Master in International Finance

RESEARCH PAPER

Academic Year 2025-2026

Fresh Losses, Stale Gains: Return Path Salience and the Cross-Section of Equity Returns

Ulysse HA THUC[1] and Orestas FREIGOFAS[2]

Under the supervision of Prof. Ciaran ROGERS

President of the Jury: Prof. Johan HOMBERT

Keywords: Return Path Salience; cross-sectional asset pricing; investor attention; chronological return ordering; short-term reversal; equity anomalies

Public Report

May 2026

[1] Eagle Point, 161 City Rd, London EC1V 1AQ, United Kingdom; +39 347 538 4764; ulysse.ha-thuc@hec.edu
[2] Eagle Point, 161 City Rd, London EC1V 1AQ, United Kingdom; +39 345 053 4188; orestas.freigofas@hec.edu

Title: Fresh Losses, Stale Gains: Return Path Salience and the Cross-Section of Equity Returns

Authors: Ulysse HA THUC and Orestas FREIGOFAS

Abstract

This thesis introduces Return Path Salience (RPS), a family of sign-conditional, recency-weighted characteristics measuring the chronology of losses and gains within recent formation windows of daily equity returns. Using U.S. common-equity data over 1985-2025 with a parallel European replication, we document that the combined RPS^{diff} long-short spread earns approximately 1.45% per month at the one-month horizon (Sharpe 1.76, NW-t=8.85), largely unchanged after standard factor adjustment. The signal survives universe restriction, replicates in Europe, and spans the unconditional chronological-return-ordering benchmark while the reverse does not hold; establishing that sign-conditional path timing is priced beyond cumulative-return and ordering measures.

Contents

1. Introduction

A central question in empirical asset pricing is whether investors process past return information in a way fully consistent with efficient-market benchmarks. Conventional predictors typically summarise recent history through cumulative performance: whether a stock has gone up or down over a given horizon. Yet the same monthly return can be reached through very different paths. One stock may suffer losses early in the month and recover gradually; another may rise first and decline sharply just before the formation date. These histories are equivalent for a path-invariant return characteristic, but they may not be equivalent if investors respond differently to the timing, sign, and salience of recent shocks.

A large behavioural literature documents systematic departures from path-invariant processing. Investors attend disproportionately to salient observations, meaning returns that are especially noticeable because of their magnitude, recency, contrast with surrounding observations, or negative sign (Barber and Odean, 2008; Da, Engelberg and Gao, 2011). They also overweight recent observations relative to equally informative older ones (Greenwood and Shleifer, 2014; Da et al., 2021) and evaluate outcomes in ways that depend on salience and comparison effects (Bordalo, Gennaioli and Shleifer, 2012, 2013). The common implication is that the temporal structure of past returns may matter for prices, rather than only the cumulative return over the formation period. Standard reversal and momentum variables do not capture this structure: they assign identical signals to two stocks with the same set of daily returns but different chronological sequences.

Recent work has begun to address this gap. Mohrschladt (2021) introduces the chronological return ordering (CRO) measure, which captures whether high or low daily returns occur earlier or later within a formation window. Cakici and Zaremba (2023) extend this evidence internationally and show that return ordering is priced across global equity markets. However, CRO is an unconditional ordering statistic: it aggregates positive and negative daily returns into a single signed correlation. A large loss near month-end and a large gain near month-end therefore enter the statistic in opposite directions, although the timing of losses and the timing of gains may contain economically distinct information. This distinction is particularly relevant in light of long-documented evidence on salience and asymmetric processing of gains and losses.

This thesis tests that conjecture. We construct a family of Return Path Salience (RPS) characteristics that partition daily returns by sign before applying a power-law recency kernel - frequency-based or magnitude-weighted - to measure the conditional chronology of losses and gains separately and as a combined signal. The latter relative measure, RPS^{diff}, compares the two: stocks with a low RPS^{diff} have recent losses combined with stale gains, while stocks with a high RPS^{diff} have the mirror pattern of recent gains and stale losses. The empirical question is whether this signed return-path structure contains cross-sectional pricing information beyond conventional return predictors and existing path-ordering measures.

The analysis is conducted on U.S. equities from 1985 onward, with a parallel European replication used as external validation. The results show that the magnitude-weighted RPS^{diff} characteristic generates large and statistically significant factor-mimicking returns. In the U.S. one-month baseline, the long-short spread earns approx. 1.45% per month, with an ann. Sharpe ratio of 1.76 and a Newey-West t-statistic close to 9. The spread remains essentially unchanged after standard factor adjustment, survives market-cap universe restrictions, is not subsumed by CRO, and replicates in Europe with smaller but still significant magnitudes. The evidence therefore suggests that the priced object is not only whether returns occur early or late within the formation window, but whether losses and gains occur early or late relative to one another.

The remainder of the thesis proceeds as follows. Section 2 reviews the related literature on attention biases, salience, recency, chronological return ordering, and standard return predictors. Section 3 defines the RPS characteristics and detailed empirical design. Section 4 presents the main empirical results. Section 5 discusses interpretation, limitations, and future research directions. Section 6 concludes. All figures and tables referenced in the body of the text can be found in the Appendix section.

2. Literature Review & Theoretical Framework

2.1. Cognitive biases and the pricing of past return information

The efficient markets hypothesis holds that asset prices incorporate all available information instantaneously and without bias (Fama, 1970). There is a large body of empirical work documenting systematic departures from this benchmark, with a foundational source of these departures being the selective and cognitively distorted way in which investors process historical return information. For example, Barber and Odean (2008) find that retail investors concentrate their purchases on attention-grabbing stocks - those that exhibit prominent news coverage, extreme returns, or unusual trading volume - suggesting that the salience of information (understood here as the prominence or attention-weight assigned to a return shock), rather than its fundamental content, also shapes investment decisions. Further, Da, Engelberg and Gao (2011) proxy investor attention through retail internet search activity and demonstrate that attention-driven demand generates predictable short-run price pressure followed by reversals, establishing a clean empirical link between cognitive engagement with past performance and subsequent cross-sectional return patterns. These findings share a common implication: how investors weight past information - not only what that information contains - has first-order consequences for asset pricing.

Among the cognitive distortions that affect return expectations, the recency effect is particularly well-documented. The recency effect is the tendency for individuals to weight recent observations more heavily than earlier ones of equal informational content, a phenomenon with deep roots in the psychology of memory and sequential learning (Murdock, 1962; Hogarth and Einhorn, 1992). We also note that in cognitive models of memory, recency is often represented by declining retention functions, with power-law decay providing a parsimonious way to capture rapid initial decay followed by slower long-run attenuation. In financial markets, this manifests as a systematic pattern in which investors extrapolate recent return performance when forming expectations, as documented through survey evidence from Greenwood and Shleifer (2014) and Da et al. (2021). The timing and ordering of past return realizations matter to investors, and this sensitivity produces measurable effects in prices.

2.2. Salience Theory and the chronological weighting of returns

A formal theoretical framework for understanding how the structure of past returns affects investor expectations is provided by the salience theory of Bordalo, Gennaioli and Shleifer (2012, 2013). In this framework, agents do not evaluate prospects by their objective probabilities. Instead, they overweight outcomes that are salient relative to a reference distribution, with attention drawn disproportionately to observations that stand out from the comparison set by reason of their magnitude or contrast. They extend this framework to asset markets, predicting that stocks whose return distributions feature comparatively salient positive payoffs will be perceived as more attractive than their fundamentals justify, leading to overvaluation and subsequent underperformance. Conversely, stocks whose distributions are characterised by comparatively salient negative payoffs will be undervalued, offering prospective positive abnormal returns.

The asset pricing implications of salience theory have been tested directly by Cosemans and Frehen (2021), who construct an empirical salience measure by comparing the level of each stock's cumulative historical returns against the cross-sectional distribution of returns in the same period. The comparison is thus across stocks at a point in time - it asks which stock's cumulative return history stands out relative to the distribution of contemporaneous returns - rather than within a single stock's sequential return path. Their evidence confirms the theory's central prediction: stocks whose cumulative return histories appear salient relative to the cross-section are mispriced in the direction salience theory would predict, with corrections yielding significant abnormal returns. Critically, this predictability is concentrated in small and illiquid firms where limits to arbitrage are most binding, a finding corroborated in an international setting by Cakici and Zaremba (2022).

This limitation motivates a re-examination of the channel through which salience and related cognitive distortions manifest in prices. The Cosemans-Frehen measure above operates cross-sectionally, comparing the level of cumulative past returns across stocks. Salience theory, however, also carries a natural temporal implication: within a single stock's return history, outcomes that are chronologically proximate to the evaluation date are more cognitively available and therefore more salient than identical outcomes from the distant past. Observing return distribution within each stock’s own formation window, the relevant comparison shifts from the cumulative distribution of returns across firms to the sequential ordering of returns within a single return path. This perspective leads naturally to the question of whether sign-conditional return ordering contains information about future returns independently of cumulative performance levels, and whether such a signal built on return ordering rather than return levels may achieve greater robustness across the firm-size spectrum.

2.3. Chronological Return Ordering

A closely related empirical framework is provided by Mohrschladt (2021), who introduces the Chronological Return Ordering (CRO) variable to measure the ordering of daily returns within a formation window. The construction is grounded in the recency-bias hypothesis: if investors overweight recent return observations relative to distant ones when assessing a stock's prospects, then a stock whose daily returns declined over the formation window - ending on a relatively weak trajectory - will be systematically undervalued relative to a stock with the same cumulative return but an improving recent path. CRO formalises this intuition as a correlation statistic computed from the daily return sequence.

Specifically, for a stock with T valid trading days in the formation window, let r_t denote the daily return on day t and let d_t denote the number of trading days remaining until the end of the window, so that $d_t = 1$ for the final trading day and $d_t = T$ for the first. CRO is then defined as the Pearson correlation between the daily return sequence and the days-remaining sequence:

$$CRO = Corr(r_t, d_t) \ , \ \ t = 1, \ldots, T$$

A positive value of CRO indicates that returns were higher on days with more time remaining (early in the window and lower near the end), meaning the stock experienced deteriorating returns over the formation period. A negative value of CRO indicates the opposite: returns strengthened toward the end of the window. Under the recency-bias hypothesis, investors in a stock with high CRO will have been disproportionately influenced by its recent weakness, driving the price below fundamental value and creating a positive expected return going forward. Mohrschladt (2021) confirms this prediction in the U.S. cross-section: high-CRO stocks outperform low-CRO stocks in the subsequent month, with the spread surviving standard factor controls. The construction has an important characteristic: by computing a single unconditional correlation over the full signed return sequence, CRO treats every trading day symmetrically regardless of whether the return was positive or negative. The sign of the daily return enters the measure only through its contribution to the overall r_t sequence. This aggregation is both the measure's strength - simple, robust, and easily computed - and the feature that motivates the extensions examined in this thesis.

Cakici and Zaremba (2023) - distinct from their earlier work on salience theory (Cakici and Zaremba, 2022) - extend CRO to 49 international equity markets. The international breadth of the evidence is significant: the return predictability of within-window ordering is not US-specific but appears to be a prevalent feature of equity markets across geographies. Interestingly, unlike the salience-level signal of Cosemans and Frehen (2021) and the international salience evidence of Cakici and Zaremba (2022), the CRO does not suffer from microcap concentration and in fact is stronger among large caps. This follows from the construction of CRO itself: as an unconditional rank correlation over the full return sequence, it does not rely on extreme or salient payoffs which are the features that confine salience-based signals to illiquid microcaps where limits to arbitrage bind most tightly.

Overall, withstanding some differences across each signal, the literature confirms the robustness of the measure and establishes that within-window return ordering contains information statistically orthogonal to the known factor structure. When tested across different formation windows (from one to six months), the chronological ordering remains informative over longer return histories. Together, these properties establish return-order-based signals as a legitimate and practically relevant class of cross-sectional predictors.

2.4. Beyond unconditional return-path measures

The CRO statistic nevertheless leaves an important question open, which motivates the contribution of the present thesis. Indeed, the unconditional correlation in the CRO formula aggregates positive and negative daily return observations into a single signed summary, making no distinction between the timing of loss days and the timing of gain days. A large negative return near month-end and a large positive return near month-end both influence the correlation, but in opposite directions - they partially cancel within the same statistic rather than being treated as separate informational events. Yet the behavioural mechanisms that motivate return-path effects may not operate symmetrically across the sign of the underlying shock. Recency, availability, and loss aversion all suggest that negative and positive observations receive different weights in investors' assessment of the same return history. The relative temporal positioning of losses versus gains within the window may therefore be a distinct and separately priced feature of the return path that a single unconditional correlation, by construction, cannot isolate. Liu and Wang (2025) provide an early indication that sign-conditional timing carries incremental information, documenting return predictability from a frequency-based measure of the recency of negative return days alone. If investor attention and memory distortion respond disproportionately to large signed shocks (as salience theory directly implies) then the economically relevant question is not just where large returns occurred, but where large losses and large gains occurred relative to one another. CRO's unconditional aggregation of signed observations motivates an extension to sign-conditional characteristics. The empirical question is whether such a refinement captures pricing information about the chronological ordering of signed return shocks that the unconditional correlation measure leaves on the table, and whether that information is priced in the cross-section of equities independently of the level of past returns.

2.5. Differentiation from standard return factors

The signal most likely to be confounded with a return-ordering measure is short-term reversal - the well-documented negative predictability of the prior month's cumulative return (Jegadeesh, 1990; Lehmann, 1990). This distinction is fundamental. Reversal is a function of the cumulative signed return over the formation window and is invariant to the ordering of daily returns within it: two stocks with identical daily return sets but different temporal sequences receive the same reversal signal. A return-ordering signal, by contrast, is explicitly permutation-sensitive; rearranging the daily sequence alters the signal while leaving the cumulative return unchanged. The two characteristics therefore capture orthogonal dimensions of the return history. Intermediate-horizon momentum, capturing return continuation over a multi-month horizon, is even more clearly distinct as it operates over a different window and measures a different economic phenomenon entirely. As established in the earlier literature, these controls establish that the pricing of within-window return ordering is a distinct empirical object and not a recombination of known factors.

3. Methodology

3.1. Empirical setting and data construction

We study whether the chronological ordering of signed return shocks within a recent formation window contains information about subsequent equity returns. The empirical object is not only the level of a stock's past return, but the path through which that return was realised. More specifically, the analysis examines whether negative and positive return observations occurring at different points within the formation window are priced differently in the cross-section.

