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: Calendar Effects and Market Events

Guilherme de Lucena Sampaio Ramos Borges

Work project carried out under the supervision of:

Afonso Januário

23/01/2026

Abstract

This thesis studied whether several documented calendar and event-driven anomalies can be combined into a practical, low-frequency trading strategy for the U.S. equity market. Using daily S&P 500 data from 1975 to 2024, the analysis evaluates the turn-of-the-month effect, pre-holiday returns, Federal Reserve meeting dates, earnings intensity and seasonal expected returns. Each effect is tested individually and then integrated into a unified strategy. While individual anomalies exhibit modest and time-varying performance, their combination concentrates returns into specific trading windows and improves risk-adjusted performance relative to buy-and-hold, showing the economic importance of predictable timing patterns in equity returns.

Keywords: Calendar Effects, Seasonal Momentum, Earnings Announcements, FOMC Meetings

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

Financial markets are often described as efficient processors of information, yet some of the earliest empirical studies in finance revealed patterns that contradict this view. One of the most striking and historically important examples is the January effect. First noticed informally in the 1940s and later documented rigorously by Rozeff and Kinney (1976), the January effect refers to the tendency for stock returns, especially among smaller firms, to be unusually high in the first trading days of the new year. Its discovery was surprising because it implied that investors could earn excess risk-adjusted returns simply by holding stocks during certain days within a year, even though no new fundamental information was being released.

Following the discovery of the January effect, a large empirical literature emerged documenting return patterns linked to the calendar. Researchers identified anomalies operating at different frequencies, including within-month effects, intra-year seasonality, and longer-horizon cycles. Among the most prominent findings are the turn-of-the-month effect, whereby stock returns tend to be unusually high around month boundaries, and the Halloween effect, which suggests that a disproportionate share of equity returns is earned between November and April. Other studies report abnormal returns associated with specific trading days, weeks, or holidays. Collectively, these findings challenged the notion that expected returns are constant through time and instead suggested that market behavior may be shaped by recurring institutional, behavioral, and liquidity-driven cycles.

Since these effects happen at predictable times without corresponding increases in measured risk, they can be hard to combine with conventional asset-pricing models. However, they are still evident in long-term historical data, suggesting that systematic patterns in investor attention, institutional trading activity, and information processing timing may be reflected in market behavior.

The purpose of this project is to investigate whether combining several such anomalies can produce a practical and transparent trading strategy that is capable of improving performance relative to the S&P 500 buy-and-hold strategy. Therefore, the goal will not be to identify new anomalies but demonstrate rather on how overlapping signals from independent factors may improve risk-adjusted performance. The analysis focuses on a set of calendar and event-oriented effects: the turn of the month, pre-holiday behavior, Federal Reserve announcement drift, and market-wide earnings-intensity effects. Each component captures a different mechanism underlying return predictability, and because these mechanisms are largely independent, overlaying them may amplify their combined effectiveness. Using daily data and simple rules, the thesis evaluates whether this multi-component approach yields economically meaningful improvements in performance when compared with buy-and-hold.

While no single anomaly is strong enough to dominate market returns on its own, the existence of multiple, distinct sources of return predictability raises the possibility that their combination may return a more stable performance. This perspective suggests that markets exhibit structured return patterns tied to time and information cycles, and that even simple rules may be able to capture these effects when applied with discipline.

2. Literature Review

One of the earliest and most influential discoveries in the study of return seasonality is the January effect, which refers to the disproportionately high stock returns observed at the beginning of the calendar year. The research indicated that January returns were significantly greater compared to other months, with this trend being especially noticeable among smaller companies. Subsequent research by Keim (1983) linked the anomaly to tax-loss selling and institutional portfolio rebalancing, while international evidence documented similar patterns across many developed and emerging markets. Later studies, however, emphasize that the

January effect is not stable through time. Using a dynamic framework, Marquering et al. (2006) show that the January effect weakened sharply following its initial academic publication and largely disappeared in subsequent decades, with only occasional and short-lived reappearances. According to their findings, calendar anomalies might have a self-destructive quality in which investor awareness and arbitrage activity lessen their economic significance once they are extensively documented. The outcome matches more extensive survey data that suggests many well-known anomalies reduce rather than instantly disappear after discovery (Schwert 2003). Despite this attenuation, the January effect remains historically central to literature because it provided the first systematic evidence that expected returns vary predictably over the calendar, motivating extensive research into seasonal and timing-based return patterns.

In addition to the January effect, researchers have identified several other calendar-driven anomalies and one of the most robust is the pre-holiday effect, documented thoroughly by Ariel (1987; 1990). Ariel demonstrated using several decades' worth of U.S. equity data that the trading days right before market holidays produce returns that are significantly higher than average daily returns and frequently account for a disproportionate portion of long-term market performance. Surprisingly, this pattern holds true for both equal-weighted and value-weighted indices, and it cannot be accounted for by other seasonal effects like the January or weekend effect. Ariel demonstrates that pre-holiday excess returns are not concentrated at the market open or close but rather are earned during the trading day before holidays as opposed to overnight. This result implies that the effect reflects broader behavioral or liquidity-related forces and challenges microstructure-based explanations. Although theories like thin trading, decreased risk aversion, and investor optimism have been put forth, no clear mechanism has been found, and the effect's economic impact varies over sample periods.

