suzanne S, and Minho Wang. 2025. 'Jumps and Post-FOMC Announcement Returns in Currency Markets'. *The Review of Asset Pricing Studies*, 10 June 2025, raaf003. https://doi.org/10.1093/rapstu/raaf003.

Lo, Andrew W. 2004. 'The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective'. SSRN Scholarly Paper No. 602222. Rochester, NY: Social Science Research Network, October 15. https://papers.ssrn.com/abstract=602222.

Lucca, David O., and Emanuel Moench. 2015. 'The Pre-FOMC Announcement Drift'. *The Journal of Finance* 70 (1): 329–71. https://doi.org/10.1111/jofi.12196.

Martineau, Charles. 2021. 'Rest in Peace Post-Earnings Announcement Drift'. SSRN Scholarly Paper No. 3111607. Rochester, NY: Social Science Research Network, March 27. https://doi.org/10.2139/ssrn.3111607.

Richardson, A. William, and Kevin Veenstra. 2022. 'The Post-Earnings Announcement Drift: A Pre-Earnings Announcement Effect? A Multi-Period Analysis'. *Abacus* 58 (4): 648–78. https://doi.org/10.1111/abac.12265.

Rozeff, Michael S., and William R. Kinney. 1976. ‘Capital Market Seasonality: The Case of Stock Returns’. *Journal of Financial Economics* 3 (4): 379–402. https://doi.org/10.1016/0304-405X(76)90028-3.

Schwert, G. William. 2002. ‘Anomalies and Market Efficiency’. SSRN Scholarly Paper No. 338080. Rochester, NY: Social Science Research Network, October 1. https://doi.org/10.2139/ssrn.338080.

Wachtel, Sidney B. 1942. ‘Certain Observations on Seasonal Movements in Stock Prices’. *The Journal of Business of the University of Chicago* 15 (2): 184–93.