The main empirical setting is the U.S. equity market, with the baseline sample from 1985 to early 2025. The sample is constructed from daily equity return data accessed through WRDS/CRSP and compiled in the empirical code as a daily stock-level panel. Empirical filters are designed to remove mechanically unreliable observations rather than to impose an ex-ante investment universe: observations with missing identifiers, missing returns, or missing prices are removed; stocks with prices below USD 5 and less than 10 trading days are excluded; and daily returns are restricted to generally economically plausible values, with observations outside $[-95\%, 300\%]$ removed. Forward returns are computed from daily returns by compounding over the next calendar month. The resulting stock-month forward returns are then cross-sectionally winsorised at the 1st and 99th percentiles within each return month. This ensures that the predictive tests are not driven by a small number of extreme observations while preserving the cross-sectional ranking of the signal.

The analysis is first conducted on the broad universe of U.S. common equities and is then repeated on more restrictive market-capitalisation universes, namely the Top 1000, Top 500, and Top 300 stocks. These restrictions assess whether the documented cross-sectional relation is driven by the full universe or remains visible in more investable subsets. The same methodology is subsequently applied to European equities as an external validation exercise.

Let i index stocks and t index calendar months. For each stock-month observation, daily returns are collected over a formation window $W_t^{(H)}$, where H denotes the lookback horizon. The baseline construction uses a one-month formation window, while robustness tests vary the horizon. Let $r_{i,\tau}$ denote the daily return of stock $i$on trading day τinside the formation window. Trading days are ordered from oldest to most recent: $\tau = 1,2, \dots, T_{i,t}$, where $T_{i,t}$is the number of valid daily observations for stock i in the formation window. The chronological age of each observation is defined as:

$$a_\tau = T_{i,t} - \tau + 1. \qquad (1)$$

Thus, $a_\tau = 1$ for the most recent trading day and $a_\tau = T_{i,t}$for the oldest trading day in the window. All signals are constructed using only information available at the portfolio formation date.

3.2. Return Path Salience: recency kernel and signed partitions

The construction of Return Path Salience (RPS) is motivated by a literature that treats the order of return realisations as economically informative, rather than reducing a month to its cumulative return. The literature section discusses the behavioural motivation in greater detail; at this stage, the relevant methodological point is that prior work provides two useful starting points. First, a frequency-based recency formulation applies a declining memory kernel to the timing of signed return days. Second, the chronological return ordering literature measures whether high or low returns occur early or late within the formation window. Building on these approaches, this thesis formalises RPS as a family of sign-conditional path-timing statistics: daily returns are partitioned into negative and positive shocks, optionally weighted by magnitude, and then compared through one-sided and relative timing measures.

To operationalise recency, we use a power-law weighting kernel. This choice is a parsimonious reduced-form way to encode declining attention or memory weight: recent observations receive the highest weight, while older observations are discounted smoothly rather than excluded discretely. The specification is not intended to impose a structural model of investor memory; it simply maps each observation's chronological age into a monotone recency weight.

$$q_\alpha(a_\tau) = a_\tau^{-\alpha}, \alpha > 0. \qquad (2)$$

Since $a_\tau = 1$ for the most recent trading day,

$$q_\alpha(1) = 1,$$

while older observations receive lower weights. The parameter α controls the steepness of the decay. A larger value of α places more relative weight on observations near the end of the formation window, while a smaller value produces a flatter weighting profile. Daily returns are then partitioned into negative and positive subsets:

$$\mathcal{N}_{i,t} = \{\tau \in W_t^{(H)}: r_{i,\tau} < 0\}, \qquad (3)$$

$$\mathcal{P}_{i,t} = \{\tau \in W_t^{(H)}: r_{i,\tau} > 0\}. \qquad (4)$$

Zero-return days are assigned to neither the positive nor the negative signed partition. When a component requires conditioning on a sign that is absent for a given stock-month, the relevant component is undefined for that observation and excluded from specifications requiring that variable. The resulting panel is unbalanced, reflecting firm entry and exit over time. We refer to the family of variables outlined above as Return Path Salience, denoted RPS. The term is intended to be descriptive rather than structural. RPS measures the chronological salience of signed return observations within the formation window. The behavioural motivation provides a basis for the construction, but the empirical object itself is a measurable path characteristic and does not require the assumption that the underlying mechanism is uniquely behavioural.

3.2.1. Frequency-based signed RPS components

The first specification measures the timing of signed return days without weighting by return magnitude. The negative-side component is defined as

$$RPS_{i,t}^{-} = 1 - \frac{\sum_{\tau \in W_t^{(H)}} q_\alpha(a_\tau)\mathbf{1}(r_{i,\tau} < 0)}{\sum_{\tau \in W_t^{(H)}} \mathbf{1}(r_{i,\tau} < 0)}. \qquad (5)$$

The fraction in Equation (5) is the average recency weight of negative-return days. Since recent observations have high values of $q_\alpha(a_\tau)$, recent losses increase the fraction and therefore reduce RPS^-. Older losses receive lower values of $q_\alpha(a_\tau)$, reducing the fraction and increasing RPS^-.

The positive-side component is defined analogously:

$$RPS_{i,t}^{+} = 1 - \frac{\sum_{\tau \in W_t^{(H)}} q_\alpha(a_\tau)\mathbf{1}(r_{i,\tau} > 0)}{\sum_{\tau \in W_t^{(H)}} \mathbf{1}(r_{i,\tau} > 0)}. \qquad (6)$$

The central feature of Equations (5) and (6) is that they are conditional timing statistics. They do not measure the total number of gains or losses, nor the cumulative return over the period. Instead, they measure the average chronological location of observations after conditioning on their sign. Formally, for a trading day drawn uniformly from the valid observations in $W_t^{(H)}$,

$$RPS_{i,t}^{-} = 1 - \mathbb{E}\left[q_\alpha(a_\tau) \mid r_{i,\tau} < 0\right], \qquad (7)$$

$$RPS_{i,t}^{+} = 1 - \mathbb{E}\left[q_\alpha(a_\tau) \mid r_{i,\tau} > 0\right]. \qquad (8)$$

The conditioning event is therefore not incidental; it defines the object being measured. $RPS^-$is the complement of the expected recency weight conditional on the return being negative. $RPS^+$is the complement of the expected recency weight conditional on the return being positive. This also explains why the negative and positive components are not symmetric complements. To briefly show this, let

$$p_{i,t}^{-} = \mathbb{P}(r_{i,\tau} < 0), p_{i,t}^{+} = \mathbb{P}(r_{i,\tau} > 0). \qquad (9)$$

Ignoring zero-return days for simplicity, the unconditional average recency weight decomposes as

$$\mathbb{E}[q_\alpha(a_\tau)] = p_{i,t}^{-}\mathbb{E}[q_\alpha(a_\tau) \mid r_{i,\tau} < 0] + p_{i,t}^{+}\mathbb{E}[q_\alpha(a_\tau) \mid r_{i,\tau} > 0]. \qquad (10)$$

Equation (10) shows that the positive and negative conditional expectations enter the unconditional average only after being weighted by the frequency of positive and negative days. By contrast,

RPS^{-} and RPS^{+} each remove the corresponding frequency through their own conditional denominator. Therefore,

$$RPS_{i,t}^{-} + RPS_{i,t}^{+} \neq C \qquad (11)$$

for any constant C in general. The two components are calculated over disjoint subsets of the return path and need not move mechanically in opposite directions.

3.2.2. Magnitude-weighted RPS

The frequency-based specification treats all positive or negative return days equally. However, large signed return shocks may be more salient than small daily fluctuations. We therefore define a more general class of RPS variables using a salience function $g(| r_{i,\tau} |)$. We considered two specifications:

$$g_0(| r |) = 1, g_1(| r |) =| r |. \qquad (12)$$

The first corresponds to the frequency-based measure. The second gives greater weight to larger absolute daily returns. The general negative-side component is

$$RPS_{i,t}^{-,g} = 1 - \frac{\sum_{\tau \in W_t^{(H)}} q_\alpha(a_\tau) g(| r_{i,\tau} |) \mathbf{1}(r_{i,\tau} < 0)}{\sum_{\tau \in W_t^{(H)}} g(| r_{i,\tau} |) \mathbf{1}(r_{i,\tau} < 0)}. \qquad (13)$$

The positive-side component is

$$RPS_{i,t}^{+,g} = 1 - \frac{\sum_{\tau \in W_t^{(H)}} q_\alpha(a_\tau) g(| r_{i,\tau} |) \mathbf{1}(r_{i,\tau} > 0)}{\sum_{\tau \in W_t^{(H)}} g(| r_{i,\tau} |) \mathbf{1}(r_{i,\tau} > 0)}. \qquad (14)$$

The magnitude-weighted version is not simply a return-weighted version of reversal. It remains a conditional timing statistic. If for the negative side we define

$$\pi_{i,\tau}^{-,g} = \frac{g(| r_{i,\tau} |) \mathbf{1}(r_{i,\tau} < 0)}{\sum_{s \in W_t^{(H)}} g(| r_{i,s} |) \mathbf{1}(r_{i,s} < 0)}. \qquad (15)$$

Then Equation (13) can be written as

$$RPS_{i,t}^{-,g} = 1 - \sum_{\tau \in W_t^{(H)}} \pi_{i,\tau}^{-,g} q_\alpha(a_\tau). \qquad (16)$$

Thus, the magnitude-weighted signal asks where the negative return mass occurred within the formation window. The positive-side version asks the same for positive return mass.

3.2.3. Approximation, parameter behaviour, and interpretation

The power-law kernel admits a useful approximation. Since

$$q_\alpha(a) = a^{-\alpha} = \exp\left[-\alpha \log(a)\right], \qquad (17)$$

a first-order expansion around $\alpha = 0$ gives

$$q_\alpha(a) \approx 1 - \alpha \log(a). \qquad (18)$$

Substituting Equation (18) into the frequency-based components yields

$$RPS_{i,t}^{-} \approx \alpha \mathbb{E}\left[\log(a_\tau) \mid r_{i,\tau} < 0\right], \qquad (19)$$

$$RPS_{i,t}^{+} \approx \alpha \mathbb{E}\left[\log(a_\tau) \mid r_{i,\tau} > 0\right]. \qquad (20)$$

For the magnitude-weighted version the same approximation holds with weighted conditional expectations:

$$RPS_{i,t}^{-,g} \approx \alpha \mathbb{E}_g[\log(a_\tau) \mid r_{i,\tau} < 0], \qquad (21)$$

$$RPS_{i,t}^{+,g} \approx \alpha \mathbb{E}_g[\log(a_\tau) \mid r_{i,\tau} > 0]. \qquad (22)$$

Equations (19) to (22) provide the main mathematical interpretation of RPS. The signal is approximately an average log-age statistic, computed conditionally on the sign of returns and optionally weighted by their magnitude. A higher value indicates that the relevant signed observations are older within the formation window; a lower value indicates that they are more recent.

The parameter α determines how strongly chronological distance affects the signal. When $\alpha \to 0$, $q_\alpha(a) \to 1$, so the numerator and denominator of Equations (5) and (6) converge and the signal approaches zero. In this limit, timing information disappears. On the other hand, as α becomes large,

$$q_\alpha(1) = 1, q_\alpha(a) \to 0 \text{ for } a > 1,$$

so the signal becomes increasingly concentrated on the final observations of the window. Very large α therefore reduces the full path statistic toward a near-endpoint indicator.

The empirical role of α will be evaluated through sensitivity tests. For small values of α, Equations (19) and (20) imply that changing α mainly rescales the signal. Since the empirical analysis relies on cross-sectional rankings and factor-mimicking portfolios, the relevant issue is whether different values of α materially alter the ordering of stocks, not so much if they change the numerical level of the signal.

3.2.4. Relative signed path timing

The negative and positive components measure different conditional timing objects. The negative component measures the age of losses, while the positive component measures the age of gains. Since these are computed over different subsets of the return path, one cannot infer one from the other. This motivates a relative signed path-timing measure. For each salience function g, define

$$RPS_{i,t}^{diff,g} = RPS_{i,t}^{-,g} - RPS_{i,t}^{+,g}. \qquad (23)$$

Using the first-order approximations above,

$$RPS_{i,t}^{diff,g} \approx \alpha[\mathbb{E}_g(\log(a_\tau) \mid r_{i,\tau} < 0) - \mathbb{E}_g(\log(a_\tau) \mid r_{i,\tau} > 0)]. \qquad (24)$$

Equation (24) is the main interpretation of the combined RPS measure. It compares the average log-age of negative return observations with the average log-age of positive return observations. The object is therefore neither a loss-timing measure nor a gain-timing measure alone; but a relative signed chronology measure. If

$$RPS_{i,t}^{diff,g} < 0, \qquad (25)$$

then negative observations have lower average log-age than positive observations. Since lower age corresponds to greater recency, losses are more recent than gains. If

$$RPS_{i,t}^{diff,g} > 0, \qquad (26)$$

then gains are more recent than losses.

The intuition is path-based. A negative value of $RPS^{diff,g}$ corresponds to a return path in which losses are positioned closer to the formation date than gains. A positive value corresponds to a path in which gains are positioned closer to the formation date than losses. The measure therefore captures whether the recent path ended on a relatively negative or relatively positive signed trajectory, after conditioning on the full formation window.

This construction is useful because the one-sided components may be incomplete when considered separately. A stock can have recent losses and recent gains if volatility is clustered near the end of the month. In that case, RPS^{-}identifies the recency of losses and RPS^{+}identifies the recency of gains, but neither component alone fully describes the net signed direction of the path. The difference signal compares the two conditional chronologies directly and assigns one scalar value to the stock.

Note that this construction is related to, but distinct from, chronological return ordering discussed in the literature review. CRO summarises the unconditional ordering of daily returns within a formation window whereas RPS instead conditions on the sign of returns before applying the timing operator. This is a central difference because CRO looks at whether returns were high or low late in the month, whereas RPS looks whether negative and positive shocks occurred at different points in the path. So our empirical analysis therefore treats CRO as a related path-ordering benchmark rather than as a substitute for RPS.

Considering all the above, our analysis treats RPS^{-}, RPS^{+}, and RPS^{diff} as related but distinct characteristics. The component signals identify one-sided signed timing effects. The difference signal tests whether the relative chronology of gains and losses provides a more complete description of the signed return path.