Another major seasonal pattern is the turn-of-the-month (TOM) effect, noted by Lakonishok and Smidt (1988) and later investigated in depth by Xu and McConnell (2006). Their study, which spans more than a century of U.S. data, demonstrates that nearly the entire equity premium is earned during a narrow four-day interval comprising the last trading day of each month and the first three trading days of the next. Returns in the remainder of the month are close to zero or negative. The TOM effect is remarkably stable across time, surviving the post-publication era and appearing in 30 out of 34 international markets examined. It also cannot be explained by standard risk factors, firm size, liquidity, interest-rate cycles, or institutional flows such as mutual fund rebalancing. The persistence and universality of the TOM effect make it one of the most puzzling findings in empirical finance, often interpreted as arising from systematic patterns in investor cash flows or institutional trading cycles, though no single explanation has gained consensus.

Seasonal patterns are complemented by event-based return anomalies, which arise not from the calendar itself but from predictable information cycles. The post-earnings-announcement-drift (PEAD), which Bernard and Thomas (1989) documented explains this effect. PEAD describes the propensity of stock prices to move in the direction of an earnings surprise for weeks or months after the announcement, as opposed to the rapid adjustment that efficient markets predict.. Later studies link post-earnings drift to information uncertainty and delayed investor response (Francis et al. 2007). While the existing literature focuses almost exclusively on firm-level reactions to earnings announcements, the underlying mechanisms it identifies, such as information uncertainty, investor inattention, and gradual information dissemination, raise the possibility that clustered earnings releases could also have market-wide effects. This hypothesis has received limited direct attention at the index level and therefore motivates part of the empirical analysis conducted in this thesis.

A parallel and increasingly influential strand of literature examines stock return behavior around monetary policy announcements, especially Federal Open Market Committee (FOMC) meetings. Lucca and Moench (2015) document the striking “pre-FOMC announcement drift,” showing that U.S. equities earn abnormally high returns in the 24 hours preceding scheduled FOMC announcements. This effect is so large and consistent that it accounts for a substantial portion of the total equity premium over several decades. Notably, the drift is concentrated in a very narrow time window and does not reverse after the announcement, suggesting that conventional risk-based explanations are insufficient. More recent work by Cieslak et al. (2019) places the pre-announcement drift within a broader monetary policy cycle. They show that equity returns are systematically concentrated in specific weeks of the FOMC cycle, with positive excess returns occurring primarily in even-numbered weeks following scheduled meetings. This evidence suggests that the FOMC effect is not confined to a narrow intraday window but reflects a broader pattern of return timing linked to the structure of monetary policy communication and investor attention over the policy cycle

Taken together, the literature demonstrates that equity returns show both seasonal structure and foreseeable responses to scheduled information events. Calendar effects such as the January, pre-holiday, and turn-of-the-month patterns reveal systematic variation tied to the flow of time and institutional cycles. Anomalies such as earnings-related drift and pre-FOMC behavior highlight the role of information processing, risk premia dynamics, and behavioral responses. Although each anomaly is small in isolation and often confined to a narrow set of trading days, their persistence across markets and decades suggests that they reflect distinct underlying mechanisms and capture independent sources of foreseeable return patterns.

It will therefore be studied if a series of anticipated calendar and event anomalies, documented largely in isolation in the literature, can be combined into a simple, low-frequency trading strategy that improves risk-adjusted performance relative to buy-and-hold investment in the S&P 500.

3. Methodology

The empirical analysis is conducted using daily data for the U.S. equity market over the period 1975–2024. Market returns are measured using the S&P 500 index, constructed from daily close-to-close returns of the ^GSPC series obtained from Yahoo Finance accessed via the *yfinance* Python library. This long sample period allows the analysis to span multiple monetary policy regimes, business cycles, and structural changes in market microstructure, which is particularly important when studying seasonal and event-driven return patterns.

All auxiliary datasets are accessed through WRDS. Earnings announcement data was extracted from the I/B/E/S database, with firm-level earnings announcement dates for U.S. equities. Firm-level market capitalization data are sourced from CRSP and are used to aggregate earnings announcements to the market level. FOMC meeting dates are collected from the Federal Reserve's official policy meeting calendar, which lists scheduled Federal Open Market Committee meetings over the sample period.

Classical studies of the post-earnings-announcement drift (PEAD), beginning with Bernard and Thomas (1989) and extended by Francis et al. (2007), examine the cross-sectional response of individual stocks to their own earnings surprises over relatively short post-announcement horizons. These studies document systematic underreaction to earnings information at the firm level, which manifests as return continuation following the announcement.

Rather than analyzing firm-level drift, a market-level perspective will be adopted on earnings-related information flow. Using earnings announcement dates from I/B/E/S and firm-level market capitalization from CRSP, a daily measure of earnings intensity is constructed. Earnings intensity is defined as the fraction of total U.S. equity market capitalization reporting earnings on a given trading day. This measure captures the degree to which aggregate corporate information is released simultaneously, instead of the magnitude of individual earnings surprises and the higher this value, the higher the share of the market reporting their earnings on that specific day.

