Improving Performance with Fast Alphas

A Tactical Overlay for Intraday Trend Trading

Carlo Zarattini, Alberto Pagani

February 2026

1 Introduction

In quantitative research, practitioners frequently encounter a class of predictive signals operating at very short horizons, ranging from seconds to several minutes. We refer to these signals as *fast-decaying alphas* (or more simply *fast alphas*), defined as return predictors whose informational content exhibits a very short half-life[1] and whose statistical edge is concentrated in the immediate periods following signal formation.

Unlike medium- or long-horizon predictors, these signals often derive their apparent profitability from microstructure effects: examples include aggressive order-flow imbalances, liquidity dislocations or consecutive directional intraday moves, all observed in high-frequency data. In backtests that assume zero trading costs and no execution frictions, such signals can exhibit outstanding gross returns.

Their primary limitation, however, lies not in statistical validity but in economic implementability. Because *fast alphas* require frequent rebalancing and immediate execution, they generate elevated turnover. Once realistic transaction costs (bid–ask spreads, fees, and slippage) are incorporated, much of the gross edge typically dissipates. In many cases, the break-even cost threshold is sufficiently low that even minimal frictions render the strategy unprofitable. As a result, *fast alphas* often appear compelling in frictionless simulations but fail to survive practical execution constraints.

The key question, therefore, is whether such signals should simply be discarded.

We argue that the inability to monetize a signal directly does not imply the absence of economic value. A *fast alpha* may fail when implemented as a standalone trading strategy, yet still contain useful information about near-future price dynamics.

Our research suggests that successfully harvesting the information embedded in *fast alphas* requires a fundamental framework shift: rather than evaluating such signals in isolation, we find considerable benefits in studying them jointly with slower, more easily *monetizable alphas*.

Stylized Example

Suppose a trading strategy operating at a two-hour horizon generates a long entry signal. Historically, the strategy delivers an average gross return of 60 basis points per trade, with an average holding period of approximately 120 minutes.

Now assume that, at the time of entry, a *fast-decaying alpha* (whose predictive content is concentrated in the next minute) predicts a short-term market drift of -1 basis point. While such a signal may exhibit positive gross expectancy, its extremely short holding period implies prohibitive turnover, rendering it economically unviable as a standalone strategy under realistic market frictions.

Nevertheless, its informational content remains relevant. If we assume, for the sake of illustration, that the 60 basis points expected from the two-hour strategy accrue approximately uniformly over the 120-minute holding period, the expected return over the next minute becomes:

$$\mathbb{E}\left[R_{1\text{m}}\right] = \frac{60 \text{ bps}}{120 \text{ min}} = 0.5 \text{ bps/min} \quad (1)$$

Under the simplifying assumption that both the slow and fast signals provide independent forecasts for the

[1] We define the half-life of a predictive signal as the time required for its expected return to decay by 50% of its initial magnitude.

next one-minute return and are assigned equal informational weight, the combined one-minute expected return can be expressed as the simple average of the two conditional forecasts. If the slow two-hour strategy implies an average per-minute drift of 0.5 bps and the fast signal predicts a -1 bp return for the minute ahead, the combined conditional expectation amounts to:

$$\mathbb{E}\left[R_{1\text{m}}\right] = \frac{1}{2}(0.5) + \frac{1}{2}(-1.0) = -0.25 \text{ bps/min} \quad (2)$$

Although the two-hour trade exhibits positive unconditional expectancy, the short-term conditional expectancy becomes negative once both predictors are considered jointly. Entering immediately therefore implies accepting a negative expected drift over the next minute. In such a setting, delaying execution until the fast signal dissipates or aligns with the slower predictor may improve entry price without materially altering the structural two-hour expected return.

More generally, this example suggests that a short-lived signal can become valuable through its marginal informational contribution within a multi-horizon framework.

To formalize this view, we distinguish between *monetizable alpha*, defined as predictive content that survives realistic transaction costs when implemented directly, and *informational alpha*, defined as predictive content that, while not independently tradable, remains economically valuable when used to condition the execution or timing of other strategies.

The appropriate question, therefore, is not whether a signal is profitable after costs when traded independently, but whether it improves the implementation efficiency of another strategy operating at a different time scale.

In the following sections of this research note, we empirically validate the distinction between *monetizable* and *informational alpha*. We first construct a fast-decaying signal on SPY that exhibited a strong gross-of-fee performance between 2007 and 2026, but becomes economically unviable once trading commissions are incorporated.