3.2.5. RPS Example visualisation

The definitions above are compact but can be difficult to interpret from the equations alone. We therefore use a simple visual example to show what RPS measures in the Appendix under *Illustrative Examples* chart. It considers three artificial stocks with the exact same set of 21 daily returns over the formation month. The cumulative return, the number of positive and negative days, and the distribution of daily return magnitudes are identical across the three panels; only the chronological ordering of the returns differs. In panel (a), the negative-return days are concentrated toward the end of the month. Since recent observations receive the highest value of the recency kernel $q_{\alpha}(a_{\tau})$, the magnitude-weighted average recency of losses is high, and $RPS^{-} = 1 -$ average recency weight is therefore low. In panel (c), the same negative-return days occur mostly early in the month, so they receive lower recency weights and $RPS^{-}$is high. Panel (b) lies between these two cases. The example shows that RPS^{-}does not measure whether the stock performed badly over the month. It measures where the loss mass occurred within the month, conditional on the return being negative.

The same logic applies to RPS^{+} (middle of chart), with the conditioning event changed from negative to positive return days. A low value of RPS^{+}indicates that gains are concentrated near the end of the formation window, while a high value indicates that gains are relatively stale.

The central empirical object, however, is the relative timing measure $RPS^{diff} = RPS^{-} - RPS^{+}$, shown at the bottom of the chart. Again, all three paths contain the exact same 21 daily returns and have the same cumulative return; only the ordering differs. In panel (g), losses are fresh and gains are stale, so $RPS^{-}$is low, $RPS^{+}$is high, and $RPS^{diff}$is negative. In panel (i), the ordering is reversed: losses are stale, gains are fresh, and $RPS^{diff}$is positive. Panel (h) interleaves gains and losses through the month, producing a value close to zero. Thus, RPS^{diff}summarises whether the recent part of the path is relatively loss-heavy or gain-heavy, after conditioning on the same full-month return realisations.

3.2.6. RPS Long-Short sorting framework

For the empirical analysis we use portfolio sorts to construct factor-mimicking returns associated with the RPS characteristics. As in standard asset-pricing practice, where characteristics such as size, value, profitability, investment, or momentum are translated into long-short (L/S) factor returns. The L/S portfolio is therefore used as a measurement device for estimating the return spread associated with exposure to the RPS characteristic, rather than as the primary object of the thesis.

At the end of each month t, stocks are sorted into deciles according to the relevant RPS variable. For a generic characteristic $X_{i,t}$, let $D1_t$denote the lowest decile and $D10_t$ the highest decile. Equal-weighted decile returns are defined as

$$R_{t+1}^{Dk,EW} = \frac{1}{N_{k,t}} \sum_{i \in Dk_t} r_{i,t+1}\,, \qquad (27)$$

and value-weighted decile returns as

$$R_{t+1}^{Dk,VW} = \sum_{i \in Dk_t} \omega_{i,t}^{VW} r_{i,t+1}, \omega_{i,t}^{VW} = \frac{ME_{i,t}}{\sum_{j \in Dk_t} ME_{j,t}}, \qquad (28)$$

where $ME_{i,t}$is market capitalisation at the formation date.
For the relative signed path-timing signal, the factor-mimicking return is formed as

$$R_{t+1}^{RPS} = R_{t+1}^{D1} - R_{t+1}^{D10} \qquad (29)$$

Under Equation (24), the low-RPS^{diff} portfolio contains stocks for which negative shocks are more recent than positive shocks, while the high-RPS^{diff} portfolio contains stocks for which positive shocks are more recent than negative shocks.

The same sorting framework is applied to the one-sided negative and positive components, with decile spreads oriented according to the economic interpretation of each signal. The full decile patterns are reported, not only the extreme-decile spread, in order to assess monotonicity and tail dependence.

3.3. Identification and robustness design

3.3.1. Benchmark characteristics and identification concerns

The earlier illustrative examples in Section 3.2.5 make this identification issue visible. When the same set of daily returns is rearranged within the formation window, the cumulative return is unchanged, so a standard reversal characteristic assigns the same value to each path. RPS changes because the signed return mass occurs at different chronological positions. The formal discussion below restates this intuition algebraically and links it to the benchmark characteristics used in the empirical tests.

The empirical tests are designed to distinguish RPS from standard return characteristics that may appear mechanically related to signed path timing. The closest alternatives by construction are formation-period reversal, terminal-window reversal, intermediate-term momentum, extreme daily return effects, and chronological return ordering. As briefly noted earlier, these variables summarise either cumulative performance, extreme observations, or unconditional return ordering, whereas RPS measures the conditional chronology of signed shocks. However, the distinction is not just that RPS is order-dependent but also that RPS is constructed from sign-conditional timing expectations, whereas reversal/CRO is constructed from an unconditional signed return sum. We illustrate this important distinction with the summarized formal working below. Let the formation-period return be

$$R_{i,t}^{form} = \prod_{\tau \in W_t^{(H)}} (1 + r_{i,\tau}) - 1. \qquad (30)$$

A standard reversal characteristic can be written as

$$REV_{i,t} = -R_{i,t}^{form}. \qquad (31)$$

Using a first-order return approximation,

$$R_{i,t}^{form} \approx \sum_{\tau \in W_t^{(H)}} r_{i,\tau}\,. \qquad (32)$$

This can be decomposed into positive and negative components as

$$R_{i,t}^{form} \approx N_{i,t}^{+}\mathbb{E}[\mid r_{i,\tau} \mid\mid r_{i,\tau} > 0] - N_{i,t}^{-}\mathbb{E}[\mid r_{i,\tau} \mid\mid r_{i,\tau} < 0] \qquad (33)$$

A reversal variable is therefore a function of the relative frequency and magnitude of gains and losses. It does not depend on where those gains and losses occurred within the formation window. RPS, by contrast, is based on

$$\mathbb{E}[\log{(a_\tau)} \mid r_{i,\tau} < 0] \text{and } \mathbb{E}[\log{(a_\tau)} \mid r_{i,\tau} > 0], \qquad (34)$$

or their magnitude-weighted equivalents.

More intuitively we can also see this distinction through permutation invariance. For any permutation π of daily returns within the formation window,

$$\prod_{\tau=1}^{T_{i,t}} (1 + r_{i,\tau}) - 1 = \prod_{\tau=1}^{T_{i,t}} (1 + r_{i,\pi(\tau)}) - 1, \qquad (35)$$

the cumulative return is unchanged when the order of daily returns is rearranged. Therefore, a reversal variable assigns the same value to two stocks with the same set of daily returns, regardless of whether losses occurred at the beginning or at the end of the window. RPS is not invariant to such permutations. Since $q_\alpha(a_\tau)$ depends on the chronological position of each observation, it follows that

$$\sum_\tau q_\alpha\,(a_\tau)\mathbf{1}(r_{i,\tau} < 0) \neq \sum_\tau q_\alpha\,(a_\tau)\mathbf{1}(r_{i,\pi(\tau)} < 0) \qquad (36)$$

in general. The same applies to positive and magnitude-weighted components. RPS therefore contains information about the ordering of signed shocks that is discarded by cumulative-return variables.

Importantly, this does not imply that RPS is empirically unrelated to reversal: a stock whose losses are concentrated near the end of the formation window may also have poor terminal-period performance. In our empirical design we therefore explicitly control for formation-period reversal, terminal-window reversal, intermediate-term momentum, and calendar effects.

Momentum is distinct for similar reasons. Standard momentum is based on cumulative returns over an intermediate past horizon (often excluding the most recent month). It is designed to capture continuation in past performance. RPS instead measures the within-window chronology of signed shocks. It asks whether losses or gains occur nearer to the formation date, not whether the stock has performed well or poorly over a medium-term horizon.

3.3.2. Cross-sectional predictive and factor regressions

To assess whether RPS factor-mimicking returns are explained by standard risk exposures, the L/S RPS return is regressed on common asset-pricing factors at the monthly frequency:

$$R_t^{RPS,k} = \alpha^k + \beta^{k\prime} F_t + \varepsilon_t^k, \qquad (37)$$

where $R_t^{RPS,k}$denotes the monthly factor-mimicking return associated with RPS specification k. The specification index k covers the one-sided negative and positive RPS components, the relative signed-timing measure RPS^{diff}, and the corresponding magnitude-weighted variants. The vector $F_t$contains standard equity factors and reversal controls. Equation (37) is estimated as a nested sequence of seven models. The first model is an intercept-only regression, so that α^k is equal to the raw mean return of the RPS factor. Subsequent models add, in sequence, the market factor, Fama-French size and value factors, momentum, short-term reversal, and long-term reversal. This design provides a regression funnel: it tests whether the RPS factor return is progressively absorbed as standard risk and anomaly factors are added, or whether the intercept remains economically and statistically similar across specifications. Regressions are estimated using monthly factor returns, with Newey-West standard errors adjusted for heteroskedasticity and autocorrelation.

Portfolio sorts are complemented by stock-level predictive regressions. For each month t, the baseline Fama-MacBeth specification is

$$r_{i,t+1} = a_t + b_t Z\big(RPS_{i,t}^{diff,g}\big) + \gamma_t' X_{i,t} + \varepsilon_{i,t+1}, \qquad (38)$$

where $r_{i,t+1}$ is the next-month stock return, $Z(\cdot)$ denotes the cross-sectional standardisation of the RPS characteristic, and g indexes the signal construction (frequency and magnitude-weighted RPS). The coefficient b_t measures the cross-sectional relation between the RPS characteristic observed at the end of month $t$and returns in month $t+1$. The reported estimate is the time-series average of b_t, with SEs corrected for time-series dependence. The baseline specification includes only the RPS characteristic. Additional specifications then add standard competing predictors, in particular prior-month reversal and intermediate-term momentum. This mirrors the time-series regression funnel, but tests the same question at the stock level: whether RPS predicts next-month returns after controlling for established cross-sectional return predictors.

To compare one-sided and relative timing effects, the Fama-MacBeth regressions are also estimated using alternative RPS sets:

$$r_{i,t+1} = a_t + \delta_t' Z\left(\mathcal{R}_{i,t}^{g}\right) + \gamma_t' X_{i,t} + \varepsilon_{i,t+1}, \qquad (39)$$

where $\mathcal{R}_{i,t}^{g}$is chosen from the relevant RPS characteristics, including RPS^{-}, RPS^{+}, RPS^{diff}, the joint pair (RPS^{-}, RPS^{+}), and nested specifications combining RPS^{diff} with one of the one-sided components. Since $RPS^{diff} = RPS^{-} - RPS^{+}$, the empirical objective is not to include all three variables mechanically in every regression, but to test whether the predictive content is better described by one-sided signed timing or by the relative chronology of losses and gains. This structure separates the component-level interpretation from the combined path-timing specification used in the main factor-mimicking tests.

3.3.3. CRO robustness

Since chronological return ordering is the closest existing path-ordering benchmark to RPS, specific tests for robustness against it are undertaken. Reminding first that the CRO characteristic is defined as the correlation between daily returns and their chronological position within the formation window,

$$CRO_{i,t} = \text{corr}_{\tau \in W_t^{(H)}}\left(r_{i,\tau}, a_\tau\right), \qquad (40)$$

where a_τ denotes the age of the daily return observation, with larger values corresponding to older observations. The sign convention follows the original CRO construction. Unlike the measure RPS that we define, which first partitions daily returns by sign and then applies the timing operator, CRO summarises the entire signed return sequence through a single unconditional correlation. The two measures are therefore expected to be empirically related, but they need not contain the same information: CRO captures unconditional return ordering, whereas RPS captures sign-conditional and, in the magnitude-weighted specification, shock-weighted chronology.

The empirical analysis evaluates this relation in four steps. First, monthly cross-sectional correlations are reported between RPS^{-}, RPS^{+}, RPS^{diff}, their magnitude-weighted equivalents, CRO, reversal, momentum, terminal-window reversal, and extreme-return controls. Second, the baseline Fama-MacBeth predictive regression is augmented with CRO:

$$r_{i,t+1} = a_t + b_t RPS_{i,t}^{diff,g} + \theta_t CRO_{i,t} + \gamma_t' X_{i,t} + \varepsilon_{i,t+1}. \quad (41)$$

The coefficient b_t measures the incremental pricing content of RPS after controlling for the existing chronological-ordering measure, while $\theta_t$measures whether CRO retains explanatory power after the sign-conditional RPS statistic is included. Third, the same comparison is implemented at the factor level by augmenting the RPS factor-alpha regression with a CRO factor-mimicking return:

$$R_t^{RPS} = \alpha + \beta' F_t + \gamma R_t^{CRO} + u_t, \quad (42)$$

and symmetrically by regressing R_t^{CRO}on the standard factor set (from section 3.3.1) augmented with R_t^{RPS}. This tests whether the two factor returns are substitutes or whether RPS captures a distinct component of return-path information.

Finally, two non-parametric identification tests are conducted. The first residualises RPS with respect to CRO each month,

$$RPS_{i,t}^{diff,g} = c_t + \lambda_t CRO_{i,t} + \eta_{i,t}, \quad (43)$$

and forms decile portfolios on the residual component $\hat{\eta}_{i,t}$. The second uses conditional double sorts: stocks are first sorted into CRO groups and then sorted by RPS within each CRO group. The reverse analysis (sorting by CRO within RPS groups) is also reported. If RPS continues to generate return spreads within CRO groups, while CRO contributes less once RPS is fixed, this would indicate that the sign-conditional RPS construction is not merely a reparametrisation of unconditional chronological return ordering.

3.3.4. Parameter, horizon, universe, and regional robustness

Further robustness tests for all results, aside from factor-based ones, are especially aimed at distinguishing RPS from known return patterns and mechanical calendar effects.

1) Calendar effects are examined because the baseline signal is formed at calendar month-end. So we include calendar fixed effects, month-of-year checks, and alternative signal-construction windows (three and six month lookback horizons H) to assess whether signed path timing is a) concentrated at short horizons; b) remains informative over longer return histories; c) remains persistent irrespective of calendar day of formation (such as known effects – rebalancing windows, liquidity events, earnings, etc.). Additional tests compare the full-window RPS measure with variables based only on the final trading day or final week to assess whether the signal is merely a terminal-return proxy.
2) The universe definition is varied. The broad universe is compared with rank-by market-capitalisation and liquidity-restricted samples (Top 1000/500/300), allowing the analysis to distinguish a broad cross-sectional pricing pattern from a small-stock or microstructure-sensitive regularity. This also opens a first link to potential persistence in closer to real-world execution implementations.

3) Finally, the full methodology is replicated on European equities. The European analysis uses the same signal definitions, portfolio sorts, factor regressions, and cross-sectional tests (but from 1986 onwards as data permits). This provides an out-of-sample test of whether RPS captures a broader signed path-timing characteristic or a US-specific empirical regularity.