Scheduled FOMC meetings represent a second source of predictable information arrival. Prior research documents abnormal equity returns around these announcements, most notably the pre-announcement drift identified by Lucca and Moench (2015), who use minute-level S&P 500 futures data to show that excess returns accumulate primarily in the 24 hours preceding the scheduled announcement time. More generally, Cieslak et al. (2019) show that equity returns vary systematically over the FOMC cycle, indicating that monetary policy communication affects return timing beyond a single intraday window.

In contrast to these high-frequency approaches, FOMC-related return predictability will be evaluated at daily frequency. The FOMC signal is constructed as a binary indicator that takes the value one on the calendar date of each scheduled meeting and remains zero otherwise. Trading strategies enter a long position in the S&P 500 on the meeting date and hold for a fixed number of trading days. This design deliberately abstracts from intraday execution timing and instead asks whether monetary-policy-related return regularities are still relevant when evaluated using close-to-close daily index returns, reflecting the perspective of a day-to-day investor.

In addition to these signals, the analysis joins seasonal effects documented in the literature, including the turn-of-the-month effect and the pre-holiday effect. These signals are constructed using simple, transparent calendar rules consistent with prior studies. The turn-of-the-month indicator captures the final trading day of each month and the first three trading days of the subsequent month, while the pre-holiday indicator identifies trading days immediately preceding market holidays. These effects are included to capture predictable institutional and behavioral patterns in return timing that are independent of scheduled information releases.

All signals are expressed as daily binary indicators applied to the same underlying asset, the

S&P 500 index. This unified framework allows seasonal and event-driven effects to be evaluated on a comparable basis and combined within a single trading strategy. By restricting attention to daily data and index-level positions, the methodology prioritizes transparency, replicability, and practical implementation over sharp intraday identification.

4. Individual Effects: Empirical Evidence

Before constructing a multi-component strategy, each seasonal and event-driven effect is examined independently. This step serves two purposes. First, it verifies whether anomalies documented in the literature remain detectable even when applied to daily index-level. Second, it establishes whether each effect provides sufficient economic and statistical justification to be included in a combined framework. Given that many prior studies rely on intraday or cross-sectional data, it is not obvious that these effects should survive aggregation to daily S&P 500 returns.

4.1. Turn-of-the-Month Effect

The turn-of-the-month effect is one of the most extensively studied calendar anomalies in the finance literature. Lakonishok and Smidt (1988) and, more recently, Xu and McConnell (2006) show that a disproportionate share of long-run equity returns is earned in a narrow window around month boundaries, typically spanning the final trading day of the month and the first three trading days of the following month.

The turn-of-the-month signal will therefore follow this standard definition: the strategy holds the S&P 500 only on the last trading day of each month and the first three trading days of the subsequent month, and remains out of the market otherwise. Applied to daily S&P 500 returns from 1975–2024, this rule produces a compound annual growth rate (CAGR) of approximately

4.23% with a Sharpe ratio of 0.60, while being invested about 48 trading days per year, provided in Table 1 (Appendix). Relative to a buy-and-hold benchmark CAGR ≈ 9.25%, Sharpe ≈ 0.60, invested essentially every trading day possible, the turn-of-the-month strategy achieves roughly 45% of the total benchmark CAGR while participating on only about 19% of trading days as shown in Table 2 (Appendix). This comparison is not meant to imply that returns are mechanically "captured" in a proportional sense, but it does highlight the concentration of positive expected returns into a small subset of days, consistent with the core claim of the turn-of-the-month literature.

One more note may be important to report regarding this anomaly: Prior to 2000, the turn-of-the-month strategy generates an annualized return of approximately 6.36% (Appendix, Table 3.1), whereas post-2000 performance declines to around 2.13% per year (Appendix, Table 3.2). While the effect is clearly weaker in the more recent sample, it remains positive albeit less. This decline matches the evidence presented by Marquering et al. (2006), who show that several calendar anomalies diminish following academic dissemination. Overall, the results confirm that the turn-of-the-month effect persists at the index level, with reduced economic magnitude in recent decades.

4.2. Earnings-Intensity Effect

The literature on post-earnings-announcement drift (PEAD), beginning with Bernard and Thomas (1989), documents that stock prices tend to adjust slowly to earnings information, leading to predictable return continuation following earnings surprises.

The main goal here is to build on these ideas by examining whether clusters of earnings announcements generate systematic return patterns at the index level. Instead of conditioning on individual earnings surprises, the analysis constructs a daily measure of Earnings Intensity

(EI), defined as the share of total U.S. equity market capitalization reporting earnings on a given day. For each trading day, the market capitalization of firms announcing earnings is aggregated and divided by the total CRSP market capitalization of U.S. equities, producing a value-weighted measure of earnings concentration, shown in Equation 1:

$$EI_t = \frac{\sum_{i \in U_t} \text{MarketCap}_{i,t}}{\sum_{i \in U} \text{MarketCap}_{i,t}}$$

Where to calculate Earnings Intensity (*EI*), for any given day(t), i represents a specific firm and U constitutes all listed companies from CRSP United States database.

This approach ensures that earnings activity by large firms receives greater weight than that of small firms, reflecting the fact that announcements by economically significant companies are more likely to affect aggregate market pricing. The resulting earnings-intensity measure varies substantially across the earnings season, with particularly high values observed during peak reporting weeks.