11. Appendix

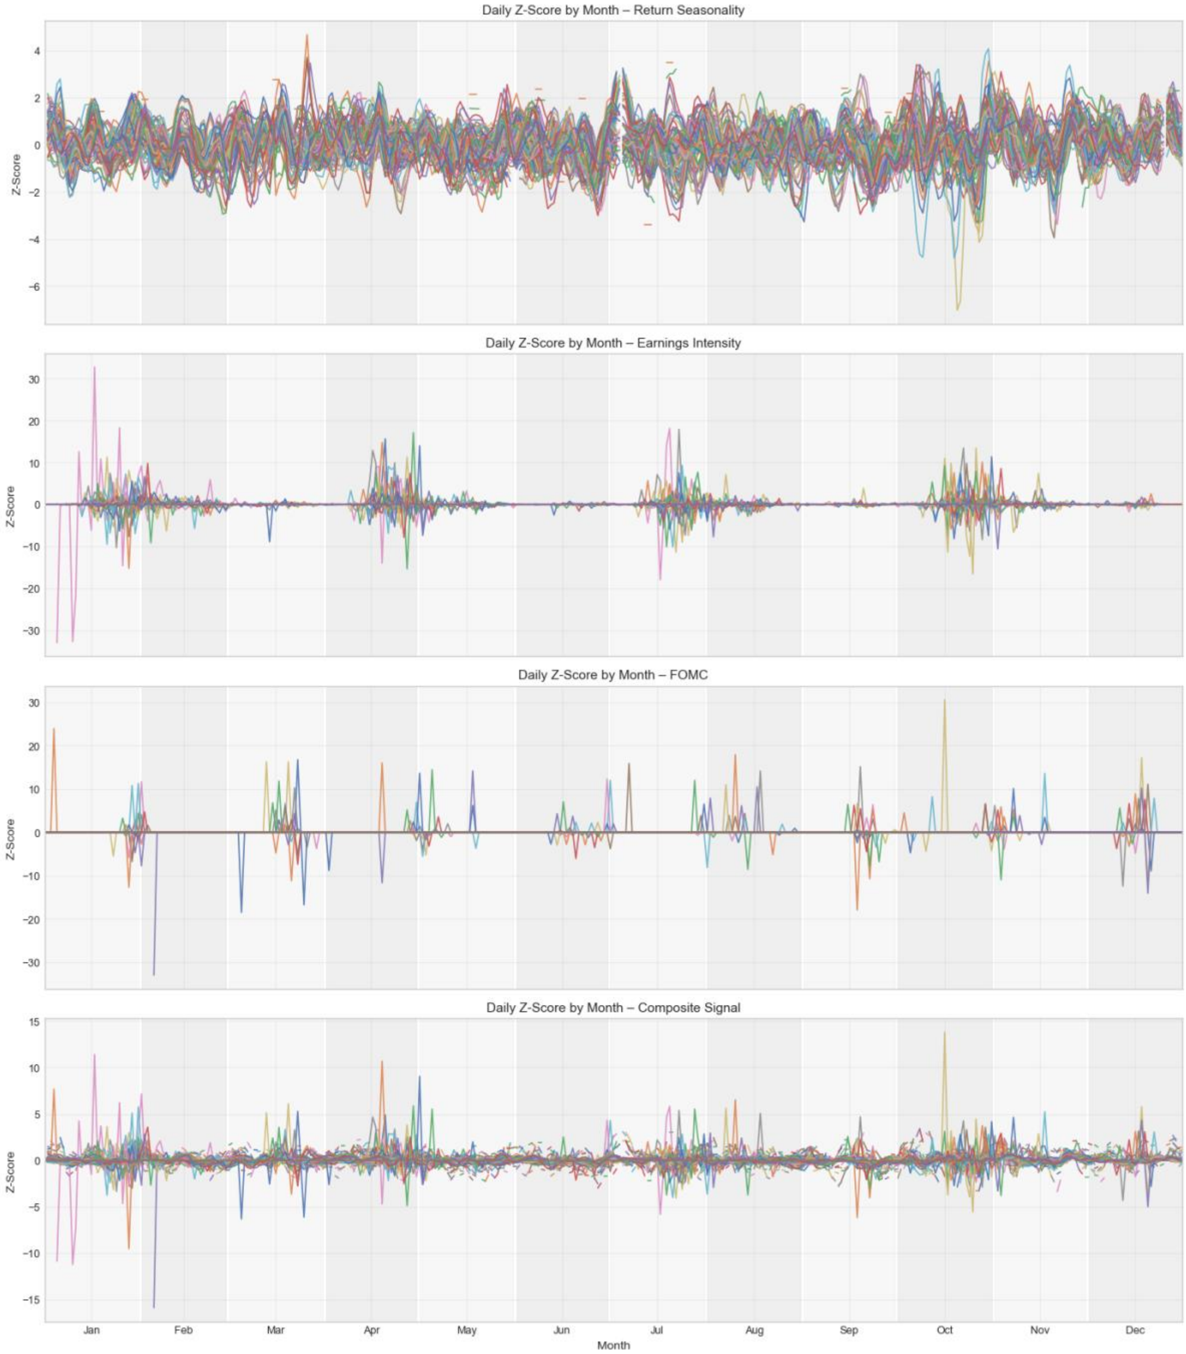

Figure 10: Daily Z-Scores of Component Signals and Composite Indicator by Month (1994 to 2025). Daily z-scores are shown for the return-seasonality, earnings-intensity, and FOMC timing components, along with the resulting composite signal, plotted by month across all years in the sample. Signals are standardised using exponentially weighted moving-average (EWMA) z-scores with three-year half-lives. Data sources: Yahoo Finance, IBES, and Federal Reserve FOMC calendar.