The methodology therefore isolates a specific empirical object: the pricing of signed return-path timing. Conventional characteristics such as reversal and momentum summarise cumulative performance over a formation period; RPS instead measures how losses and gains are positioned within that period. RPS^{-} and RPS^{+} capture the timing of negative and positive return days separately, while RPS^{diff} compares the two conditional chronologies and summarises the relative direction of the signed path. The resulting measure is reduced-form: it is compatible with a hypothesis of behavioural interpretations based on salience, but does not require the pricing relation to arise from investor memory alone. The empirical analysis therefore tests whether signed path salience is priced after controlling for conventional return predictors, factor exposures, universe restrictions, and calendar mechanics.

4. Results

4.1. Signal properties and cross-sectional variation

The U.S. baseline panel contains 1,841,703 stock-month observations spanning 479 calendar months from February 1985 to 2025, with an average of 3,845 stocks per month after data filters of Section 3.

The raw RPS measure is bounded mechanically by the recency exponent and the formation window: at α=0.05 and a one-month window, the theoretical range is [0, 0.141], with the corresponding bounds widening to [0, 0.187] at three months and [0, 0.215] at six months. This compression preserves the cross-sectional ranking but renders the numerical level of the signal incomparable across α values and lookback horizons, so all empirical tests below use a within-month z-score of the signal before the decile sort. Note that the rescaling is irrelevant for the rank-based sorts and Fama-MacBeth regressions that follow.

Cross-sectionally, the RPS characteristics are essentially orthogonal to the conventional return predictors. In the U.S. sample, the time-series mean of the monthly cross-sectional correlations between RPS^{diff} (magnitude) and the standard reversal proxy is +0.03; between RPS^{diff} and twelve-month minus one-month momentum, -0.02; and between RPS^{diff} and the maximum daily return over the formation window, +0.00 (see Figure 1). The picture is the same for RPS^{-} and RPS^{+} taken separately, and across the magnitude- and frequency-weighted variants. The signal therefore does not duplicate any well-known characteristic in the cross-section; whatever it predicts in the next month is not predicted by the level of past returns, the medium-term momentum tilt, or the salience of extreme daily observations. The one substantive non-zero correlation in Figure 1 is between RPS^{diff} and the chronological return-ordering measure CRO (Mohrschladt, 2021): the time-series mean cross-sectional correlation is -0.84 for the magnitude-weighted variant and -0.62 for the frequency-weighted variant, with the negative sign reflecting our opposite orientation convention adopted in the construction. Dedicated identification tests in Section 4.5 are designed to establish whether the residual cross-sectional content of RPS after CRO is partialled out remains priced.

The following analysis section focuses on the magnitude-weighted RPS measures since, as summarised in Table 1, the measures always outperform the frequency-weighted counterparts.

4.2. Signed components and the combined measure

The two one-sided components measure different conditional timing objects and, empirically, they are not mechanical mirror images of one another. Figure 2 documents the asymmetry directly. Panel A reports the full-sample annualised Sharpe ratio of the L/S spread on each component: in the U.S., RPS^{-} delivers a Sharpe of 1.42 against 1.53 for RPS^{+}, and the corresponding European Sharpes are 0.76 and 1.14. The positive-side signal is therefore at least as strong as the negative-side signal in raw univariate terms, in both markets, and the asymmetry is somewhat more pronounced in Europe than in the U.S. Panels B and C plot the 12-month rolling absolute monthly mean of each component for the U.S. and European panels respectively: the two components share a broadly similar level but their rolling profiles do not move in lockstep, with extended periods (the late 1990s and mid-2010s in the U.S., the early-to-mid-2010s in Europe) in which the two series visibly decouple. This is the empirical analogue of Equation (11) in the methodology, which established that the negative- and positive-side components do not sum to a constant and need not move mechanically in opposite directions. Panel D plots the rolling 36-month Sharpe of RPS^{-}, RPS^{+} and RPS^{diff} in the U.S.: the combined characteristic tracks above both one-sided components for the bulk of the sample, with the wedge between RPS^{diff} and the better one-sided component widening in the 1995-2000 and 2008-2012 episodes, and the three series converge in the post-2018 low-Sharpe environment.

The decile profile (all reported in Figure 3) of RPS^{-} in the U.S. sample is monotone and economically large. Mean next-month returns decline almost linearly from the D1 decile (stocks whose negative-return mass is concentrated near the end of the formation month) to the D10 decile (stocks whose negative-return mass is older), with a D1-D10 spread of 1.0% per month under the magnitude-weighted construction. The decile profile of RPS^{+} is symmetric in shape but oriented in the opposite direction: stocks with stale gains (high RPS^{+}) outperform stocks with recent gains (low RPS^{+}), with the high-minus-low spread on the magnitude version also close to 1.0% per month. The patterns are monotonic in both signs.

The cross-sectional information content of the two components is independent in the joint Fama-MacBeth regression. The bivariate structure motivates the combined characteristic $RPS^{diff} = RPS^{-}$-RPS^{+} , which compresses both pieces of information into a single scalar that simultaneously rewards recent losses and stale gains. The results indicate that gain-side and loss-side timing capture largely separate information.

A direct test of the combined signal against its components is the joint Fama-MacBeth regression with RPS^{diff} and RPS^{-} on the right-hand side. In the U.S. one-month panel, the RPS^{diff} coefficient is -0.402 (NW-t=-9.3) while the RPS^{-} coefficient collapses to -0.030 (NW-t=-1.0). The combined measure therefore subsumes the loss-side component at the one-month horizon. At three- and six-month horizons the result is more nuanced: the RPS^{diff} coefficient strengthens to -0.533 (NW-t=-11.6) at three months and -0.489 (NW-t=-11.8) at six months, while RPS^{-} flips sign and becomes positive at +0.139 and +0.158 (both significant at the 1% level). The natural reading is that at longer horizons the two components carry partially independent information that the simple difference cannot capture. Taken together, RPS^{diff} (magnitude) offers the most parsimonious summary of the signed path-timing content -combining both sides of the asymmetry containing different information into a single interpretable scalar while delivering the highest univariate Sharpe of the three variants - and is adopted as the primary characteristic for the remainder of all Section 4.

4.3. Baseline factor-mimicking portfolio returns

The L/S decile portfolios on the three signal families produce large and statistically significant return spreads in the U.S. baseline. Table 1 reports the full panel of mean monthly returns, annualised volatilities, Sharpe ratios, and Newey-West t-statistics for the negative-side, positive-side, and combined signals in both the frequency and magnitude-weighted variants, at one-, three- and six-month lookback windows. The U.S. one-month figures for the magnitude-weighted variants are 1.025% per month for RPS^{-} (Sharpe 1.43, NW-t=+7.6), 1.002% per month for RPS^{+} (Sharpe 1.53, NW-t=+9.0) and 1.451% per month for RPS^{diff} (Sharpe 1.76, NW-t=+8.85). Maximum drawdowns of the cumulative L/S series over the full 479-month sample are -10.0%, -9.5%, and -11.3% respectively (Figure 4 plots the cumulative series performance).

A natural concern is whether the spread is being generated symmetrically by both legs of the portfolio or whether one side dominates. The raw long-leg of RPS^{diff} (the D1 decile, containing stocks with recent losses and stale gains) earns +20.1% annualised, while the raw short-leg (D10, stocks with recent gains and stale losses) earns +2.7%. A direct interpretation of these raw figures should be prudent because both legs are long-equity positions and both inherit the unconditional equity premium of roughly 8-9% per year over this sample; the only object that strips out the common time trend is the factor-model alpha of each leg separately. The summary alpha-of-spread results in Section 4.4 show that the M6 alpha of the L/S portfolio is +1.45% per month against a raw L/S return of +1.4% per month, indicating that the cross-sectional tilt is preserved almost intact under factor adjustment.

As mentioned above, the full decile profile is monotone for both signed components and for the combined characteristic across all three lookback horizons, although the monotonicity weakens at the long horizons for RPS^{-} alone, where the slope of the decile profile flattens. Subperiod analysis shows that the RPS^{diff} spread is positive and significant at the 5% level in every decade-long subperiod: pre-2000 Sharpe 2.89, 2000-2009 Sharpe 1.71, 2010-2017 Sharpe 0.96, 2018+ Sharpe 1.01 (all U.S. one-month, see Figure 5). The same subperiod cut applied to RPS^{-} alone is materially weaker in the most recent decade, with the 2018+ Sharpe falling to 0.60 (NW-t=+1.7); the combined characteristic is therefore more stable to the late-sample period than the loss-side component alone. The pattern is consistent with a regime story rather than an arbitrage-erosion story: the signal is strongest in periods of high cross-sectional return dispersion (the dot-com unwind and 2008 crisis) and weakest during the low-volatility 2010-2017 expansion, with partial recovery thereafter.

4.4. Factor controls and Fama-MacBeth identification

Table 2 reports the raw mean and time-series factor alphas of the magnitude-variant L/S portfolios for the three headline RPS signals - RPS^{-}, RPS^{+} and RPS^{diff} - at one-, three-, and six-month formation horizons, under six nested asset-pricing specifications running from the raw spread through CAPM, Fama-French 3-factor, Carhart 4-factor (Carhart, 1997), the FF3+short-term reversal augmentation, and the fully augmented FF3+UMD+ST_Rev. The final long-term reversal extension (M7) is reported in Table 3 alongside the factor loadings discussed below, with the same qualitative conclusions. For the U.S. one-month RPS^{diff} baseline, the raw mean of +1.451% per month and the fully adjusted M6 alpha of +1.447% per month are essentially identical - a compression of under 1ppt, significant at the 1% level. RPS^{diff} remains highly significant and economically large at the three- and six-month horizons, with M6 alphas of +1.450% (NW-t=+9.0) and +1.259% (NW-t=+8.2) respectively; the compression from raw to M6 grows modestly with horizon but the spread never loses statistical significance. The one-sided variants tell a horizon-dependent story already noted in Section 4.2: RPS^{+} delivers stable factor-adjusted alphas across all three horizons, while RPS^{-} flattens to insignificance at six months. The RPS^{diff} spread is therefore not materially explained by the standard market, size, value, momentum, or reversal exposures across the one- to six-month horizon range.

From Table 3, for the U.S. one-month signal in the full universe, the market beta is small and weakly positive (β=+0.06, t=+1.9), SMB and UMD loadings are statistically zero, HML is mildly negative (β=-0.16, t=-2.1), the short-term reversal loading is statistically insignificant (β=-0.11, t=-1.2), and the only sizeable significant loading is on long-term reversal (β=+0.21, t=+2.8). In the restricted universes the ST_Rev loading becomes meaningfully negative, reaching β between -0.34 and -0.36 (t between -3.0 and -3.3) across Top-1000/-500/-300 - so among institutionally accessible stocks the signal appears more contrarian to the conventional reversal premium; this pattern is taken up in Section 4.7. The M7 R^2 rises modestly from 6.9% in the full universe to between 9.1% and 11.0% in the restricted samples, confirming that standard factors explain only a limited share of the monthly variation in the RPS factor return.

The factor-model tests use traded portfolios as controls; the Fama-MacBeth regression tests whether RPS^{diff} subsumes known stock-level characteristics directly. With standardised RPS^{diff} magnitude as the only regressor, the U.S. one-month coefficient is -0.425 (×100, NW-t=-9.2). Adding the prior-month return as a stock-level reversal control changes the coefficient by less than 3bps to -0.423 (NW-t=-9.9). Adding 12-month minus one-month momentum yields -0.417 (NW-t=-9.7), and adding both controls jointly yields -0.413 (NW-t=-10.4); the t-statistic on RPS^{diff} rises in the joint specification because the controls absorb idiosyncratic variation in the dependent variable without absorbing the signal itself. Both control variables enter with the expected signs and standard magnitudes (reversal coefficient negative; momentum coefficient positive). RPS therefore prices the cross-section in the presence of, rather than instead of, the existing stock-level predictors.

A non-parametric counterpart to the Fama-MacBeth result above is the conditional sort within reversal quintiles. The RPS^{diff} L/S spread, formed independently within each prior-month-return quintile, eliminates by construction any common variation with the reversal characteristic. Consistent with the FM finding that the RPS^{diff} coefficient barely shifts when prior-month return is added as a stock-level control, the conditional spread is significantly positive in all five quintiles in the U.S. full universe: 22.9% annualised in Q1 (worst prior return), 15.5% in Q2, 12.5% in Q3, 13.9% in Q4, and 18.8% in Q5 (best prior return), with all NW-t statistics above 6.5 (Table 4). The same exercise yields significant spreads in all five quintiles in the Top-1000 universe and in four of five (including the highest-prior-return bucket) in the Top-500. The signal is therefore not the within-month component of reversal: it prices stocks just as strongly among recent winners as among recent losers, confirming through this distribution-free design the result the stock-level regression already showed.

4.5. CRO identification

The closer empirical question is whether RPS^{diff} is functionally redundant with the existing path-ordering characteristic, chronological return ordering (Mohrschladt, 2021; Cakici and Zaremba, 2023). The two measures are mechanically related but not identical. In the U.S. sample, the time-series mean of the monthly cross-sectional correlation between CRO and RPS^{diff} is -0.84 (the negative sign follows the opposite orientation convention: high CRO = stale gains and recent losses, while low RPS^{diff}= recent losses and stale gains; up to sign, the two measures share approx. 70% of their cross-sectional variation). Their respective correlations with the conventional reversal proxy are +0.03 for CRO and +0.04 for RPS^{diff}, so both are orthogonal to reversal in the same way and to roughly the same degree. The question is therefore whether the residual 30% of the cross-sectional content of RPS^{diff} matters for pricing.

The first test is residualisation. Each month, the cross-section of RPS^{diff} is regressed on the cross-section of CRO, and decile portfolios are formed on the residual. The average monthly cross-sectional beta of RPS on CRO is 0.846 in the full universe (0.854 in both concentrated universes), confirming a tight but incomplete mechanical link. After orthogonalisation, the residualised RPS^{diff} L/S spread retains 65% of its original raw return in the full universe (11.4% annualised against 17.4%, compounded 11.6% vs 18.3%, Sharpe 1.27 vs 1.76, NW-t=+7.16), 85% in the Top-1000 universe (8.7% vs 10.3%, compounded 8.4% vs 9.9%), and 92% in the Top-500 universe (8.0% vs 8.6%, compounded 7.5% vs 8.0%). The residualised signal is significant at the 1% level in every universe. To further isolate the incremental component, a complementary series is constructed as the cross-sectional difference in standardised ranks: $\left(RPS_z^{diff} - CRO_z\right)$. This simple z-score differential, which directly extracts the signed conditional timing content of RPS net of the unconditional chronological ordering, earns 6.9% annualised in the full universe (SR=0.81, NW-t=+5.01), 6.1% in Top-1000 (SR=0.54, NW-t=+2.96), and 5.7% in Top-500 (SR=0.47, NW-t=+2.89). Taken together, RPS^{diff} therefore prices a non-trivial component of the cross-section that is orthogonal to chronological return ordering, and this is tradeable and statistically reliable across the market cap spectrum.