Empirically, the strategy enters a long position on the day which earnings intensity exceeds a fixed threshold which was set at 0.5% of the market reporting in any given day so as to capture the activity of large firms and so that it does not signal a trade whenever any company in the United States reports their earnings. When applied to daily S&P 500 returns, this rule generates a CAGR of approximately 4.07%, with a Sharpe ratio slightly above 0.40 (Appendix, Table 2), while being invested on roughly 93 days per year.

The subperiod analysis indicates that the effect is present both before and after 2000, with similar magnitudes across regimes and slightly higher CAGR but also days invested per year post 2000 (Appendix, Tables 3.1; 3.2)

Taken together, these findings suggest that a significant share of market returns is earned during periods of concentrated earnings activity. By shifting the focus from individual firm reactions to market-wide earnings intensity, this analysis provides evidence that the gradual assimilation of corporate earnings information operates not only at the micro level but also at the level of broad equity indices. As such, the earnings-intensity effect represents a distinct and valid source of systematic differences in expected returns that complement traditional calendar-based anomalies.

As it has been stated, literature argues that markets react positively to positive earnings events and in this scenario and while that remains the case, there is a correlation, in these specific conditions, between information being released and markets reacting positively.

4.3. Seasonal Expected Return

Unlike anomalies such as the January effect, turn-of-the-month, or pre-holiday, literature does not document strong day-of-the-year patterns. However, related research suggests that calendar time can still structure expected returns in systematic ways. For example, Heston and Sadka (2008) demonstrate that stock returns in the cross-section exhibit persistent calendar-month seasonality, with firms tending to earn higher returns in the same calendar months across years. Although this evidence is cross-sectional and monthly rather than index-level and daily, it motivates the broader hypothesis that expected returns may vary across calendar dates in subtle but persistent ways.

Building on this, the present thesis explores whether weak calendar-date regularities can be detected at the index level using historical daily market data. Instead of imposing fixed seasonal effects, the analysis constructs a moving average estimate of expected returns conditional on the day of the year. This approach allows seasonal patterns to evolve gradually over time and accommodates the possibility that any these effects can potentially erode, shift, or disappear across regimes.

It is crucial to mention that the construction of the seasonal expected return signal is fully backward-looking. For each trading day, expected returns are estimated using only historical observations of the same calendar date from prior years, with greater weight placed on more recent data through exponential decay. No information from future periods is used in the estimation of seasonal patterns, ensuring that the strategy does not rely on any form of look-ahead bias.

This design choice is essential from both a methodological and economic perspective. Calendar-based strategies are particularly vulnerable to inadvertent forward-looking bias, as seasonal averages computed using the full sample can mechanically embed information from future returns. By restricting the seasonal estimate to data available at each point in time, the analysis evaluates whether historical calendar regularities would have been observable and actionable to an investor in real time.

With this in mind, the strategy takes long positions only on calendar dates for which the recency-weighted expected return exceeds a fixed threshold. This threshold was added because as the S&P 500 Index has increased severalfold in the past 50 years, the strategy would simply go long on the large majority of days. Thus, a threshold of 0.2% was introduced as it was a good equilibrium between a reasonable mean and meaningful results. When applied to daily S&P 500 returns, this rule generates a compound annual growth rate of approximately 3.28%, with a Sharpe ratio of 0.41, while being invested on roughly 64 trading days per year present in Figure 2(Appendix).

Once more, analysis indicates that the seasonal expected return effect slightly weakens in the post-2000 period, with lower Sharpe ratios (Appendix, Figure 3.1;3.2) and reduced statistical significance.

4.4. Pre-Holiday Effect

The pre-holiday effect, documented by Ariel (1990), refers to the propensity for unusually high stock returns on trading days before market holidays.. This effect is notable because it is not explained by news releases, risk premia, or standard calendar effects. Ariel's intraday analysis shows that returns accumulate steadily throughout the trading day, pointing toward behavioral or liquidity-based explanations.

Applying the pre-holiday rule where we invest on the day before holiday but contrary to what we have talked about, we will hold it for 2 days. By extending the holding period beyond the holiday we test whether this effect spills over into subsequent sessions. Applying this idea with our daily S&P 500 data yields a CAGR of approximately 1.49%, with a Sharpe ratio of 0.41 and being statistically significant. Over the full sample, the strategy is invested on approximately 17 trading days per year, which is somewhat low but makes sense give the number of U.S. market holidays (Appendix, Table 2). In this scenario, the effect from an economic point of view is not drastic, but given the number of days invested, it is a fair result.

Subperiod analysis does not show a clear decline after 2000, in contrast to previous single-day implementations. Despite very little market exposure, strategy-level performance is consistent across subsamples, with annualized returns of roughly 1.37% prior to 2000 and 1.60% after it as shown in both Table 3.1 and 3.2 (Appendix).

4.5. FOMC Announcement Effect

The monetary-policy event component is based on scheduled Federal Open Market Committee (FOMC) meetings. The academic literature documents abnormal equity returns around these announcements, most famously the "pre-FOMC announcement drift" described by Lucca and Moench (2015), which is largely concentrated in a narrow 24-hour window preceding scheduled meetings and is typically identified using intraday S&P 500 futures data. Because this thesis operates only at daily frequency, it does not attempt to isolate the precise intraday drift window nor to replicate what was done in that specific paper. Instead, it tries to assess whether a practically implementable, meeting-date rule retains serious predictable return behavior when translated into close-to-close index returns.