We then introduce a baseline intraday trend-following model operating at a slower frequency. Finally, we show empirically that conditioning the execution of the intraday trend model on the *fast-decaying alpha* significantly enhances long-term performance. In particular, the fast signal acts as an execution overlay that aims to improve entry timing, reduce short-term adverse drift and increase risk-adjusted returns net of costs.

2 The Challenge of *Fast Alphas*

To illustrate the nature of *fast alphas*, we construct a simple high-frequency signal using 5-minute intraday bars of the SPY ETF. Our sample spans from January 2007 through January 2026. Intraday regular trading hours (RTH) data are sourced from IQFeed, while multi-day time series are obtained from Norgate Data.

We focus on a parsimonious and transparent specification designed to capture short-term mean-reversion effects. Specifically, we examine whether a streak of consecutive 5-minute bars in the same direction contains predictive information about the return of the subsequent 5-minute bar.

Let R_t denote the return of the 5-minute bar ending at time t. Define a streak of length N as a sequence of at least N consecutive bars sharing the same return sign. Conditional on the occurrence of such a streak at time t, we evaluate the average return of bar $t+1$.

In the simplest case $N = 1$, the signal reduces to a one-bar reversal rule:

- If the most recent 5-minute return is positive, a short position is taken for the next bar.
- If the most recent 5-minute return is negative, a long position is taken for the next bar.

Formally, the trading signal at the closure of bar t can be expressed as:

$$S_t = -\operatorname{sign}(R_t) \quad (3)$$

This specification intentionally represents an extreme instance of *fast alpha*: the signal may flip direction at the close of every bar, implying elevated turnover and immediate execution requirements. As such, it provides a natural stress test for the economic viability of ultra-short-horizon predictors. We begin by evaluating the statistical properties of this streak-based signal, abstracting from transaction costs.

Figure 1 reports the average return of bar $t+1$ conditional on streak lengths of at least N consecutive bars moving in the same direction. The results reveal a consistent mean-reversion pattern: following upward streaks, the expected return of the subsequent 5-minute bar is negative; following downward streaks, the expected return is positive.

Importantly, this behavior is observed across multiple values of N, suggesting that the effect is not confined to a particular streak length. The magnitude of the reversal increases moderately with longer streaks, although the economic size remains small in absolute terms.

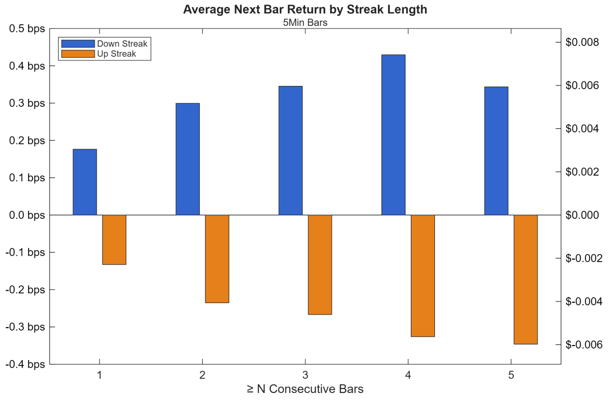

Figure 1: Average return (in bps and $/share) of bar *t+1* following a streak of at least *N* consecutive bars in the same direction, illustrating the presence of a mean-reversion effect.

As an illustrative example, after four consecutive positive bars, the average return of the next 5-minute interval is approximately -0.33 basis points, indicating a modest pullback. Conversely, after four or more consecutive negative bars, the subsequent return averages roughly $+0.43$ basis points, reflecting a short-term rebound.

These findings are consistent with microstructure-driven mean-reversion effects commonly observed at very short horizons. While the predictive edge is economically small, it appears statistically persistent across the sample period.

Turnover and Economic Fragility

From a practitioner's perspective, the key concern is not statistical predictability but economic implementability.

Even a visual inspection of the average next-bar dollar PnL per share presented in Table 1 suggests that the expected return per trade is modest relative to typical trading costs. Generally, when expected returns are measured in fractions of a basis point, execution frictions become a first-order consideration.

Table 1: Different specifications of the streak-based signal. Average next-bar returns (in bps and $/share) and trading-activity statistics are grouped by streak length and direction.