	Return-Seasonality	FOMC	Earnings	Combined
Return-Seasonality	1.0000	0.0032	0.0158	0.6380
FOMC	0.0032	1.0000	0.2489	0.3804
Earnings	0.0158	0.2489	1.0000	0.6365
Combined	0.6380	0.3804	0.6365	1.0000

Table 4: Correlation Matrix of Component and Composite Signal Returns (1994 to 2025). The table shows the correlation structure of daily returns from the seasonality, FOMC, earnings-intensity, and combined signals. Data sources: Yahoo Finance, IBES, Federal Reserve.

Statistic	Return Seasonality	Earnings Intensity	FOMC	Composite Signal
Mean	0.040	0.001	-0.000	0.029
Std. Dev.	0.831	1.337	1.038	0.715
Minimum	-7.024	-33.010	-33.010	-15.892
Median	0.043	-0.025	-0.024	-0.004
Maximum	4.676	32.738	30.605	13.799
Skewness	-0.118	-0.634	3.504	0.281
Kurtosis	1.274	154.428	291.551	53.830
% Positive	52.300%	11.000%	11.600%	28.400%

Table 5: Summary Statistics for Component and Composite Signals (1994–2025). Descriptive measures are computed from daily signal values using EWMA transformations for smoothing and standardisation. Data sources: Yahoo Finance, IBES, Federal Reserve.

Parameter	Symbol	Default Value	Description
Signal half-life	λ_signal	3 years	EWMA decay for seasonal smoothing
Z-score half-life	λ_zscore	3 years	EWMA decay for standardization
Position threshold	θ	1.0	Signal magnitude for full position

Table 6: Overview of Model Parameters Used in Signal Construction. Inital parameter choices for smoothing the raw signals, computing rolling z-scores, and determining when the trading rule takes a full long or short position.

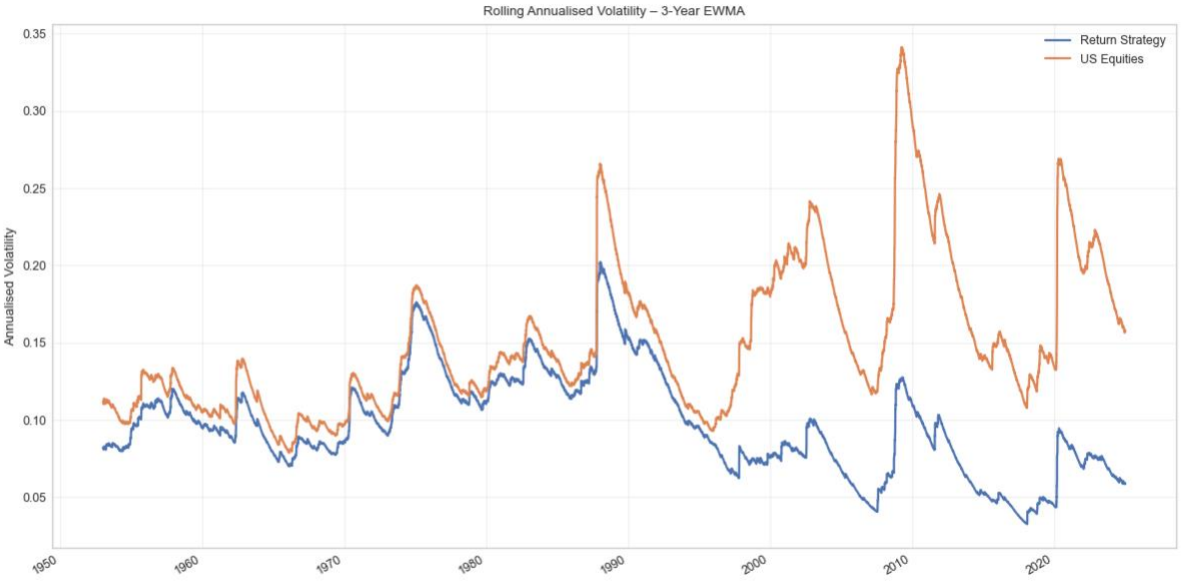

Figure 11: Rolling 3-Year Annualised Volatility (EWMA) of Trading-Strategy vs. S&P 500 (1950 to 2025). Annualised volatility is derived from daily log returns using a three-year exponentially weighted moving-average window. Data source: Yahoo Finance.