The FOMC calendar was built to include only meeting dates, eliminating non-policy dates that would otherwise dilute the effect by introducing noise. Under this specification, the strategy trades approximately 8.4 days per year, which matches the typical frequency of scheduled FOMC meetings.

Two implementations are considered. The first enters a long position on the meeting day and holds for one trading day (hold = 1). On table 4 (Appendix) this specification produces an annualized CAGR of around 1.13%, with a Sharpe ratio of 0.34. This effect is not only statistically significant, but it shows crucial information: As with the Earnings-Intensity effect, when information is released, markets on average tend to go up and, in this case, its standalone CAGR is about 12% of the buy-and-hold CAGR, despite being invested ~8.4 days/year.

The second specification is instead holding the position for two trading days, entering on the meeting day and exiting after the second day. This slightly increases exposure to approximately 16.9 days per year and improves overall performance to a CAGR of approximately 1.67% with a Sharpe ratio of 0.39(Appendix, Table 2). Conditional returns remain positive and statistically significant. This two-day variant is therefore adopted as the baseline FOMC component: it remains easy to implement and interpret, is in line with the idea that policy information may be incorporated over more than a single close-to-close return and provides stronger risk-adjusted performance per unit of market exposure than longer holding rules. Once more, the effect is quite limited but it is understandable given the number of days invested.

When looking at this last strategy, results suggest that the effect is present both before and after

2000, although statistical strength and returns are slightly weaker in the post-2000 part (Appendix, Table 3.1; 3.2).

As it was with the Earnings Intensity strategy, there seems to be a positive correlation once more where markets tend to react positively to information being released. While this is not proven here, investors may, as market conditions become more clear, be more willing to invest in the stock market, therefore placing pressure on prices and leading it to ultimately increase which we have observed with the results.

5. Combined Strategies

In assessing whether calendar-based and event-related indicators provide useful information beyond passive market exposure, several combined specifications are evaluated against a buy-and-hold investment in the S&P 500. Buy-and-hold serves as the natural benchmark, delivering a compound annual growth rate of approximately 9.25% with a Sharpe ratio close to 0.60 (Appendix, Table 1) while remaining invested on essentially all trading days. This benchmark is important not only because it represents the standard alternative available to investors, but also because it sets a high bar: any rules-based approach that spends meaningful time out of the market must justify itself either by improving risk-adjusted performance, by concentrating returns into a smaller subset of trading days, or by achieving both.

Before constructing combined strategies, it is important to distinguish between two conceptually different approaches to capturing seasonality in returns. The turn-of-the-month and pre-holiday effects represent discrete calendar events that occur at clearly defined and infrequent points in time. These effects have been linked in the literature to institutional cash flows, portfolio rebalancing, liquidity conditions, and behavioral patterns, and they concentrate expected returns into a small number of trading days each year. As such, they naturally function as high-conviction signals that identify narrow windows of elevated expected returns.

In contrast, the seasonal expected return strategy based on day-of-the-year averages is inherently more complex. By construction, it identifies calendar dates that have exhibited higher average returns historically, but it does not correspond to a specific economic event or institutional mechanism occurring on those days. It should, in theory, capture some of the trading days that the pre-holiday effect and TOM effect also do as well as some others. As a result, the seasonal expected return signal tends to activate more frequently and increases overall market exposure, rather than sharply isolating high-return periods. And while it is interesting to demonstrate that calendar-date regularities persist at the index level, its economic interpretation is less direct than that of discrete calendar effects. Due to this, the turn-of-the-month and pre-holiday effects are treated as the primary seasonal components in the combined strategy, while the seasonal expected return signal is considered separately and its effect is only studied in isolation.

5.1. First Strategy

The first combined specification strategy four signals, turn-of-the-month behavior, pre-holiday effects, earnings intensity, and Federal Reserve meeting dates, such that the market is entered whenever at least one signal is active. According to the results on Table 5.1 (Appendix) under this rule, the strategy is invested on approximately 145 trading days per year, substantially fewer than buy-and-hold. Despite this reduced exposure, the strategy delivers a CAGR of 8.84% with a Sharpe ratio of 0.70, exceeding the benchmark on a risk-adjusted basis only with a Sharpe Ratio of around 0.10 higher. This result indicates that a large fraction of long-run market returns is earned during trading days associated with anticipated calendar and information events, while exposure on lower-expected-return days can be avoided without materially reducing performance. Importantly, we go long nearly 100 days less than the benchmark strategy and obtain a CAGR that is only 0.41% lower.

5.2. Second Strategy: Increasing exposure

Allowing exposure to increase on days in which multiple signals coincide further strengthens this concentration effect. To prove this Table 6.1 in the Appendix reports the average number of trading days per year on which each signal is active, together with their pairwise overlap. While earnings-related signals are active on roughly 93 days per year and the turn-of-the-month effect on approximately 48 days, other signals such as FOMC meetings and pre-holiday effects are far more infrequent. What is important to note here is that overlapping between signals is quite limited and as such it could be assumed that overlapping suggests that the signals capture distinct sources of return. So, increasing exposure when multiple signals coincide concentrates capital on a small subset of days which on average are 28.6 days per year.