N ≥	Streak	Trades / Day	Next-Bar Ret (bps)	Next-Bar PnL ($)	Daily Turnover
1	Down	39.32	0.18	0.0038	39.3x
2	Down	18.33	0.30	0.0065	18.3x
3	Down	8.24	0.35	0.0061	8.2x
4	Down	3.65	0.43	0.0077	3.7x
5	Down	1.58	0.34	0.0025	1.6x
1	Up	39.46	−0.13	−0.0030	39.5x
2	Up	19.26	−0.24	−0.0044	19.3x
3	Up	9.05	−0.27	−0.0054	9.0x
4	Up	4.19	−0.33	−0.0057	4.2x
5	Up	1.87	−0.35	−0.0047	1.9x

To evaluate economic viability more rigorously, we isolate the highest-frequency configuration ($N = 1$), which produces a binary signal capable of reversing position at the close of each 5-minute bar: given its extremely elevated turnover, it represents a conservative test of cost sensitivity.

We then incorporate the following transaction cost assumptions:

- **Commissions:** Interactive Brokers (IBKR) tiered rate of $0.0035 per share.
- **Regulatory fees:** Sell-side costs doubled to approximate SEC-related charges.
- **Slippage:** No slippage assumed.
- **Other constraints:** Fractional share trading permitted; no minimum commission threshold imposed.

Position sizing targets full notional exposure. At the start of each trading day, the number of shares traded per signal is determined using the session's opening price:

$$\text{Shares}_t = \frac{\text{AUM}_{t-1}}{\text{Session Open}_t} \tag{4}$$

This quantity remains fixed throughout the entire trading session.

Gross Versus Net Performance

A gross-of-fee simulation of the $N = 1$ specification yields striking results. Over the sample period, the strategy produces a CAGR (*Compound Annual Growth Rate*) of approximately 31.9% and a Sharpe ratio exceeding 2. These figures illustrate that even extremely simple high-frequency predictors can display substantial gross performance when transaction costs are ignored.

However, the inclusion of realistic commission costs dramatically alters the outcome. As shown in Figure 2, once standard IBKR commissions are applied, the strategy's net performance becomes negative. This outcome reveals a structural fragility: the signal's expected per-trade return is of the same order of magnitude as per-trade execution costs: as a result, even modest transaction costs are sufficient to erode profitability.

Given that our signal does not survive as a standalone strategy, it may be better interpreted as *informational alpha* rather than *monetizable alpha* within the framework introduced in Section 1.

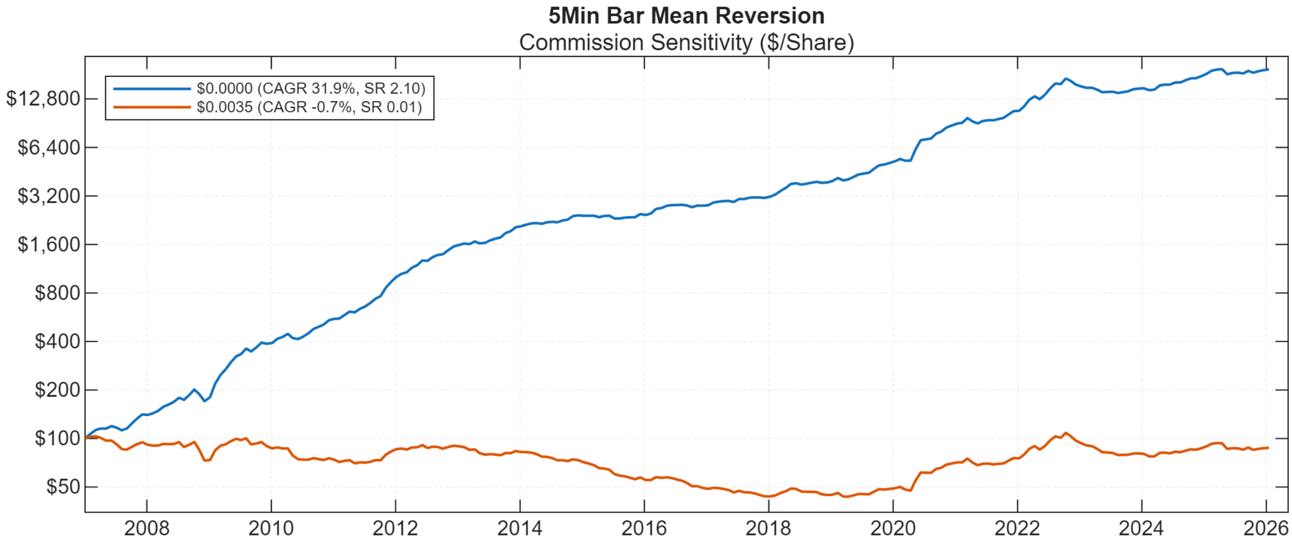

Figure 2: Sensitivity of the fast alpha signal ($N = 1$) to IBKR standard-tier commission rate. The signal uses the sign of the most recent bar's return as a mean-reversion cue, taking the opposite position on the following 5-minute bar. Despite strong gross-of-fee performance, transaction costs completely erode its profitability.