Figure 12: Rolling 10-Year Annualised Sharpe Ratio: Trading-Strategy vs. US Equities (1950 to 2025). Sharpe ratios are computed using daily excess returns and annualised with standard conventions. Data source: Yahoo Finance.

Decade	Avg Total (%)	Avg Beta (%)	Avg Alpha (%)	Positive Alpha Years
1960	13.38	1.83	11.27	6/7
1970	9.00	1.56	7.29	8/10
1980	17.01	3.75	12.98	8/10
1990	1.90	3.92	-2.04	2/10
2000	0.51	0.44	-0.03	4/10
2010	-0.53	3.10	-3.59	3/10
2020	-3.22	3.52	-6.72	0/5

Table 7: Decade-Level Decomposition of Annual Returns Into Beta and Alpha Components. Average annual return components and counts of positive-alpha years by decade, computed using Yahoo Finance return data.

Metric	Value
Total	62
Positive Alpha	31 (50.0%)
Total Return (Annual Ø)	5.75%
Beta Contribution (Annual Ø)	2.55%
Alpha Contribution (Annual Ø)	3.09%

Table 8: Summary of Annual Return Decomposition Over the Full Sample. The table presents total sample size, frequency of positive-alpha years, and average annual total, beta, and alpha returns, computed from annual Yahoo Finance data.

Statistic	Value
Dependent Variable	excess_ret_m
R-squared	0.165
Adjusted R-squared	0.051
F-statistic	2.516
Prob (F-statistic)	0.0351
Method	Least Squares
No. Observations	51
Df Residuals	44
Df Model	6
Covariance Type	HAC
AIC	–254.9
BIC	–241.4
Log-Likelihood	134.45
Durbin-Watson	1.906

Table 9: OLS Regression Results (FF6 Model). The table reports summary statistics and diagnostic measures for the regression of monthly excess returns on the six Fama–French factors, estimated using HAC standard errors. Data sources: Yahoo Finance and Kenneth French Data Library.

Test	Statistic	Prob.
Omnibus	1.708	0.426
Jarque-Bera (JB)	1.450	0.484
Skew	–0.256	—
Kurtosis	2.351	—
Condition No.	63.4	—

Table 10: Residual Diagnostics for the Fama–French Regression Model. The table reports normality tests, skewness, kurtosis, and the condition number based on the residuals from the six-factor regression estimated with HAC standard errors. Data sources: Yahoo Finance and Kenneth French Data Library.

Statistic	Value
Number of Observations	10,501
Degrees of Freedom (Model)	2
Degrees of Freedom (Residuals)	10,498
R-squared	0.0004
Adjusted R-squared	0.0002
F-statistic	2.330
Prob (F-statistic)	0.0974
Log-Likelihood	-12,982
AIC	25,970
BIC	25,990

Table 11: Model Fit Statistics for the Return Seasonality Regression Model. The table reports key goodness-of-fit measures, test statistics, and information criteria for the OLS regression of return seasonality on FOMC and earnings indicators estimated with HAC standard errors. Source: Yahoo Finance, IBES, Federal Reserve.

Variable	Coefficient	Std. Error	z-Statistic	p-Value	95% Confidence Interval
Constant	0.0034	0.015	0.228	0.820	[-0.025 , 0.032]
FOMC	0.0060	0.009	0.695	0.487	[-0.011 , 0.023]
Earnings	0.0112	0.006	1.897	0.058	[-0.000 , 0.023]

Table 12: Coefficient Estimates for the Return Seasonality Regression Model. The table reports coefficient estimates, HAC standard errors, z-statistics, p-values, and 95% confidence intervals for the OLS regression of return seasonality on FOMC and earnings indicators. Source: Yahoo Finance, IBES, Federal Reserve.