Furthermore, to assess whether the individual signals capture distinct sources of return predictability, Table 6.2 (Appendix) reports pairwise correlations between strategy returns. Correlations between signals are all lower 0.25 proving that there are no major correlation between the effects. These low correlations suggest that the strategies are driven by different reasons and that combining them provides diversification benefits.

As such, during these days, the strategy increases position size to a maximum of 2 when two or more signals overlap on the same trading day, meaning that we would double our original investment these days. This specific value was chosen as a conservative compromise that allows increased exposure during high-conviction periods without approaching continuous or leveraged market participation for an everyday investor. Because the individual components are motivated by distinct economic mechanisms, days on which multiple signals coincide represent periods in which several independent sources of return predictability align simultaneously. Therefore, increasing exposure during these intersections concentrates capital precisely when expected returns are highest.

In doing this we get a CAGR of around 11.29% and the Sharpe ratio improves to 0.72, while effective market exposure remains well below that of buy-and-hold strategy at roughly 174 trading days per year (Table 5.2, Appendix). These 29 additional trading days are not the strategy investing in 29 other days but rather increasing exposure on 29 out of the 145 days the strategy signals investment. Thus, this proves that days characterized by overlapping calendar and event-based conditions must exhibit systematically higher expected returns than days in which only a single effect is present. Looking at it economically, this matches the interpretation that independent sources of predictable excess return components reinforce one another when they occur simultaneously.

These results highlight an important distinction between discrete calendar and event-based effects and smoother seasonal patterns. Effects such as turn-of-the-month and pre-holiday behavior identify narrow, well-defined periods of elevated expected returns and therefore contribute strongly to capital efficiency. In contrast, the seasonal expected return signal spreads exposure across a larger fraction of the trading calendar, increasing raw returns but diluting the concentration of returns per unit of exposure. For this reason, the analysis favors a parsimonious combination of discrete calendar and event-driven effects as the core strategy, while treating seasonal expected returns as informative but secondary. This choice emphasizes interpretability, robustness, and economic intuition over maximizing in-sample returns through increased market exposure.

Figure 1- Equity Curves with 2 Different strategies and benchmark with logarithmic scale

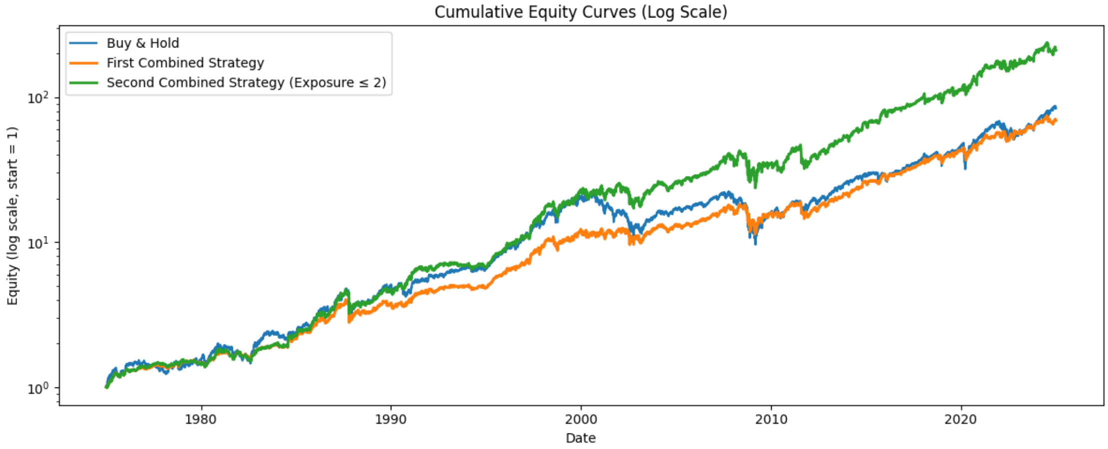

Figure 1 presents the cumulative equity curves of the buy-and-hold benchmark and the two combined strategies on a logarithmic scale. The logarithmic scale facilitates comparison of long-run growth rates by putting equal proportional changes on a common visual footing. The graph illustrates that the second combined strategy has a consistently steeper growth trajectory than the first combined specification and the buy-and-hold benchmark, suggesting a higher long-term growth rate as opposed to sporadic outperforming periods. For completeness, the linear-scale version of the same figure is presented in the Appendix, Figure 2; however, because of the significant variation in terminal wealth across strategies, it is less instructive.

Although this improvement may be good and the strategy itself may be more effective in terms of CAGR and Sharpe ratio than simple buy-and-hold strategy, it must be noted that these are the results of this specific 50-year sample and does not mean that this is a better alternative for passive investment.

To understand the source of the combined strategy's performance, the return profile of the complement set of trading days are examined. That is: days on which the strategy remains out of the market. If the strategy's gains were driven primarily by leverage or mechanical exposure effects, the excluded days would be expected to exhibit similar average returns. However, this is not the case. On figure 7 of the Appendix, the complement strategy delivers an annualized return of only 0.38% with a Sharpe ratio of 0.09, while investing for over 106 days. This indicates that the avoided days are characterized by poor risk–return trade-offs, contributing to very low long-run growth while still exposing investors to substantial volatility.