3 Baseline Intraday Trend Model

To test our *fast alpha* signal as execution overlay, we first require a robust and economically viable baseline strategy. For this purpose, we adopt a classic intraday breakout system inspired by *Trading Systems and Methods* (Kaufman, 2013).

The strategy leverages the Average True Range (ATR)[2] indicator to construct dynamic volatility bands around the session's opening price, allowing positions to be initiated in the direction of meaningful intraday price moves.

At the beginning of each trading session, we compute an upper and a lower volatility band. The upper band is defined as the session opening price plus one-half of the 14-period ATR, while the lower band is defined as the session opening price minus the same amount:

$$\begin{aligned} \text{Upper Band}_t &= \text{Session Open}_t + \frac{1}{2}\text{ATR}(14), \\ \text{Lower Band}_t &= \text{Session Open}_t - \frac{1}{2}\text{ATR}(14) \end{aligned} \tag{5}$$

These bands serve as dynamic breakout thresholds, illustrated in Figure 3 through a representative intraday trading session.

A long position is initiated whenever the price closes above the upper band, signaling an upward intraday breakout. Conversely, a short position is initiated when the price closes below the lower band, signaling a downside breakout.

[2]The Average True Range (ATR) is a volatility indicator introduced by Wilder (1978), computed as a moving average of the True Range over a specified lookback period. The True Range is defined as the maximum of: (i) the current high minus the current low, (ii) the absolute value of the current high minus the previous close, and (iii) the absolute value of the current low minus the previous close.

Figure 3: Illustrative example of a single trading session: the volatility bands (black) define the breakout thresholds, while the session open level (purple) acts as the stop price for open positions.

Once a position is opened, it is closed if the price returns to the session's opening level, which acts as a stop indicating that the breakout has failed. To eliminate overnight exposure, all open positions are flattened at the end of the trading day.

Following an approach similar to Zarattini et al. (2024), execution is restricted to 15-minute intervals (specifically at *HH:00*, *HH:15*, *HH:30*, and *HH:45*) in order to mitigate the risk of overtrading caused by immediate price fluctuations. Moreover, position sizing follows the same volatility-targeting framework, using the realized volatility of daily SPY returns over a 14-day lookback window to target a daily portfolio volatility of 2%.

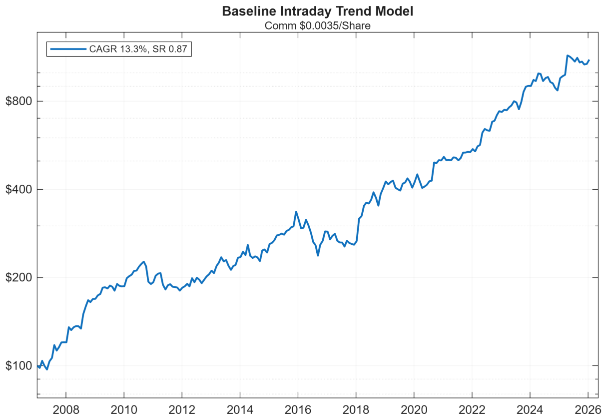

Figure 4: Evolution of a $100 investment in the baseline intraday trend strategy (performance shown is net of trading fees).

Importantly, once the target position size is determined at the session open, it remains fixed for the entire trading day. Each trade initiated during the session is executed using this predetermined quantity.

Despite its simplicity, the strategy demonstrates long-term robustness. As shown in Figure 4, a $100 initial investment grows steadily over time on a net-of-fee basis. Over the sample period, the strategy delivers a double-digit CAGR (*Compound Annual Growth Rate*) above 13% and a Sharpe ratio of approximately 0.87. Given its structural simplicity and absence of regime filters, these results are economically meaningful.

Within the terminology introduced in Section 1, this baseline framework qualifies as *monetizable alpha*: it survives realistic transaction costs and generates persistent net returns. It therefore provides a suitable foundation upon which the *fast alpha* execution overlay can be evaluated in the following section.

4 Enhancing Execution via *Fast Alpha*

We now integrate the 5-minute mean-reversion signal (introduced in Section 2) as a tactical execution overlay to the baseline intraday trend strategy described in Section 3.

The integration logic is conceptually straightforward. The baseline trend system determines directional exposure by identifying the dominant breakout direction, effectively providing a lower-frequency estimate of expected price direction. The *fast alpha*, in contrast, provides short-horizon information about immediate price drift.

Rather than executing immediately upon breakout confirmation, entry timing is conditioned on the *fast alpha*. Specifically, when a breakout is detected on the 15-minute time frame, execution is postponed until the *fast alpha* indicates a short-term pullback in the opposite direction. For long entries, this entails observing a negative 5-minute return prior to execution; for short entries, a positive 5-minute return is required.

Importantly, the *fast alpha* overlay is applied not only to entries but also to exits. In the baseline strategy, positions are closed when price reverts to the session opening level. Under the enhanced framework, liquidation is conditioned on the *fast alpha* signal: rather than exiting immediately at the point of maximum short-term dislocation, the strategy waits for a brief counter-move before closing the position. For example, if a long position reaches the stop level, the strategy delays liquidation until a 5-minute upward move is observed. The same logic applies symmetrically to short positions.

This approach does not alter the structural stop rule itself; rather, it refines the timing of execution around stop triggers. In other words, the objective is not to avoid necessary exits, but to improve exit price quality by conditioning liquidation on short-term counter-moves.

Figure 5 illustrates both the buy- and sell-side adjustments. Under the standard breakout rule, entry would occur at the close of the first bar exceeding the volatility band. With the overlay in place, execution is delayed until a micro-pullback materializes, potentially resulting in a more favorable entry price. The exit logic follows the same principle.

This overlay materially changes the execution profile of the baseline strategy. Traditional trend-following systems are inherently liquidity-consuming: they enter positions during expanding volatility and directional acceleration, often paying full bid-ask spread and accepting short-term adverse selection in exchange for immediate exposure.

By incorporating a mean-reversion overlay, the strategy may introduce elements of liquidity provision. Rather than competing aggressively for liquidity during breakouts or liquidating immediately at stop breaches, the strategy waits for temporary counter-moves. This shift may:

- Improve average entry prices,
- Enhance stop execution quality,
- Reduce exposure to microstructure noise.

Importantly, the overlay does not attempt to monetize the *fast alpha* directly. Instead, it uses short-horizon information to refine the conditional execution of a lower-frequency *monetizable alpha*.

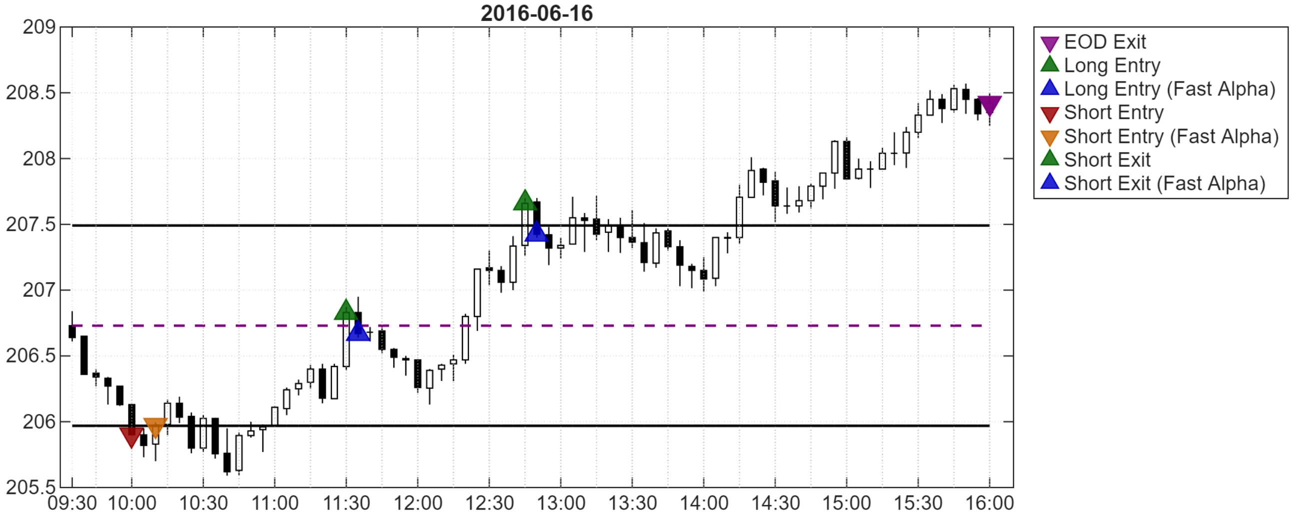

Figure 5: Illustrative example of execution enhancement using a *fast alpha* signal. Standard execution logic would force entry at the close of the first 15-minute breakout bar; instead, by waiting for a 5-minute bar in the opposite direction, the entry price is improved. The end-of-day (EOD) exit remains unaffected by this adjustment. Vertical grid lines indicate the 15-minute timestamps at which execution of the base strategy is permitted.

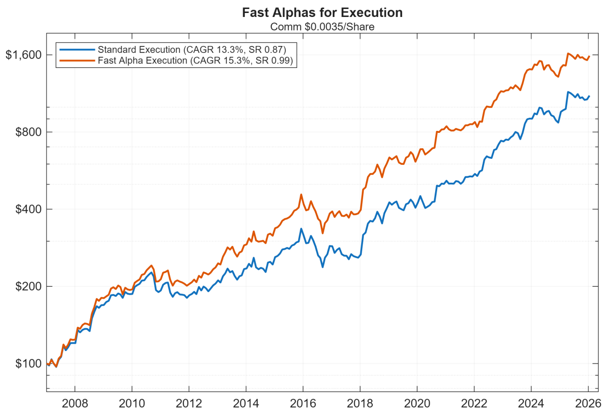

Figure 6: Comparison of the baseline intraday trend strategy and the version enhanced with fast-alpha execution logic.

Note — The fast alpha *overlay is applied symmetrically to both entries and exits. While this approach may occasionally delay participation in exceptionally strong trends, the pullback requirement is intentionally mild (a single opposite 5-minute bar), making missed trades relatively infrequent. More sophisticated urgency controls could be introduced to force execution after a predefined time window; however, such refinements fall outside the scope of this research note.*

To quantify the impact of the *fast alpha* overlay, we conduct a comparative simulation between the baseline intraday trend strategy and the enhanced variant incorporating conditional execution. The results, reported in Figure 6, indicate a meaningful improvement in net performance. The enhanced strategy increases net-of-fee CAGR (*Compound Annual Growth Rate*) by approximately 200 basis points and improves the Sharpe ratio from 0.87 to 0.99.

Crucially, the performance improvement remains observable after transaction costs. By conditioning the high-turnover *fast alpha* signal on the lower-frequency trend strategy, the incremental predictive value is captured without inheriting the prohibitive turnover associated with a standalone implementation.

These findings provide empirical support for the distinction introduced in Section 1: while the short-horizon signal does not qualify as *monetizable alpha* when traded independently, it contributes economically meaningful value when deployed as *informational alpha* within a multi-horizon framework.

5 Conclusion

This research highlights a central matter in quantitative trading: while predictive signals at very short horizons are often statistically detectable, their economic viability frequently collapses once realistic transaction costs are incorporated.

Our empirical analysis confirms that *fast alphas* may exhibit strong gross-of-fee performance yet fail to survive trading frictions when implemented as standalone strategies. Their expected per-trade return is typically of the same order of magnitude as execution costs, rendering direct monetization economically fragile.

However, the absence of standalone profitability does not imply the absence of economic value.

When integrated as execution overlays for robust lower-frequency strategies, *fast alphas* can enhance implementation efficiency, improve entry and exit price quality, and increase risk-adjusted performance net of costs. Rather than serving as independent return engines, they function as informational filters that refine the execution of *monetizable alpha*.

These findings suggest that evaluating predictive signals solely on their standalone net performance may understate their true economic relevance. Within a multi-horizon framework, signals can generate value not by being traded directly, but by improving how other strategies are traded.

In light of this evidence, we believe that the distinction between *monetizable alpha* and *informational alpha* provides a practical lens through which high-frequency predictors should be assessed.

References

Perry J Kaufman. *Trading Systems and Methods.* John Wiley & Sons, 5th edition, 2013.

Carlo Zarattini, Andrew Aziz, and Andrea Barbon. Beat the Market: An Effective Intraday Momentum Strategy for S&P500 ETF (SPY). *Available at SSRN 4824172*, 2024.

CONTACT US

Concretum SA
Viale Carlo Cattaneo 1
CH-6900 Lugano
Switzerland

info@concretumgroup.com
www.concretumgroup.com