The second test is the incremental factor alpha. Augmenting the M6 regression of the RPS^{diff} spread with the CRO L/S return as an additional explanatory variable reduces the alpha from +1.447% per month to +0.560% per month, but it remains highly significant (NW-t=+6.15); the CRO factor loading is +0.81 (t=+15.5). And the symmetric exercise - regressing the CRO spread on the standard factors plus the RPS^{diff} L/S - eliminates the CRO alpha as we see the intercept drop from +1.101% per month (NW-t=+6.92) to -0.090% per month (NW-t=-1.12), a complete absorption of the CRO premium by the RPS^{diff} factor. This result holds in narrower universes: in Top-1000 the residual RPS^{diff} alpha after controlling for CRO is +0.280% (NW-t=+2.97), with CRO alpha after controlling for RPS^{diff} at -0.026% (NW-t=-0.26); in Top-500 the figures are +0.242% (NW-t=+2.84) versus +0.028% (NW-t=+0.27). In every universe tested, RPS^{diff} spans CRO but the reverse is not true, confirming that RPS contains information that CRO does not.

The third and final test is the double sort. Conditioning first on CRO quintile and then sorting on RPS^{diff} within each CRO bin produces a significant RPS^{diff} L/S in all five CRO quintiles, with annualised spreads between 8.3% and 11.6% and NW-t statistics between +5.7 and +7.5. The reverse procedure - conditioning first on RPS^{diff} and sorting CRO within bin - yields CRO spreads that are statistically indistinguishable from zero in four of the five RPS^{diff} quintiles and marginally negative in the fifth (Table 5).

Read together, the three identification tests indicate that RPS^{diff} prices a distinct cross-sectional dimension that CRO does not capture, whereas the converse cannot be sustained; the sign-conditional decomposition of the return path carries pricing information beyond what the unconditional chronological return correlation provides.

4.6. Calendar and formation-day robustness

The signal is formed at the calendar month-end in the baseline, so a natural concern is that month-end microstructure such as turn-of-the-month liquidity effects, settlement-driven flows, or window-dressing, could explain or account for part of the spread. The relevant test rolls the formation date across the trading days of the month and computes a one-month forward Sharpe ratio at every day-of-month bin. For the U.S. one-month signal, the within-month Sharpe profile is +0.87 in trading days 1-5, +0.63 in days 6-10, +0.53 in days 11-15, +0.80 in days 16-20, +0.98 in days 21-25, and +1.44 in days 26-31 (Figure 7). The end-of-month bin is the strongest, with an increment of +0.68 over the mean of the earlier five bins; but the signal is statistically significant in every formation bin (each at NW-t above 3.0), and the implied Sharpe in the earliest bin is comparable to that of many well-established equity factors.

The implication is that the path-timing effect is present throughout the month and is not a creature of the calendar month-end. At the same time, the one-month bin does carry a genuine month-end premium that the data permit; this premium is much smaller at longer formation windows. The within-month Sharpe profile for the U.S. three-month signal is essentially flat between days 1-5 (+1.19) and days 21-25 (+1.15), with an end-of-month bin of +1.52 - a month-end increment of only +0.40, and for the six-month signal the corresponding increment is +0.27. The interpretation is that the baseline one-month signal mixes a path-timing effect with a modest end-of-month microstructure effect, while the longer-window signals capture only the former. What can be concluded from this profile is that the path-timing component is robust to within-month formation date; what cannot be concluded is that the one-month spread contains no microstructural element at all.

The lookback-horizon comparison (Figure 8) shows that the combined RPS^{diff} signal is robust across formation windows from two weeks to twelve months: the U.S. mean Sharpe averaged across all 23 within-month formation days is 0.88 at one month, 1.21 at three months, 1.10 at six months, and 0.81 at twelve months. It is worth noting that the negative-side component RPS^{-} degrades sharply at long horizons (Sharpe falls from 0.96 at one month to 0.28 at six months and turns negative at twelve months in the U.S.), and the positive-side component RPS^{+} improves slightly with horizon (0.62 at one month, 0.91 at three months, 0.95 at six months). This pattern reinforces the view that the two components are not mechanical mirror images: loss timing and gain timing display different horizon sensitivities. One possible interpretation is that recent losses are most informative over shorter windows, while gain timing may require a longer history to separate stale from recent positive performance. Formally, although the theoretical range of the $q_{\alpha}(a)$ kernel widens with the lookback length, the marginal effect of any single recent signed observation on the sign-conditional average is diluted when the window contains more observations. The combined RPS^{diff} signal therefore provides a more stable summary of relative signed chronology and supports its use as baseline specification across formation windows.

4.7. Universe and investability

The full U.S. cross-section is dominated by small and microcap stocks. To address the standard concern that path-ordering anomalies may be concentrated in the least liquid corner of the market (Cosemans and Frehen, 2021; Cakici and Zaremba, 2022), as shown by some results discussed earlier, the entire pipeline is rerun on 3 liquidity-restricted universes defined by lagged month-end market capitalisation: the Top-1000, the Top-500, and the Top-300 U.S. firms by market value. Figure 9 presents the full panel of factor-adjusted alphas for the one-, three-, and six-month signals across the four universes.

The picture is monotone but not destructive. The U.S. one-month M6 alpha on RPS^{diff} (magnitude) falls from +1.447% per month in the full universe (NW-t=+8.78) to +0.918% in Top-1000 (NW-t=+5.50), +0.870% in Top-500 (NW-t=+5.22), and +0.825% in Top-300 (NW-t=+4.99). The compression is approximately a halving of the alpha between Full and Top-500, but every restricted universe retains a NW t-statistic >4.5 in the one-month specification. The same pattern holds at the three- and six-month horizons.

Three additional patterns are noteworthy. First, the ST_Rev loading of the RPS^{diff} spread becomes meaningfully negative in the restricted universes (β between -0.34 and -0.36 in Top-1000/500/300, t around -3), whereas in the full universe it is statistically indistinguishable from zero (β=-0.11, t=-1.10). The interpretation is that in the cleaner sample, the residual covariation of the RPS spread with the conventional reversal factor is mildly negative - the strategy is more contrarian to the reversal premium among large firms than the full-universe spread suggests. Second, the within-reversal-quintile test continues to identify a significant RPS^{diff} spread in the highest reversal quintile (Q5 = best prior returners) in every universe: +12.7% annualised (t=+4.2) in Top-1000, +15.1% (t=+5.1) in Top-500, +14.3% (t=+4.6) in Top-300. The signal is therefore not just a small-cap pricing effect dressed up in a path-timing label; among institutionally accessible large stocks, it continues to identify a tilt that is independent of recent return level. Third, the relative ranking of the three signal variants changes across universes: in the full universe the difference signal dominates, but in the Top-500 and Top-300 universes the positive-side signal RPS^{+} has the highest Sharpe (Sharpe 0.68 and 0.65 compared to 0.63 and 0.61 for RPS^{diff}). The stale-gains effect is more robust to universe restriction than the fresh-losses effect or the combined characteristic.

4.8. European replication

As an external validation, the entire baseline U.S. pipeline is replicated on the European common-equity panel of 1.83 million stock-months spanning 1985-2025. The U.S. and European panels are constructed identically. The European RPS^{diff} (magnitude) L/S delivers +0.83% per month at the one-month formation horizon (Sharpe 1.38, NW-t=+8.40), +1.12% per month at three months (Sharpe 1.62, NW-t=+9.33), and +1.06% per month at six months (Sharpe 1.59, NW-t=+9.49). The pre-2000 European subsample is comparatively weak (one-month Sharpe 0.81), but each of the three later subperiods is significant at the 1% level, including the most recent (2018+ Sharpe 1.50, NW-t=+4.7). The component asymmetry documented in Section 4.2 also carries over: the positive-side component delivers a higher univariate Sharpe than the negative-side (1.14 vs 0.76, Figure 2 Panel A), if anything more sharply than in the U.S. The European magnitudes are roughly half the U.S. magnitudes, consistent with the international heterogeneity in behavioural anomalies documented by Cakici and Zaremba (2023).

The qualitative cross-sectional patterns documented on the U.S. panel carry over to Europe: the joint Fama-MacBeth regression with both one-sided components yields coefficients of -0.142 on RPS^{-} (NW-t=-5.1) and +0.209 on RPS^{+} (NW-t=+8.0), with the same near-identity to their univariate counterparts as in the U.S., confirming that the two timing components carry independent cross-sectional information on the European panel as well. The European evidence is sufficient to conclude that the central RPS^{diff} result is not a U.S.-specific empirical regularity.

5. Discussion

The empirical evidence presented in Section 4 establishes that the chronological ordering of signed return shocks, within a recent formation window, is a meaningful dimension of the cross-section of equity returns. Our signal is moreover independent of the conventional risk factors and of the existing unconditional path-ordering benchmark. This section interprets that evidence, situates it within the theoretical literature reviewed in Section 2, and registers the limitations of the design and the most promising directions for future work.

5.1. Interpretation of the empirical evidence

The central result of the thesis is that sign-conditional path timing carries pricing information that the cumulative-return characteristics in standard factor models do not capture, and that closest competing path-ordering measure does not span. Three features of this finding deserve interpretive emphasis.

First, the robustness of the spread to factor adjustment is quite clean. Across all six signal variants (the frequency and magnitude-weighted constructions) and all three formation horizons, the maximum compression from the raw L/S return to the M6 factor alpha is below 5%. Such pattern appears to be from a characteristic that is genuinely orthogonal to the joint span of market, size, value, momentum, and short-term reversal - not the signature of a small abnormal return that survives because of low statistical power. The conventional reading is that the cross-sectional information in the signal is not redundant with the standard tradeable risk factors. Based on the results, we do not claim more than this; in particular, we do not claim that the spread cannot be absorbed by some yet-to-be-discovered traded factor, only that it is not absorbed by the empirically established and most documented ones that are tested in this paper.

Second, the identification tests of our signal against CRO are asymmetric and informative. Indeed, the sign-conditional decomposition spans the unconditional chronological ordering - the RPS difference factor reduces the CRO alpha from highly significant to indistinguishable from zero - while the converse does not hold. The 30% of cross-sectional variation in the difference signal that is orthogonal to CRO continues to deliver a statistically significant and economically meaningful spread in every universe tested (referring to Figure 6). The natural reading is that there is information in whether losses occurred recently and whether gains occurred recently, jointly, that an unconditional correlation cannot resolve, and that this information is priced. This is the principal contribution of the thesis to the path-ordering literature initiated by Mohrschladt (2021) and extended internationally by Cakici and Zaremba (2023): not that path ordering is priced which was already established but that the relevant or augmented statistic is sign-conditional rather than unconditional.

Third, the within-reversal-quintile tests show that the spread is not the within-month component of short-term reversal. The signal prices the cross-section just as strongly among the best recent performers as among the worst, including in the universe-restricted samples where it remains significant but compresses. Tying it together with the orthogonality of the signal to the standard reversal proxy at the cross-sectional level, the evidence is that the path-timing characteristic identifies a tilt in the cross-section that operates on a different dimension from the level of past returns.

What the evidence does not allow is the stronger conclusion that the observed mispricing is caused by investor memory or salience-driven attention. The empirical design measures the cross-sectional pricing of a path characteristic and not investor beliefs or attention directly. The signal could be priced because of recency-weighted memory, because of salience-driven extrapolation, because of investor inattention to old information of either sign, because of slow-moving mispricing that happens to manifest in path-timing form, or because of some combination of these channels with non-behavioural arbitrage-cost asymmetries. RPS is framed throughout as a reduced-form characteristic, and that framing is the appropriate one given the evidence: the signal is robust, but the precise mechanism behind its pricing content remains an open question that survey or experimental work would be better positioned to address than a cross-sectional asset-pricing test.

5.2. Integration with theoretical frameworks

The evidence is consistent with predictions that flow from the salience-theoretic framework of Bordalo, Gennaioli and Shleifer (2012, 2013) and the recency-bias literature on extrapolation in expectations (Greenwood and Shleifer, 2014; Da et al., 2021). The natural mapping is that recent signed shocks are more cognitively available than older shocks, and that within a stock's return history the chronological position of large losses and large gains is informative about how investors are likely to be currently perceiving the security. Under this reading, the long leg of the RPS difference portfolio, that is stocks with recent losses and stale gains, collects securities whose return distributions, as currently weighted by investors, are dominated by adverse recent observations relative to the salience their fundamentals would otherwise lead (and the short leg collects the mirror case). The cross-sectional spread is then the compensation as the path-driven distortion in perceptions corrects.

One feature of the data deepens but also nuances this interpretation. The first is the subperiod regime pattern: the RPS difference spread is strongest in the pre-2000 period and during 2000-2009, which are periods that include the dot-com unwind and the GFC, both characterised by elevated cross-sectional return dispersion and large signed daily shocks. And the spread is materially weaker in 2010-2017, a period of low realised volatility and unusually persistent equity-market trending. The signal recovers partially in the post-2018 period that contains the COVID dislocations and the 2022 rate shock. If the pricing of within-window path timing depends on the presence of large signed shocks that differentiate stocks in the cross-section, then a regime of compressed dispersion provides less raw material for the signal to operate on. This is consistent with a salience or attention-amplification reading - large signed shocks are precisely what salience theory predicts will be over-weighted in investor perception - but it is also consistent with simpler explanations in which the spread reflects cross-sectional dispersion in path characteristics that is mechanically smaller in low-volatility regimes. The data don’t separate the two, and so the thesis doesn’t either.

Finally, the contribution of the thesis to the literature on cross-sectional salience signals deserves explicit framing. From the studies examined, the salience-level measure of Cosemans and Frehen (2021) and its international extension by Cakici and Zaremba (2022) is concentrated in small and illiquid firms, posing a fundamental challenge to its economic relevance at scale. The path-timing signal documented in this thesis is materially more robust to universe restriction: the M6 alpha attenuates by less than half but remains highly significant in every liquidity-restricted U.S. cross-section examined, and the same qualitative pattern holds in the European replication. This is not a contradiction of the salience literature - it is consistent with the view that a signal built on the within-stock chronology of signed shocks captures a different and more broadly priced dimension than one built on the cross-sectional salience of cumulative returns. The two should perhaps be read as complements rather than substitutes.

5.3. Limitations and future research directions

The contribution of this thesis is bounded by a few design choices and questions to still explore given the scope of the paper. Some of the more important of these are discussed below.