As was the case with other effects explained by Marquering et al. (2006), effects tend to lose strength and some even disappear after publication as markets become more efficient. Ultimately, this is what happens with the combined strategy with increased exposure as even though results look interesting on paper, the effect dies down as happens with most effects that as the building blocks of the strategy. On Table 8 of the Appendix, we can notice that prior to 2000 CAGR was 13.33% and fell to 9.27% after 2000 and Sharpe ratio also suffered, falling from 1.00 before 2000 to 0.54 after the turn of the century.

6. Limitations

There are several important limitations that should be acknowledged when interpreting the results. First, the analysis relies exclusively on daily close-to-close returns of the S&P 500 index. While this choice reflects the perspective of a realistic, low-frequency investor, it necessarily abstracts from intraday dynamics that are central to the identification of several event-driven anomalies in literature. Documented effects surrounding Federal Reserve announcements and earnings releases often occur within narrow intraday windows. Aggregating returns to the daily level may therefore reduce the magnitude of these effects.

Second, the empirical framework assumes frictionless trading. Transaction costs, bid–ask spreads, and market impact are not explicitly modeled. Although the S&P 500 Index is quite liquid and has a very small bid-ask spread, the transaction costs of opening and closing positions as many times as the strategy requires would probably reduce the earnings quite substantially.

The results should therefore be interpreted as upper bounds on achievable performance instead of precise estimates of implementable returns as these costs would in theory damage a big part of the strategy's returns.

Third, the construction of trading signals necessarily involves several design choices and variables, including threshold levels, holding periods, and precise signal definitions. Although these choices are motivated by economic intuition, they introduce the risk of specification and selection bias. Some parameter values may perform well partly because they were selected after examining historical data, such as the threshold for the Earnings Intensity signal or the number of holding days in the FOMC strategy, and alternative specifications could lead to materially different results. As a result, the reported performance may overstate the effectiveness of the strategies.

Fourth, although the analysis documents statistically significant return patterns associated with calendar-based and event-driven signals, these effects are identified using historical daily data and may reflect in-sample regularities rather than stable sources of out-of-sample excess returns. This raises the possibility that some of the detected predictability reflects survivorship or data-mining effects, whereby patterns that are strong in historical samples wear out as markets evolve and investor awareness increases. One thing is clear: these results do not imply that timing the market is superior to maintaining continuous exposure for a typical long-term investor. The strategy should then be interpreted as showing how return predictability can arise from recurring time and information cycles, and not as a robust substitute for passive investment. For many investors, sustained time in the market should remain a more reliable approach instead of trying to time the market.

Finally, many of the anomalies examined have been shown in the literature to lose strength after their initial discovery, as increased investor awareness and arbitrage activity reduce their profitability over time. The results presented here are therefore subject to the risk that the documented effects reflect historical regularities that may not persist in the future, the most noticeable of these is the TOM results where before 2000, as seen in Table 3.1 of the Appendix its CAGR was 6.36% and Sharpe Ratio of 0.99 which reduced quite a lot after 2000 with CAGR of 2.13% and Sharpe Ratio of 0.30. Even if some of the individual anomalies remain statistically detectable, their economic magnitude may continue to decline, limiting their usefulness in real-time trading. This concern is particularly relevant for combined strategies, which rely on the continued presence of multiple effects whose strength has also declined throughout time.

7. Conclusion

This thesis examines whether return regularities documented in academic literature can be identified at the aggregate index level and integrated into a coherent binary trading framework. Inspired by established research on calendar-based anomalies and scheduled information events, the analysis focuses on four effects: the turn-of-the-month, pre-holiday behavior, Federal Reserve policy meetings, and periods of concentrated earnings announcements.

Using daily S&P 500 data for the past 50 full years, the thesis first looks at each effect individually and finds that several well-known anomalies are significant when translated into close-to-close index returns, albeit with varying strength across subperiods. These findings are in line with the literature documenting both the persistence and weakening of return anomalies over time, and as such they should be interpreted with caution.

It was also crucial to see how these effects interact when combined within a single trading strategy. The results show that concentrating market exposure on days associated with multiple independent sources of return predictability improves risk-adjusted performance relative to passive buy-and-hold investment, even without continuous market participation. Importantly, the strongest gains arise not from increased exposure, but from selectively and carefully allocating capital to periods characterized by overlapping calendar and information-driven signals.

Taken as a whole, the results show that equity returns exhibit systematic variation linked to recurring calendar patterns and scheduled information events. No individual anomaly considered in this thesis delivers a dominant or standalone source of excess returns. However, when several economically distinct effects are combined, they can produce modest but persistent improvements in performance relative to a passive benchmark. Mainly, the results show that insights from the academic literature on return predictability can be implemented in a transparent framework that can be relevant for long-only investors but should not be viewed as an alternative to a buy-and-hold strategy.

References

Ariel, Robert A. 1987. “A Monthly Effect in Stock Returns.” *Journal of Financial Economics* 18 (1): 161–74. https://doi.org/10.1016/0304-405X(87)90066-3.

Ariel, Robert A. 1990. *High Stock Returns before Holidays: Existence and Evidence on Possible Causes*. https://doi.org/10.1111/j.1540-6261.1990.tb03731.x

Bernard, Victor L., and Jacob K. Thomas. 1989. “Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?” *Journal of Accounting Research* 27: 1. https://doi.org/10.2307/2491062.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen. 2019. “Stock Returns over the FOMC Cycle.” *The Journal of Finance* 74 (5): 2201–48. https://doi.org/10.1111/jofi.12818.

Francis, Jennifer, Ryan Lafond, Per Olsson, and Katherine Schipper. 2007. “Information Uncertainty and Post-Earnings-Announcement-Drift.” *Journal of Business Finance & Accounting* 34 (3–4): 403–33. https://doi.org/10.1111/j.1468-5957.2007.02030.x.

Heston, S, and R Sadka. 2008. “Seasonality in the Cross-Section of Stock Returns☆.” *Journal of Financial Economics* 87 (2): 418–45. https://doi.org/10.1016/j.jfineco.2007.02.003.

Keim, Donald B. 1983. “Size-Related Anomalies and Stock Return Seasonality.” *Journal of Financial Economics* 12 (1): 13–32. https://doi.org/10.1016/0304-405X(83)90025-9.

Lakonishok, Josef, and Seymour Smidt. 1988. “Are Seasonal Anomalies Real? A Ninety-Year Perspective.” *Review of Financial Studies* 1 (4): 403–25. https://doi.org/10.1093/rfs/1.4.403.

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.

Marquering, Wessel, Johan Nisser, and Toni Valla. 2006. “Disappearing Anomalies: A Dynamic Analysis of the Persistence of Anomalies.” *Applied Financial Economics* 16 (4): 291–302. https://doi.org/10.1080/09603100500400361.

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. 2003. “Chapter 15 Anomalies and Market Efficiency.” In *Handbook of the Economics of Finance*, vol. 1. Elsevier. https://doi.org/10.1016/S1574-0102(03)01024-0.

Xu, Wei, and John J. McConnell. 2006. “Equity Returns at the Turn of the Month.” *SSRN Electronic Journal*, ahead of print. https://doi.org/10.2139/ssrn.917884.

Appendix

Table 1 – S&P Index buy-and-hold benchmark results

Strategy	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
Buy & Hold	9.25	0.60	17.34	-56.78	252.1

Table 2- Results per strategy component

Strategy	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
Earnings Intensity	4.07	0.41	11.42	-41.05	93.5
FOMC	1.67	0.39	4.55	-14.72	16.9
Turn-of-Month	4.23	0.60	7.39	-20.00	48.0
Pre-Holiday	1.49	0.42	3.68	-11.53	17.2
Seasonal (DOY)	3.28	0.41	8.89	-28.63	64.5

Table 3.1- - Results per strategy component Before 2000

Strategy Pre 2000	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
Earnings Intensity	3.82	0.46	9.06	-35.49	66.6
FOMC	1.64	0.44	3.84	-13.34	17.8
Turn-of-Month	6.36	0.99	6.46	-15.79	48.0
Pre-Holiday	1.37	0.44	3.19	-11.53	16.2
Seasonal (DOY)	3.38	0.46	7.82	-18.24	71.6

Table 3.2- Results per strategy component After 2000

Strategy Post-2000	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
Earnings Intensity	4.32	0.38	13.38	-41.05	120.4
FOMC	1.70	0.35	5.16	-14.72	15.9
Turn-of-Month	2.13	0.30	8.22	-20.00	48.0
Pre-Holiday	1.60	0.41	4.10	-9.28	18.2
Seasonal (DOY)	3.18	0.37	9.84	-28.63	57.5

Table 4- FOMC strategy results

Strategy	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
FOMC (1 holding day)	1.13	0.34	3.44	-12.40	8.42
FOMC (2 holding days)	1.67	0.39	4.55	-14.72	16.9

Table 5.1- First Combined Strategy results with maximum exposure capped at 1

Strategy	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
First Combined Strategy	8.84	0.70	13.50	-37.99	145.2

Table 5.2- Second Combined Strategy results with maximum exposure capped at 2

Strategy	CAGR (%)	Sharpe	Vol Ann (%)	MaxDD (%)	Days Invested per Year
Second Combined Strategy	11.29	0.72	16.90	-44.66	173.8

Table 6.1- Daily overlapping signals per effect per year on average

	Earnings	Turn-of-Month	FOMC	Pre-Holiday
Earnings	93.48	14.12	6.82	3.18
Turn-of-Month	14.12	48.00	2.86	5.00
FOMC	6.82	2.86	16.86	0.26
Pre-Holiday	3.18	5.00	0.26	17.24

Table 6.2- Return correlations between signals per year on average

	Earnings	Turn-of-Month	FOMC	Pre-Holiday
Earnings	1.000	0.204	0.225	0.068
Turn-of-Month	0.204	1.000	0.107	0.130
FOMC	0.225	0.107	1.000	0.003
Pre-Holiday	0.068	0.130	0.003	1.000

Table 7- Results for days not investing in the Strategy

CAGR (%)	Sharpe	Vol Ann (%)	Days Invested per Year
0.3825	0.0896	10.885	106.98

Table 8- Second Combined Strategy results across different periods

Period	CAGR (%)	Sharpe	Vol Ann (%)	Days Invested per Year
Pre-2000	13.33	1.00	13.30	147.32
Post-2000	9.27	0.54	19.87	200.20

Figure 2- Equity Curves with 2 Different strategies and benchmark with linear scale