The framework does not allow a final conclusion on mechanism. The empirical object is the cross-sectional pricing of a path characteristic; the thesis does not measure investor beliefs, attention allocation, or memory weighting directly. The earlier interpretation in Section 5.2 is offered as the most natural reading of the empirical patterns, but a behavioural mechanism is not formally identified by the tests reported. The most direct extensions would combine RPS with survey measures of return expectations, in the spirit of Greenwood and Shleifer (2014) and Da et al. (2021), or use high-frequency attention proxies, such as search volume, news intensity, or retail-flow measures, to test whether the RPS spread is stronger precisely when investor attention is more likely to be concentrated on recent signed shocks or sequential information.

The reversal controls are limited to the standard one-month cumulative return benchmark and intermediate-term momentum. Faster reversal speeds, such as a 5-day prior-window cumulative return and exponentially decayed past-return controls with short half-lives, have not been included in the factor specifications due to the scope. To the extent that the RPS spread captures a component of higher-frequency microstructure-driven reversal that operates at daily or weekly frequencies, the current M7 controls may not fully absorb it. Adding more granular reversal controls to the Fama-MacBeth specifications is a natural close extension of the identification strategy.

Further work could explore the conditioning on fundamental information or flow/liquidity events, periods that influence investor salience. Whether the RPS spread is concentrated in stock-months whose formation window contains an earnings announcement (some period during which new fundamental information arrives, price reactions occur, and investor misinterpretation is most plausible) compared to stock-months without such events is something that may shed light on explainability behind the observed anomaly. Similarly, whether RPS^{-} is amplified among stocks that experienced a recent earnings surprise, or whether the signal decays differently following fundamental-information shocks versus pure liquidity events, would directly test whether the path-timing effect is tied to the arrival and processing of hard information. Conditional analysis around corporate earnings announcements or liquidity/flow events may help identify the economic source of the signal.

One major extension is exploring different asset classes than equities. The behavioural literature on salience, recency, and extrapolation does not imply that path-dependent belief formation should stop at the equity boundary. Fixed-income, FX, commodity, and crypto-asset cross-sections are natural settings for an RPS-style measure, particularly because return histories in these markets are often interpreted through macro narratives, trend extrapolation, and anchoring to recent shocks. Such an extension would also discipline the interpretation of the equity results. If sign-conditional path timing were priced in asset classes where conventional reversal and momentum effects operate at different horizons, the evidence would be harder to attribute to a purely mechanical or equity-specific phenomenon. It would also clarify how the short-horizon, stock-level recency channel studied here relates to longer-horizon behavioural mechanisms, such as the experiential-learning channel of Malmendier and Nagel (2011, 2016), in which investors' lifetime macroeconomic experiences shape expectations over much longer horizons.

Each of these directions extends rather than overturns the central conclusion of the thesis: that the sign-conditional chronology of recent return shocks is a distinct and broadly priced dimension of the equity cross-section.

6. Conclusion

We examine whether the chronological ordering of signed return shocks within a recent formation window is priced in the equity cross-section. The central finding is that it is. Stocks whose recent return paths are characterised by fresh losses relative to stale gains subsequently outperform stocks whose recent paths are characterised by fresh gains relative to stale losses. The pattern is economically large, statistically robust across a standard panel of identification tests, and not absorbed by the cross-sectional return predictors most likely to be mechanically related to it. A return history contains information not only in its cumulative level but also in the order in which positive and negative shocks occur.

The negative-side and positive-side RPS components are distinct empirical objects rather than mechanical mirror images, with independent cross-sectional information content in joint FM regressions and horizon profiles that run in opposite directions. The combined characteristic, RPS^{diff}, compares the two conditional chronologies directly. In the full U.S. stock universe baseline, the magnitude-weighted RPS^{diff} spread delivers approx. 1.45% per month at the one-month horizon, with a Sharpe ratio of 1.76 and a Newey-West t-statistic of 8.85; the result holds at 3- and 6-month horizons and the decile profiles are monotonic in both signs.

The RPS^{diff} spread is essentially unchanged after factor adjustment and stock-level FM regressions confirm the same conclusion with prior-month reversal and intermediate-term momentum as controls. Conditional sorts within reversal quintiles further establish that RPS prices both recent losers and recent winners. So RPS measures the chronology of signed shocks, not the level of past returns.

The deeper contribution comes from the comparison with chronological return ordering. CRO already establishes that within-window ordering is priced; the relevant question is whether sign-conditioning adds anything beyond that unconditional path statistic. The evidence is unambiguous and asymmetric: controlling for CRO reduces but does not eliminate the RPS alpha, while controlling for RPS eliminates the CRO alpha. The sign-conditional decomposition spans the unconditional ordering measure while the reverse fails. Our contribution to the existing literature extends beyond general return timing to establish that the relative chronological positioning of losses and gains is the true priced object.

The robustness exercises qualify the result. The concentration around month-end suggests that calendar-related trading dynamics may contribute to the effect, although they do not explain it away. Universe restriction monotonically attenuates the alpha across the Top 1000/500/300 cross-sections but it remains highly significant in each, with compression driven by the short leg. The European replication delivers the same qualitative pattern with magnitudes roughly half the U.S.

The evidence is consistent with a behavioural interpretation based on recency, salience, and asymmetric attention to losses and gains, under which investors overweight recent signed shocks when forming beliefs about a stock, generating temporary mispricings that subsequently correct. However it does not claim to prove this mechanism is unique. The evidence is therefore best interpreted as reduced-form: signed path timing is priced, and the pattern is consistent with behavioural mechanisms, although the design does not separately identify beliefs, attention, memory, or arbitrage frictions as unique channel.

These findings have academic and practical implications. For empirical asset pricing, they suggest that summarising recent performance through cumulative returns alone leaves out a meaningful dimension of cross-sectional information. For portfolio construction, RPS is most useful as a short-horizon overlay on slower-moving multi-factor strategies rather than as a stand-alone strategy. Natural extensions include a broader set of factor controls such as alternative reversal speeds, profitability and investment factors, liquidity proxies, and direct attention measures; and event-conditional analysis around earnings announcements to identify whether the path-timing effect is concentrated near information arrivals. Most importantly, applying the RPS construction beyond equities to fixed income, FX, commodities, and crypto-asset cross-sections would test whether sign-conditional path timing is an equity-specific regularity or a broader feature of asset prices, and would bridge the micro-level evidence presented here with the macro-level evidence on behavioural anchoring in long-duration assets.

To conclude, the sequence of returns matters. A stock's recent history is not fully summarised by how much it went up or down, nor by whether returns generally improved or deteriorated over time; what also matters is the relative chronology of losses and gains. By isolating this sign-conditional dimension, RPS captures a return-predictive characteristic that standard factors and controls and unconditional path-ordering measures do not fully absorb.

7. References

Barber, B.M., Odean, T., 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. Review of Financial Studies 21, 785-818.

Bordalo, P., Gennaioli, N., Shleifer, A., 2012. Salience theory of choice under risk. Quarterly Journal of Economics 127, 1243-1285.

Bordalo, P., Gennaioli, N., Shleifer, A., 2013. Salience and asset prices. American Economic Review 103, 623-628.

Cakici, N., Zaremba, A., 2022. Salience theory and the cross-section of stock returns: International and further evidence. Journal of Financial Economics 146, 689-725.

Cakici, N., Zaremba, A., 2023. Recency bias and the cross-section of international stock returns. Journal of International Financial Markets, Institutions and Money 84, 101738.

Carhart, M.M., 1997. On persistence in mutual fund performance. Journal of Finance 52, 57-82.

Center for Research in Security Prices, 2026. CRSP daily stock file. Wharton Research Data Services, University of Chicago.

Cosemans, M., Frehen, R., 2021. Salience theory and stock prices: Empirical evidence. Journal of Financial Economics 140, 460-483.

Da, Z., Engelberg, J., Gao, P., 2011. In search of attention. Journal of Finance 66, 1461-1499.

Da, Z., Huang, X., Jin, L.J., 2021. Extrapolative beliefs in the cross-section: What can we learn from the crowds? Journal of Financial Economics 140, 175-196.

Fama, E.F., 1970. Efficient capital markets: A review of theory and empirical work. Journal of Finance 25, 383-417.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3-56.

Kenneth R. French Data Library, 2026. Tuck School of Business at Dartmouth College. Accessed March 2026. http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. Review of Financial Studies 27, 714-746.

Hogarth, R.M., Einhorn, H.J., 1992. Order effects in belief updating: The belief-adjustment model. Cognitive Psychology 24, 1–55.

Jegadeesh, N., 1990. Evidence of predictable behavior of security returns. Journal of Finance 45, 881-898.

Lehmann, B.N., 1990. Fads, martingales, and market efficiency. Quarterly Journal of Economics 105, 1-28.

Liu, X., Wang, Y., 2025. Motivated memory bias. SSRN working paper No. 5566638.

Malmendier, U., Nagel, S., 2011. Depression babies: Do macroeconomic experiences affect risk taking? Quarterly Journal of Economics 126, 373-416.

Malmendier, U., Nagel, S., 2016. Learning from inflation experiences. Quarterly Journal of Economics 131, 53-87.

Mohrschladt, H., 2021. The ordering of historical returns and the cross-section of subsequent returns. Journal of Banking and Finance 125, 106064.

Murdock, B.B., Jr., 1962. The serial position effect of free recall. Journal of Experimental Psychology 64, 482-488. [3]

[3] LLMs were used for limited editorial assistance, including improving sentence clarity, structure, and readability, as well as for support in reviewing and debugging code.

8. Appendix/Figures/Tables

<u>Illustrative Examples for Methodology:</u>

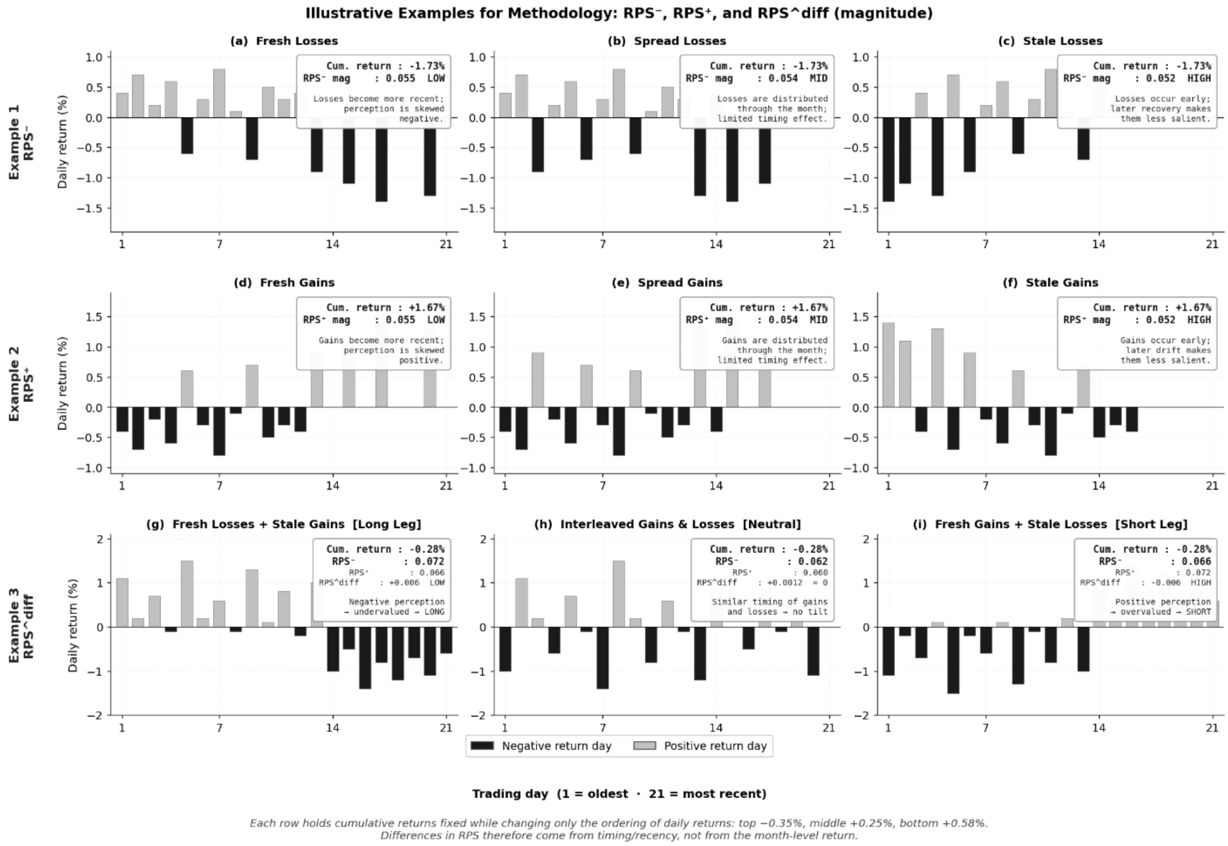

<u>Figure 1:</u>

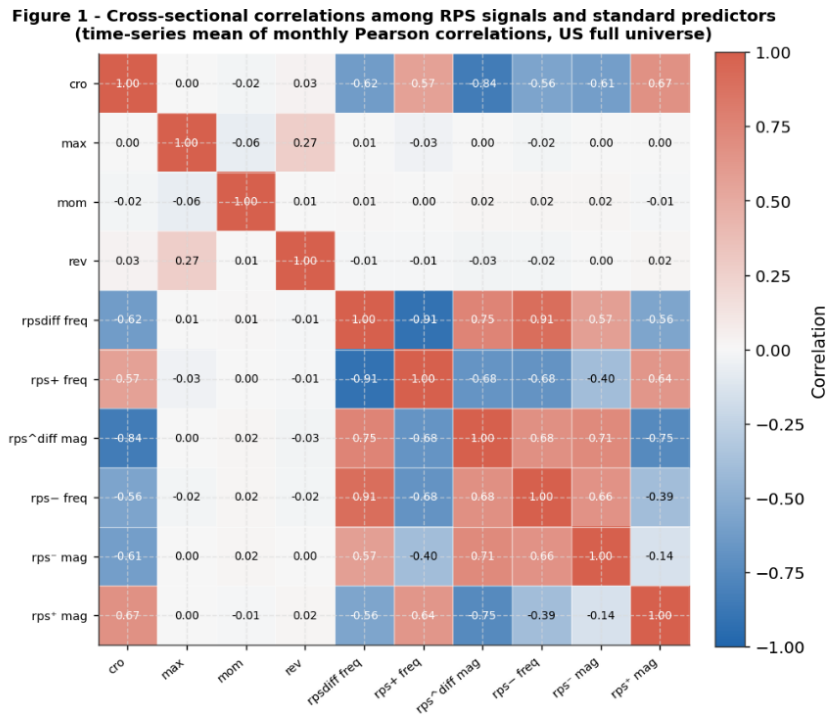

Figure 2:

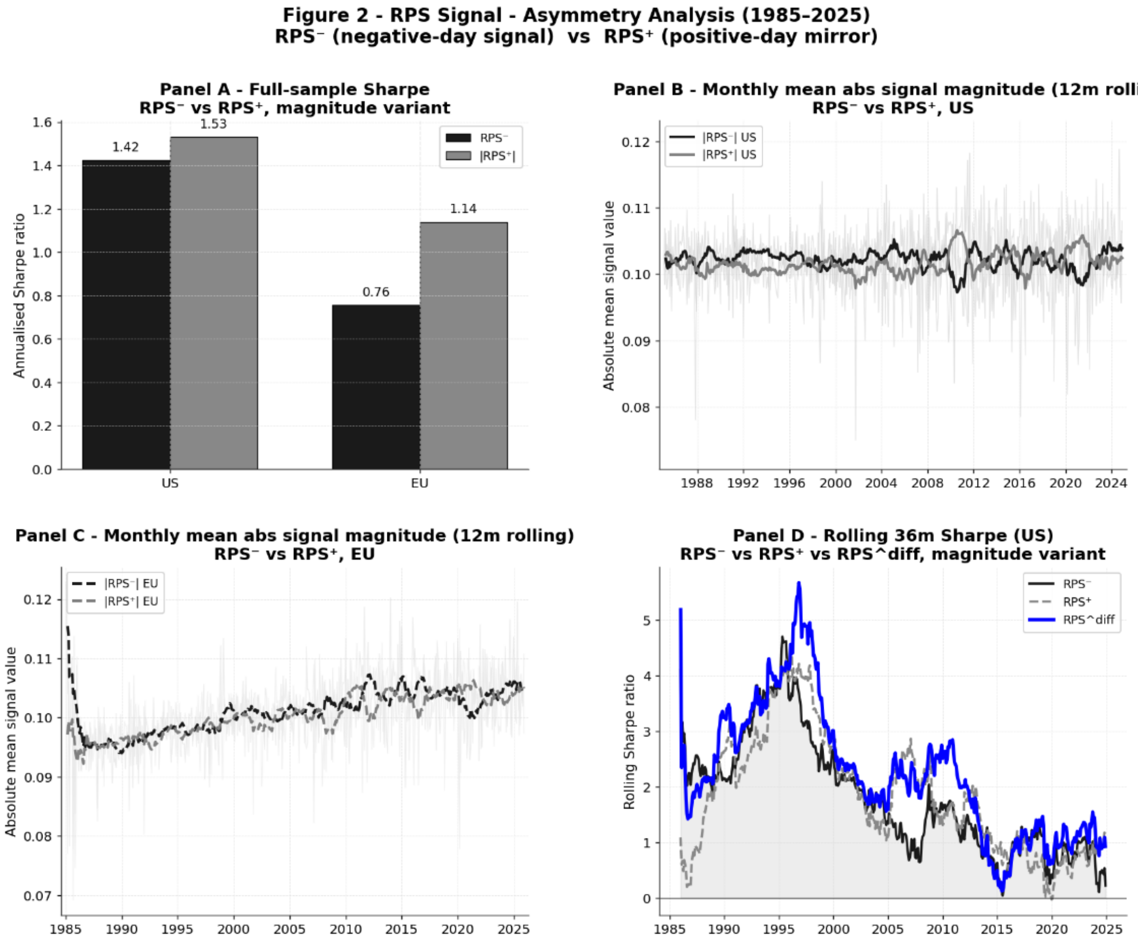

Figure 3:

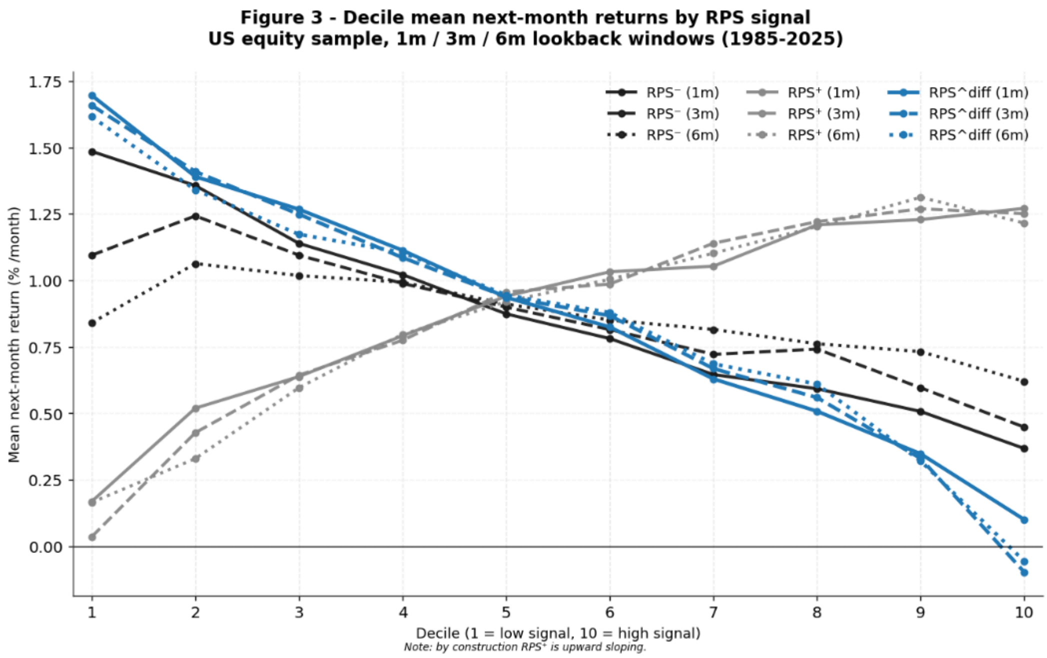

Table 1:

Table 1 – L/S performance summary (U.S.)

signal_lb	ann_%	vol_%	sharpe	nw_t	maxdd_%
RPS⁻ (mag) | 1-month	12.295	8.629	1.42	7.63	-9.981
RPS⁻ (mag) | 3-month	7.283	10.602	0.69	4.48	-42.471
RPS⁻ (mag) | 6-month	2.185	13.317	0.16	1.01	-73.352
RPS⁺ (mag) | 1-month	12.019	7.851	1.53	8.97	-9.471
RPS⁺ (mag) | 3-month	13.541	8.375	1.62	9.55	-10.351
RPS⁺ (mag) | 6-month	12.048	9.662	1.25	6.67	-34.272
RPS^diff (mag) | 1-month	17.415	9.887	1.76	8.85	-11.323
RPS^diff (mag) | 3-month	19.486	10.841	1.80	9.80	-21.229
RPS^diff (mag) | 6-month	18.407	11.816	1.56	9.50	-25.289
RPS⁻ (freq) | 1-month	10.358	8.664	1.20	6.96	-13.990
RPS⁻ (freq) | 3-month	7.590	10.758	0.71	4.24	-48.604
RPS⁻ (freq) | 6-month	3.214	13.173	0.24	1.45	-64.514
RPS⁺ (freq) | 1-month	11.632	7.051	1.65	10.15	-9.052
RPS⁺ (freq) | 3-month	12.099	8.103	1.49	8.34	-19.232
RPS⁺ (freq) | 6-month	10.472	10.196	1.03	5.09	-45.531
RPS^diff (freq) | 1-month	12.939	8.423	1.54	8.92	-10.085
RPS^diff (freq) | 3-month	14.926	9.210	1.62	9.64	-19.115
RPS^diff (freq) | 6-month	13.554	9.933	1.36	8.68	-22.859

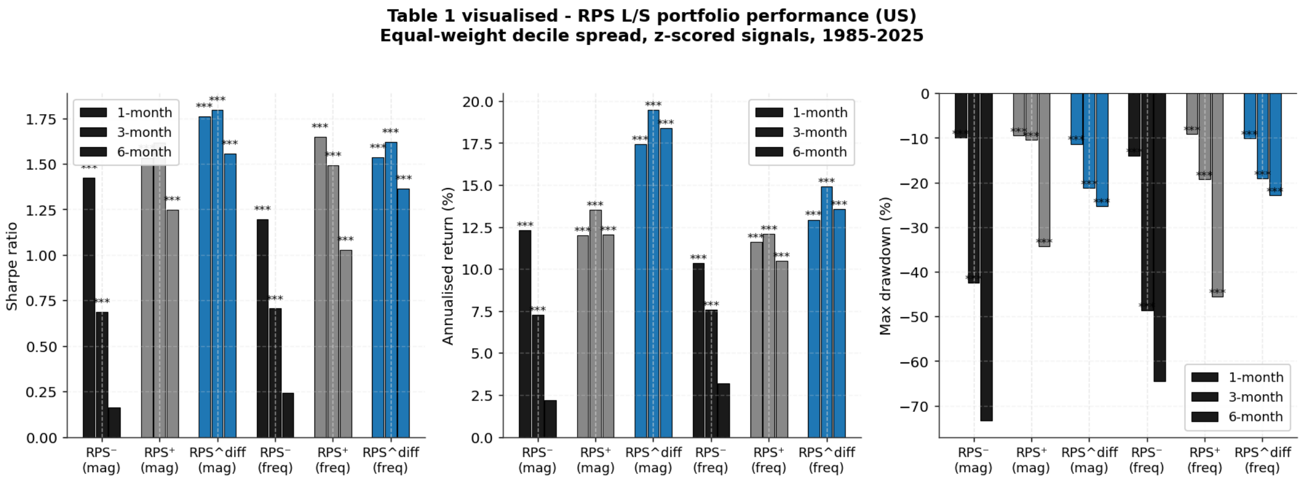

Figure 4.1:

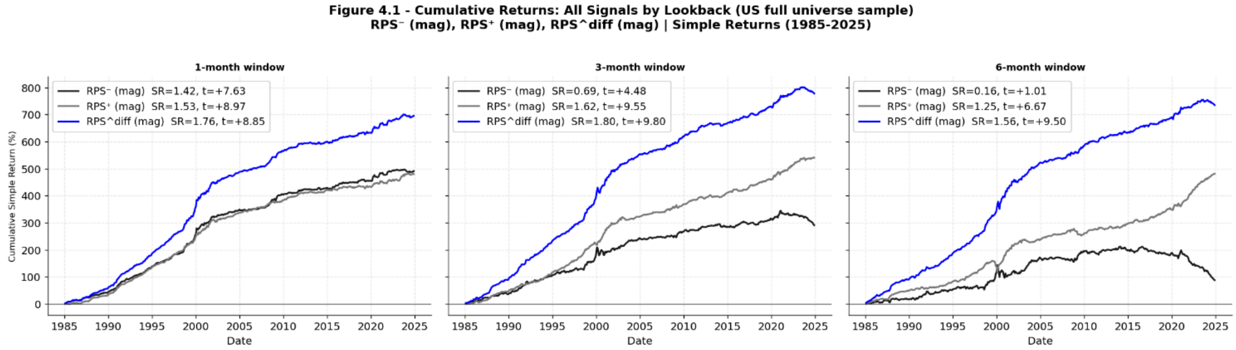

Figure 4.2:

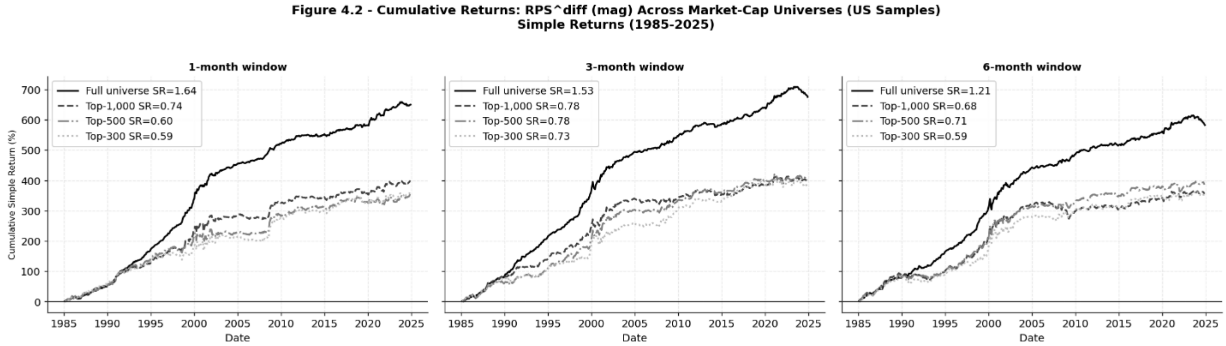

Figure 5.1:

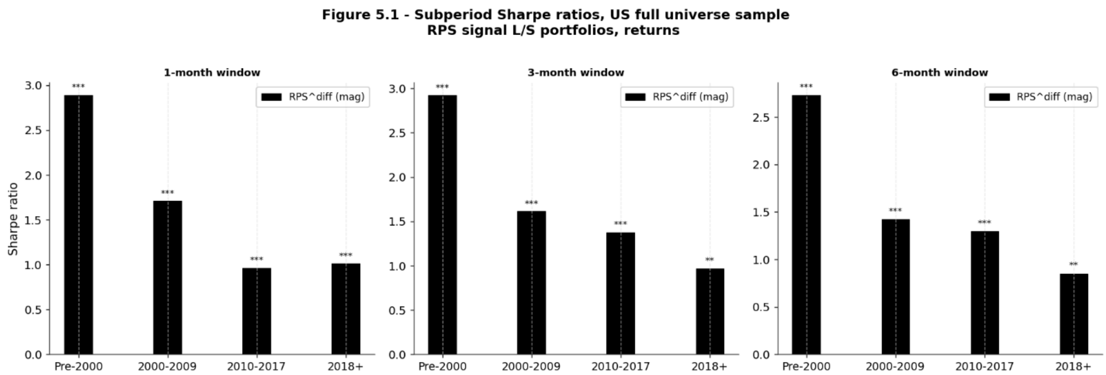

Figure 5.2:

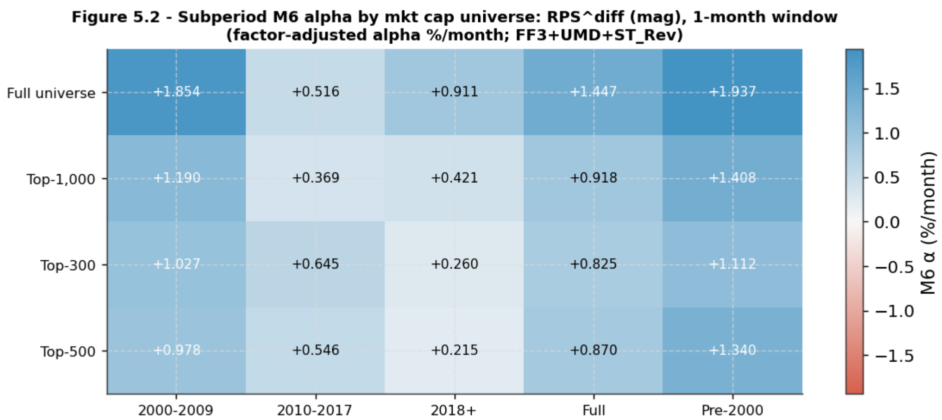

Table 2:

Table 2 – U.S. Factor alpha summary (Full U.S. universe: 1m, 3m, 6m)

signal	model	1m alpha_%m	1m t_alpha	3m alpha_%m	3m t_alpha	6m alpha_%m	6m t_alpha
RPS^{-} (mag)	Raw	+1.025	7.63	+0.607	4.48	+0.182	1.01
	CAPM	+0.949	7.20	+0.414	3.15	-0.123	-0.70
	FF3	+0.986	7.27	+0.482	3.89	-0.010	-0.07
	FF3+UMD	+1.002	7.67	+0.469	3.69	-0.019	-0.11
	FF3+ST_Rev	+0.986	7.19	+0.476	3.79	-0.018	-0.12
	FF3+UMD+ST_Rev	+1.001	7.45	+0.438	3.45	-0.057	-0.34
RPS^{+} (mag)	Raw	+1.002	8.97	+1.128	9.55	+1.004	6.67
	CAPM	+0.966	8.57	+1.175	10.01	+1.104	7.43
	FF3	+0.968	8.32	+1.143	10.19	+1.029	7.89
	FF3+UMD	+0.966	8.02	+1.129	9.30	+1.038	7.16
	FF3+ST_Rev	+0.970	8.34	+1.133	10.62	+1.019	8.77
	FF3+UMD+ST_Rev	+0.975	7.84	+1.084	9.56	+0.989	7.79
RPS^diff (mag)	Raw	+1.451	8.85	+1.624	9.80	+1.534	9.50
	CAPM	+1.396	8.62	+1.526	9.26	+1.357	8.31
	FF3	+1.422	8.43	+1.549	9.35	+1.402	8.50
	FF3+UMD	+1.431	8.98	+1.506	8.91	+1.345	7.65
	FF3+ST_Rev	+1.426	8.47	+1.538	9.31	+1.385	9.02
	FF3+UMD+ST_Rev	+1.447	8.78	+1.450	9.03	+1.259	8.15

Table 3:

Table 3- U.S. M7 Factor Loadings Summary

Period	universe	alpha_%m	t_alpha	r2	MktRF	MktRF_t	SMB	SMB_t	HML	HML_t	UMD	UMD_t	ST_Rev	ST_Rev_t	LT_Rev	LT_Rev_t
1m	Full universe	1.458	9.06	0.069	0.06	1.89	0.041	0.48	-0.162	-2.13	-0.044	-0.92	-0.11	-1.19	0.21	2.79
	Top-1,000	0.920	5.52	0.091	0.09	2.03	0.025	0.37	-0.125	-1.32	-0.086	-1.58	-0.34	-3.31	0.04	0.40
	Top-500	0.868	5.18	0.110	0.06	1.25	0.075	0.98	-0.168	-1.71	-0.157	-1.97	-0.36	-2.99	-0.04	-0.33
	Top-300	0.822	4.95	0.092	0.06	1.28	0.017	0.25	-0.126	-1.32	-0.136	-1.73	-0.35	-3.30	-0.05	-0.50
3m	Full universe	1.459	9.24	0.207	0.03	0.87	0.079	0.99	-0.086	-1.00	0.106	2.23	0.37	4.73	0.16	2.01
	Top-1,000	0.803	5.07	0.149	0.02	0.45	0.125	1.50	-0.099	-0.93	0.115	2.21	0.37	2.70	0.02	0.21
	Top-500	0.849	5.03	0.171	-0.02	-0.36	0.131	1.52	-0.161	-1.51	0.038	0.70	0.42	3.01	-0.01	-0.07
	Top-300	0.764	4.43	0.127	-0.01	-0.22	0.094	1.11	-0.156	-1.44	0.049	0.89	0.36	2.61	0.03	0.30

Table 4:

Table 4 – U.S. RPS^diff L/S spread within REV quintiles (1m, 3m, 6m)

universe	quintile	1m ann_%	1m sharpe	1m nw_t	3m ann_%	3m sharpe	3m nw_t	6m ann_%	6m sharpe	6m nw_t
Full universe	Q1(worst)	22.885	1.725	8.40	28.352	2.105	11.21	29.516	1.989	10.49
	Q2	15.526	1.537	9.18	20.005	1.798	9.83	16.326	1.417	8.34
	Q3	12.460	1.253	8.07	15.061	1.389	8.23	12.963	1.139	7.18
	Q4	13.906	1.247	7.51	16.900	1.406	8.65	14.027	1.144	7.52
	Q5(best)	18.764	1.359	6.72	16.277	1.159	6.76	13.868	0.918	5.82
Top-1000	Q1(worst)	11.640	0.690	4.42	9.762	0.556	3.58	17.556	0.925	5.75
	Q2	5.953	0.447	3.12	9.187	0.654	3.76	11.738	0.845	4.20
	Q3	8.957	0.676	4.26	11.033	0.840	4.95	6.038	0.444	2.57
	Q4	8.747	0.644	3.90	11.431	0.823	4.99	7.167	0.515	3.16
	Q5(best)	12.675	0.731	4.18	13.847	0.818	4.83	10.659	0.651	3.75
Top-500	Q1(worst)	11.106	0.546	3.74	8.193	0.387	2.41	11.363	0.554	3.73
	Q2	5.235	0.345	2.40	8.056	0.492	3.51	12.517	0.790	5.24
	Q3	6.449	0.409	2.73	9.652	0.640	4.52	6.819	0.436	2.49
	Q4	5.126	0.341	2.32	10.792	0.767	4.67	10.073	0.641	3.51
	Q5(best)	15.139	0.857	5.09	13.632	0.790	4.76	12.446	0.698	4.37
Top-300	Q1(worst)	14.004	0.642	5.18	9.038	0.384	2.43	13.327	0.572	4.06
	Q2	5.099	0.296	2.01	6.438	0.353	2.11	8.412	0.502	2.96
	Q3	5.585	0.333	2.16	9.400	0.549	3.46	7.717	0.422	2.17
	Q4	5.774	0.353	2.41	11.494	0.703	4.58	10.382	0.582	3.87
	Q5(best)	14.341	0.726	4.56	10.695	0.555	3.14	6.483	0.342	2.28

Table 5:

Table 5 – U.S. double Sort: RPS^diff and CRO

Panel	universe	base_quintile	sort	annualised_return_%	nw_t
	Full universe	Q1	RPS^diff	8.337	5.68
		Q2	RPS^diff	10.316	7.12
		Q3	RPS^diff	10.316	6.77
		Q4	RPS^diff	11.637	7.52
		Q5	RPS^diff	10.133	6.06
		Q1	RPS^diff	3.309	1.68

RPS^diff within CRO	Top-1000	Q2	RPS^diff	9.381	4.54
		Q3	RPS^diff	8.561	4.74
		Q4	RPS^diff	7.586	3.34
		Q5	RPS^diff	6.295	3.40
	Top-500	Q1	RPS^diff	5.517	2.47
		Q2	RPS^diff	6.546	3.53
		Q3	RPS^diff	8.355	3.89
		Q4	RPS^diff	6.303	2.73
		Q5	RPS^diff	5.367	2.74
CRO within RPS^diff	Full universe	Q1	CRO	-0.871	-0.75
		Q2	CRO	-4.596	-3.58
		Q3	CRO	-1.751	-1.33
		Q4	CRO	-1.095	-0.78
		Q5	CRO	3.551	2.92
	Top-1000	Q1	CRO	-1.528	-0.72
		Q2	CRO	-3.560	-1.68
		Q3	CRO	-1.366	-0.82
		Q4	CRO	-0.346	-0.18
		Q5	CRO	-0.821	-0.58
	Top-500	Q1	CRO	-2.425	-0.89
		Q2	CRO	-3.319	-1.73
		Q3	CRO	-1.566	-0.74
		Q4	CRO	-1.492	-0.70
		Q5	CRO	1.212	0.74

Figure 6:

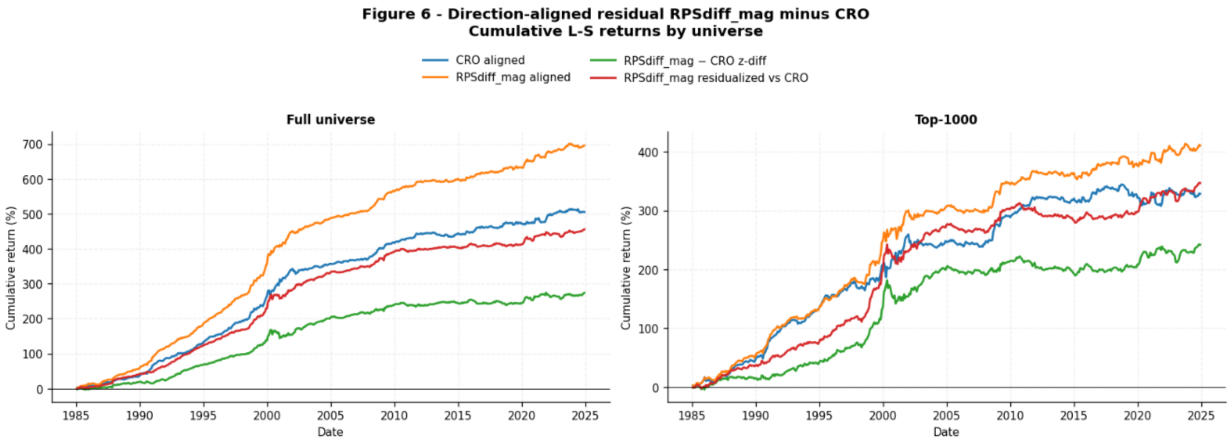

Figure 7:

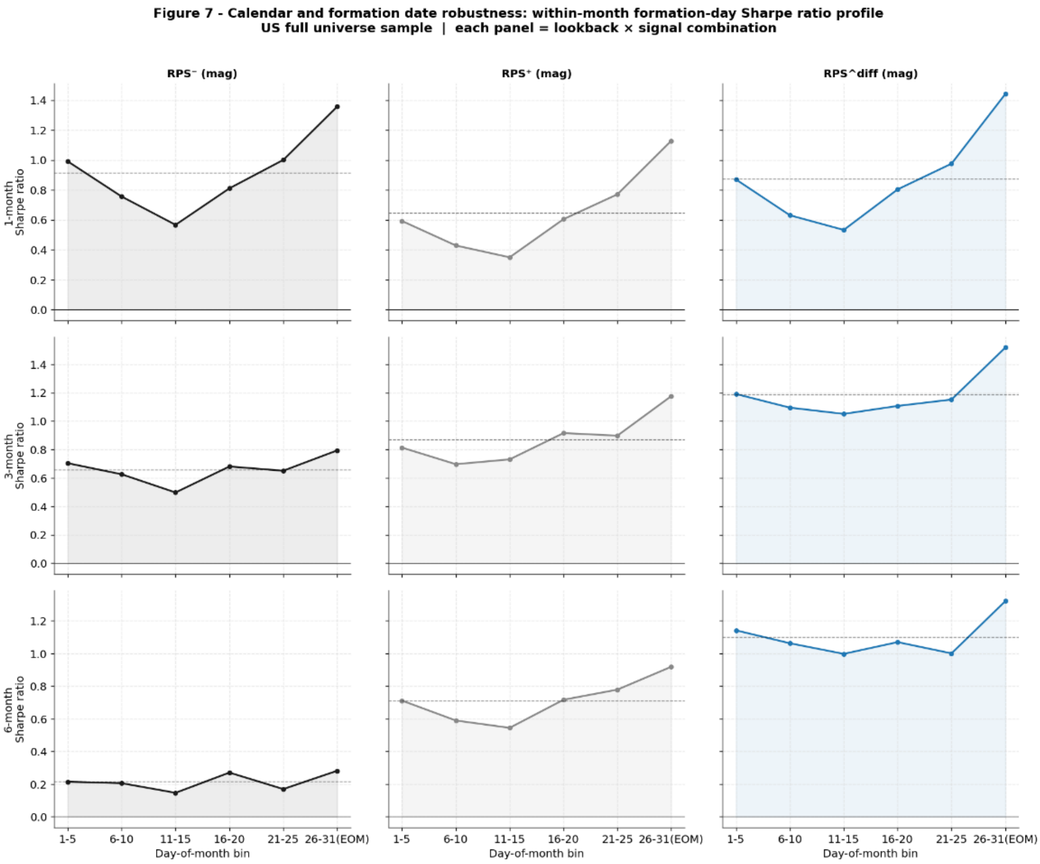

Figure 8:

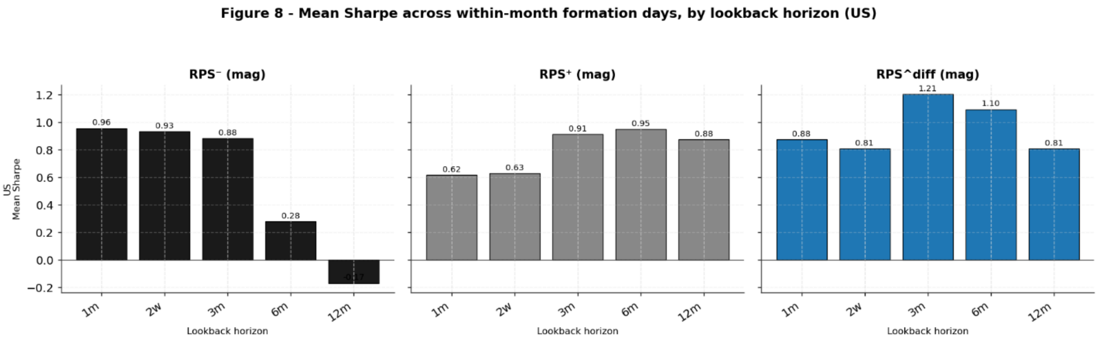

Figure 9: