Why Trade an Unpredictable FX Market? Intraday Evidence Based on Technical Trading Rules

Augusto César de Freitas Gurgel* Alex Luiz Ferreira†

Abstract

This paper asks whether high-frequency technical trading rules generate statistically robust and economically meaningful net returns in the Brazilian real/U.S. dollar futures market, and whether such profits can be reliably exploited *ex ante*. Using transaction-level data for U.S. dollar futures (`DOL`) traded on B3, we evaluate 31,986 candidate rules per horizon at five intraday aggregation levels under Stepwise Superior Predictive Ability inference, with explicit transaction costs and strict day-trading constraints. Significant rules peak at 774 at the 60-second horizon, with median annualized excess returns reaching 49.74% and Sharpe ratios exceeding 4 among selected rules. Rule significance is state-dependent, consistent with a model in which episodic predictability arises from a latent directional component that dissipates at longer horizons. Despite this in-sample evidence, fewer than 4% of significant rules retain significance out of sample, supporting a conditional-viability interpretation: profitable rules exist, but identifying and applying them *ex ante* remains elusive.

*University of São Paulo, Department of Economics, FEA-RP, Ribeirão Preto, Brazil, `augustocfg@usp.br`

†University of São Paulo, Department of Economics, FEA-RP, Ribeirão Preto, Brazil, `alexferreira@usp.br`

1 Introduction

Can short-horizon trading strategies systematically profit from exchange-rate movements? This question has become increasingly relevant in modern financial markets. Retail trading platforms and brokerage firms actively promote algorithmic and rule-based systems, and participation in intraday trading has increased rapidly in recent years.[1]

At the same time, a large academic literature shows that exchange rates are notoriously difficult to forecast, especially at short horizons. Even structural macroeconomic models rarely outperform a random-walk benchmark out-of-sample (Meese and Rogoff, 1983; Engel and West, 2005). This coexistence of widespread trading activity and limited predictability raises a natural question: do technical trading rules actually generate robust net profits once realistic trading costs and data-snooping issues are taken into account?

This paper examines whether high-frequency technical trading rules can generate statistically robust and economically meaningful net returns in the Brazilian real/U.S. dollar (BRL/USD) futures market. Our primary question is whether any technical rules deliver positive expected returns once transaction costs and the joint search across a large universe of candidate strategies are taken into account. Additional questions concern the structure of this profitability: which rule families drive performance, at which intraday aggregation horizons predictability concentrates, and under what market conditions trading rules become profitable or cease to be viable. At last, another important matter is whether the trading rules that deliver the highest in-sample returns are robust and remain persistent across different market conditions, particularly across days marked by heterogeneous intraday volatility and drift dynamics.

These questions relate to two strands of the literature. The first examines exchange-rate predictability. A large body of evidence finds that exchange rates are difficult to forecast, while research in FX microstructure emphasizes the role of order flow and trading behavior in shaping short-horizon exchange-rate dynamics (Evans and Lyons, 2002; Osler, 2003; Menkhoff and Taylor, 2007). The second strand evaluates trading-rule profitability directly. Early studies document significant returns from simple technical rules in equity markets (Brock et al., 1992), while subsequent work emphasizes the importance of correcting for data-snooping bias when evaluating large rule universes (Sullivan et al., 1999). More recent

[1] Public reports and exchange disclosures indicate that retail trading activity has expanded rapidly in Brazil. According to B3 statements reported in financial media, the number of individuals classified as day traders surpassed one million in 2024 (InfoMoney, 2024). The technological infrastructure supporting this activity has also expanded. Trading platforms widely used in Brazil, such as Nelogica's Profit platform, provide integrated modules for automated rule-based strategies, while brokerage firms including XP Investimentos and BTG Pactual regularly distribute intraday technical-analysis reports and automated trading tools to retail clients (XP Investimentos, 2024). For instance, BTG Pactual organized the 2025 Copa BTG Trader in a simulated trading environment, distributed R$1.25 million in prizes, and public information from BTG and related coverage indicates participation by more than 12 thousand traders. Source: `https://lp.btgpactual.com/copa-btg-trader`.

studies apply these approaches to foreign exchange markets and other asset classes (Hsu et al., 2016; Coakley et al., 2016; Ghanem et al., 2024). Much of this literature, however, relies on daily data and often treats transaction costs only partially, leaving open whether technical rules remain profitable at high intraday frequencies.

In the Brazilian context, the existing literature remains limited in its coverage of technical trading rules in emerging-market FX environments. Hsu et al. (2016) report only seven profitable rules when applied to the BRL/USD market. Although informative, these findings do not fully settle the broader question of whether profitable technical strategies may persist in Brazil under alternative market conditions or trading horizons. This issue is particularly relevant because the Brazilian FX market is both highly liquid and subject to active central bank intervention, two features that may materially affect price dynamics and, consequently, the performance of technical trading strategies (Kohlscheen and Andrade, 2014). Research on day trading in Brazil has also focused primarily on the outcomes of individual traders, mostly in equity markets, and has produced heterogeneous results. Using detailed transaction-level records from B3, Chague et al. (2020) show that the vast majority of day traders lose money. While these findings provide important evidence on trader-level outcomes, they do not directly address whether specific trading strategies may nonetheless generate positive returns under particular market conditions, nor whether a meaningful distinction should be drawn between retail day trading and high-frequency trading.

The BRL/USD futures market traded on B3 provides a particularly informative environment in which to study these issues. Currency derivatives concentrate a large share of trading activity in Brazil. Two U.S. dollar futures contracts are available in this segment, `WDO` and `DOL`. `WDO` is typically the most liquid contract, with higher trading volume, while `DOL` also records substantial daily turnover. Because `WDO` and `DOL` are tightly linked through arbitrage, their intraday price dynamics are closely connected, so evidence from `DOL` is informative about the broader BRL/USD futures environment. For example, according to the *Boletim Diário do Mercado* of B3 for Friday, March 13, 2026 (No. 49), `DOL` registered 32,506 trades and 318,570 contracts negotiated in a single trading session, corresponding to a financial volume of approximately R$84.4 billion (about US$16.2 billion).[2] Beyond liquidity, the BRL/USD market is characterized by high volatility, episodic central bank intervention, and heterogeneous participation, all of which are relevant for intraday trading-rule performance.

Our empirical analysis uses transaction-level records of U.S. dollar futures (`DOL`) traded on B3, aggregated into five intraday horizons ranging from 5 to 120 seconds. The sample covers 115 trading days and is divided into a 30-day in-sample period and an 85-day

[2]See B3, *Boletim Diário do Mercado*, Friday, March 13, 2026, No. 49, available at `https://www.b3.com.br/en_us/market-data-and-indices/data-services/market-data/reports/daily-bulletin/public-listed-and-otc-products-data/`.

out-of-sample evaluation window. In brief, the in-sample period is used for model search and statistical selection, while the out-of-sample period is reserved for evaluating whether selected rules remain economically and statistically relevant outside the estimation window. From these data we construct a mechanically generated universe of 31,986 candidate trading rules per horizon spanning seven families of technical strategies: RSI, Filter, Moving Average, Channel Breakout, Support–Resistance, Bollinger Bands, and MACD. This count is the number of admissible parameter combinations implied by the pre-specified grids and inequality restrictions for each rule variant (for example, constraints such as $y < x$, $d_y \leq d_x$, $p < q$, and $n < p < q$), summed across variants within one horizon; details are reported in Appendix Table 6. By family, the per-horizon totals are: RSI (600), Filter (3,360), Moving Average (12,870), Channel Breakout (3,000), Support–Resistance (2,160), Bollinger Bands (9,408), and MACD variants (588), which sum to 31,986.

To address the multiple-testing problem inherent in large rule universes, statistical inference relies on the Stepwise Superior Predictive Ability (Stepwise SPA) procedure of White (2000); Hansen (2005); Romano and Wolf (2005). Intuitively, instead of asking whether one pre-chosen rule works, the Stepwise SPA asks whether any rule in a large pool truly outperforms a benchmark once we correct for the fact that many alternatives were tried. We incorporate explicit transaction costs, impose strict day-trading constraints with end-of-session position closure, and evaluate performance both in-sample and out-of-sample.

We contribute to the literature by showing that technical trading rules can generate statistically significant profits in the BRL/USD futures market, but that these profits are strongly horizon- and regime-dependent. After correcting for data-snooping bias via Stepwise SPA, significant rules are concentrated at intermediate intraday horizons and vary systematically across rule families, indicating that design choices and sampling frequency are first-order determinants of performance. Moreover, rule-level profitability proves highly unstable over time: fewer than 4% of in-sample significant rules retain significance out of sample, so investors cannot reliably exploit historical best strategies going forward. Although some rules are profitable within the estimation window and survive multiple-testing correction, identifying and applying them *ex ante* remains elusive, supporting a conditional-viability interpretation rather than evidence of persistent alpha.

The remainder of the paper is organized as follows. Section 2 reviews the literature on technical trading and exchange-rate predictability. Section 3 presents the theoretical framework. Section 4 describes the data. Section 5 outlines the empirical methodology. Section 6 reports the results, and Section 7 concludes.

2 Literature Review

A central question in financial economics is whether persistent abnormal trading profits are compatible with informationally efficient markets. In the classical framework of Fama (1970), weak-form efficiency rules out systematic predictability from past prices, while stronger forms further limit gains from private information. Evidence of returns associated with insider information (e.g. Seyhun, 1986), together with the persistence of technical trading in practice, motivates broader interpretations of efficiency. Along these lines, Lo (2004) proposes the Adaptive Markets Hypothesis, in which efficiency is time-varying and profitability is episodic rather than constant. This view is central to our paper, because our main hypothesis is that predictability and profitability are regime-dependent.

Exchange-rate predictability. The exchange-rate predictability literature asks whether exchange rates can be forecast with economically useful precision at relevant horizons. A large body of work shows that exchange rates are difficult to predict, with structural macroeconomic models often failing to outperform a random-walk benchmark out-of-sample (Meese and Rogoff, 1983; Engel and West, 2005). For Brazil specifically, Marçal and Junior (2016) revisit this benchmark and show how difficult it remains to beat the random walk in BRL/USD forecasting, while Gaglianone and Marins (2017) evaluate point and density forecasts for the Brazilian exchange rate using both macroeconomic and financial information. Evidence for emerging markets is somewhat more nuanced: Pincheira and Selaive (2011) document predictability at intermediate horizons in an emerging-economy setting, and Baku (2019) show that financial and macroeconomic variables can help predict exchange-rate returns across emerging markets. Taken together, this literature suggests that while unconditional predictability remains weak, forecastability may vary across horizons, model classes, and market environments. This motivates our focus on whether short-horizon predictability can be economically exploited in BRL/USD futures through technical trading rules. This perspective is also consistent with evidence that FX practitioners rely heavily on technical tools precisely at the horizons where structural models are least informative (Menkhoff and Taylor, 2007).

Technical trading rules in financial markets. This strand evaluates whether rule-based strategies generate returns after accounting for model search and implementation frictions. In FX, Hsu et al. (2016) tests a very large rule set and documents profitable episodes, with Brazil among the strongest markets; related evidence appears in Coakley et al. (2016) and Ghanem et al. (2024). At higher frequencies and in other asset classes, Hudson et al. (2017), Narayan et al. (2015), and Chu et al. (2020) also report profitable technical strategies, although profits weaken once realistic costs and trading frictions are included. As synthesized by Farias Nazário et al. (2017), key concerns remain cost measurement, risk adjustment, and external validity across markets and samples; in addition, long-horizon abnormal returns may reflect regime-dependent risk premia (Hřebačka, 2025). This motivates our design choice to implement a large rule universe with explicit costs, risk-aware metrics, and multiple-testing-robust inference.

High-frequency trading and market microstructure. This literature studies how trading technology and market design shape short-horizon price formation and tradability. Velu (2020) links regulatory changes to tighter spreads and a shift toward speed-based automated trading, while Menkveld (2013) documents high-frequency trading (HFT) activity consistent with electronic market-making and liquidity provision. Baron et al. (2012) further shows that HFT order flow contains short-run predictive content. In Brazilian FX, futures markets are especially relevant because they concentrate liquidity and often lead spot prices (Fernandes, 2008). This microstructure context supports our choice to study intraday `DOL` futures, where second-by-second dynamics and turnover costs are first-order determinants of rule performance.

Taken together, the literature points to four persistent gaps: limited evidence for emerging markets at intraday horizons, incomplete treatment of market-specific transaction costs, uneven control for data-snooping and model multiplicity, and weak integration between regime-dependent theory and empirical rule selection. Our paper addresses these gaps by combining a large mechanically generated rule universe, realistic Brazilian cost calibration, Stepwise SPA inference, and an explicit mapping from market conditions to rule-selection outcomes (e.g. Carver, 2023).

3 A Simple Model of Episodic Predictability and Trading Profitability

This section develops a minimal framework illustrating how exchange-rate returns may appear unpredictable on average while still displaying short-lived episodes of conditional predictability that support trading profitability. The model features a latent directional component that generates temporary predictability while averaging to zero in the long run. Observation noise ensures that unconditional predictability remains weak, even though directional pressure may occasionally emerge. The objective is not to build a full structural model of exchange-rate determination, but rather to provide a simple mechanism capable of reconciling unconditional unpredictability with episodic trading profitability. We also show how a directional trading signal can be extracted from observed prices and variance information alone, and how the resulting profitability condition naturally allows multiple trading rules to be profitable simultaneously.

Model Setup. Let F_t denote the futures exchange rate and define the log price $f_t = \ln F_t$. Log returns satisfy

$$\Delta f_t = \mu_t + \sigma z_t, \tag{1}$$

where $z_t \sim \text{i.i.d.}(0,1)$ represents observation noise with constant variance σ^2, and μ_t is a latent drift component capturing temporary order-flow imbalances, news effects, or other forces generating short-lived directional pressure in the market.

The latent component evolves according to

$$\mu_t = \rho\mu_{t-1} + u_t, \qquad u_t \sim \text{i.i.d.}(0, \sigma_u^2), \tag{2}$$

with $|\rho| < 1$. The shocks z_t and u_t are mutually independent and $\mathbb{E}[\mu_t] = 0$. By construction,

$$\mathbb{E}[\Delta f_t] = 0, \tag{3}$$

so exchange-rate changes are unpredictable in unconditional mean, consistent with the classical martingale-difference benchmark. Predictability arises only conditionally through the latent drift μ_t, which introduces temporary directional components.

Information Assumptions. Agents observe the history of past returns

$$\mathcal{J}_t = \{\Delta f_s : s \le t\},$$

but do not observe the latent drift μ_t directly. Agents are assumed to know the variance parameters σ^2 and σ_u^2, which are in principle estimable from observed returns, but they do not need to know the persistence parameter ρ. This is a deliberately weak informational assumption: agents understand that a latent directional component exists and have knowledge of its magnitude relative to observation noise, but they do not know its exact dynamic structure.

Signal Extraction. Given knowledge of the relative magnitudes of the latent component and observation noise, agents attempt to infer the unobserved drift from observed returns. To obtain a tractable representation of this inference problem, consider the linear projection of the latent drift onto the observed return.

Let the extracted signal be

$$\widehat{\mu}_t = \lambda\,\Delta f_t, \tag{4}$$

where λ is the coefficient of the population OLS projection of μ_t on Δf_t. The projection coefficient is defined by the moment condition

$$\mathbb{E}\Big[(\mu_t - \lambda \Delta f_t)\Delta f_t\Big] = 0. \tag{5}$$

Using the return decomposition $\Delta f_t = \mu_t + \sigma z_t$ and the orthogonality condition $\mathbb{E}(\mu_t z_t) = 0$, the projection coefficient becomes

$$\lambda = \frac{\mathbb{E}(\mu_t^2)}{\mathbb{E}(\mu_t^2) + \sigma^2} = \frac{\mathrm{Var}(\mu_t)}{\mathrm{Var}(\mu_t) + \sigma^2}. \tag{6}$$

Under the AR(1) specification in (2), the stationary variance of the latent drift is

$$\mathrm{Var}(\mu_t) = \frac{\sigma_u^2}{1 - \rho^2}, \tag{7}$$

so that the signal weight can be written as

$$\lambda = \frac{\sigma_u^2/(1-\rho^2)}{\sigma_u^2/(1-\rho^2) + \sigma^2}. \tag{8}$$

The extracted signal $\widehat{\mu}_t$ therefore provides a noisy but informative estimate of the latent directional component. When the variance of the latent component is large relative to the observation noise, the return contains substantial information about the underlying drift. When observation noise dominates, the signal is heavily attenuated. In such environments the signal-to-noise ratio is low, so unconditional predictive regressions may have little statistical power even though economically meaningful directional pressure occasionally arises.

Trading Rule and Profitability Condition. Consider a simple directional trading rule

$$S_t = \mathrm{sign}(\widehat{\mu}_t)\,\mathbf{1}\{|\widehat{\mu}_t| \geq \tau\}, \tag{9}$$

where $\tau > 0$ is a threshold that prevents trading when the estimated signal is weak relative to transaction costs.

The net return from holding this position over one period is

$$R_{t+1} = S_t\,\Delta f_{t+1} - \kappa|\Delta S_t|, \tag{10}$$

where $\kappa > 0$ denotes proportional transaction costs.

Since the latent component is persistent and agents observe neither μ_t nor ρ directly, the feasible approximation replaces both with their observable counterparts - the extracted signal $\widehat{\mu}_t$ and a sample estimate $\widehat{\rho}$ - yielding

$$\mathbb{E}[\Delta f_{t+1} \mid \mathcal{J}_t] \approx \widehat{\rho}\, \widehat{\mu}_t. \tag{11}$$

The conditional expected net return is therefore

$$\mathbb{E}[R_{t+1} \mid \mathcal{J}_t] \approx S_t\, \widehat{\rho}\, \widehat{\mu}_t - \kappa\, \mathbb{E}[|\Delta S_t| \mid \mathcal{J}_t], \tag{12}$$

and a sufficient condition for conditional profitability is

$$|\widehat{\rho}\, \widehat{\mu}_t| > \kappa\, \mathbb{E}[|\Delta S_t| \mid \mathcal{J}_t]. \tag{13}$$

This condition holds episodically—precisely when the estimated drift is large enough to overcome transaction costs. Because $\widehat{\rho}$ scales the signal uniformly, the threshold τ in the trading rule implicitly absorbs its effect: a stronger estimated persistence raises the expected return from any given signal strength, which is equivalent to lowering the effective threshold. When $\widehat{\mu}_t \approx 0$, the condition fails and any active strategy incurs losses on net. The model therefore illustrates how unconditional unpredictability can coexist with episodic trading profitability: the market behaves on average like a martingale, yet short-lived episodes of directional pressure generate exploitable signals.

Multiple Rules and the Same Latent Signal. An important implication of this framework is that multiple distinct trading rules may be simultaneously profitable during episodes of elevated drift. Any rule $\widehat{s}_t^{(j)}$ that is positively correlated with $\operatorname{sign}(\mu_t)$ when $|\mu_t|$ is large will satisfy condition (13) with positive probability. Different technical rules—moving averages, momentum filters, oscillators, channel breakouts—represent alternative approximations to the same underlying signal extraction problem, each applying a different functional transformation of past prices. Their simultaneous profitability therefore need not reflect market inefficiency, but rather the presence of episodic directional pressure that can be detected imperfectly by many heuristic filters.

Why Profitability Disappears on Average and at Longer Horizons. Two forces limit profitability. First, because $\mathbb{E}[\mu_t] = 0$ and drift episodes are transient, unconditional expected returns from any active strategy tend toward zero over long samples. Second, consider returns aggregated over h periods:

$$\Delta f_t^{(h)} = \sum_{i=0}^{h-1} \Delta f_{t-i} = \sum_{i=0}^{h-1} \mu_{t-i} + \sigma \sum_{i=0}^{h-1} z_{t-i}. \tag{14}$$

Because the latent drift is mean-reverting, its cumulative contribution grows more slowly than the variance of the noise component, which increases proportionally with $h\sigma^2$. As a result, the signal-to-noise ratio deteriorates as the horizon increases. This mechanism provides a simple theoretical rationale for the empirical finding that trading-rule profitability tends to decline at lower sampling frequencies.

Extensions. The framework presented here is intentionally parsimonious. A natural extension would allow the observation variance to vary over time, for example through a GARCH-type specification for σ_t^2. In such a setting the signal-to-noise ratio

$$\lambda_t = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_t^2}$$

would fluctuate across volatility regimes, causing the optimal signal weight to change over time. Trading rules that adapt to volatility conditions—such as Bollinger Bands or volatility-scaled momentum filters—could then be interpreted as approximations to this time-varying optimal filter. More generally, any additional latent structure in the data-generating process opens space for additional families of trading rules to become simultaneously profitable, each approximating a different aspect of the unobserved state. The large universe of rules explored in the empirical analysis of this paper can be understood in this light.

In this interpretation, technical rules are heuristic signal-extraction devices for the latent directional component μ_t. Moving averages, momentum filters, breakouts, and oscillators correspond to alternative weighting schemes applied to past returns. When drift episodes make μ_t large relative to observation noise, these distinct filters can recover the same directional sign, so multiple technically different rules may be simultaneously profitable. This mechanism is directly consistent with the empirical finding that many rules appear significant in the data.

4 Data

The data analyzed in this paper are transaction-level records of all trades executed during the continuous trading session, obtained from B3's website. For each trading day, the analysis is restricted to the most actively traded `DOL` maturity, defined as the `DOL` contract with the highest traded volume on that day. Although the WDO mini contract is the most traded, its dynamics are highly correlated with DOL, and to save processing power we used the smaller tick-by-tick sample. The full sample spans July 10, 2024 to December 23, 2024, totaling 115 trading days. The first 30 trading days, from July 10, 2024 to August 21, 2024, define the in-sample period, while the remaining 85 trading days, from August 22, 2024 to December 23, 2024, define the out-of-sample period, totaling 3,493,627 transaction trades. On average, each trading day contains 30,379 transactions, with a median of 28,223, a minimum of 10,579, and a maximum of 66,009. The in-sample period contains 1,038,805 transactions across 30 trading days, averaging 34,627 trades per day, while the out-of-sample period contains 2,454,822 transactions across 85 trading days, averaging 28,880 trades per day. All timestamps are treated in B3 local market time and transaction prices are aggregated into five intraday frequencies: 5 seconds, 15 seconds, 30 seconds, 60 seconds, and 120 seconds.

Figure 1. Daily open-high-low-close (OHLC) candlestick chart for the `DOL` futures series. The in-sample period (July 10, 2024 to August 21, 2024) covers DOLQ24 and DOLU24, while the out-of-sample period (August 22, 2024 to December 23, 2024) covers the remaining DOLU24,then DOLV24, DOLX24, DOLZ24, and DOLF25. Green (red) candles indicate a close above (below) the open, and the wicks show the intraday low-high range. The dashed vertical line marks the in-sample/out-of-sample (IS/OOS) split.

Source: Authors' calculations based on B3 transaction-level data.

Figure 1 displays the daily OHLC candlesticks for the contracts used in the analysis. OHLC candlesticks are constructed by Open (First) Price, High (Maximum) Price, Low (Minimum) Price and Close (Last) Price. The in-sample segment covers DOLQ24 and DOLU24, whereas the out-of-sample segment covers the remaining DOLU24 observations and then rolls through DOLV24, DOLX24, DOLZ24, and DOLF25; the dashed vertical line marks the split between the two periods. The sample is heterogeneous, with trading days exhibiting different trends and volatility levels, which makes it representative of actual market conditions. The in-sample portion of the data serves as the basis for applying the Stepwise-SPA method to identify significant rules during that period, while the out-of-sample portion will be used to answer if robust in the in-sample period can continue to give positive results in periods with different market conditions.

Table 1 reports descriptive statistics of intraday log returns across five aggregation frequencies and reveals several stylized facts that are central to the empirical design. First, return volatility increases monotonically with the sampling interval, as shown by the systematic rise in the standard deviation from 5-second to 120-second bars in the full sample, in-sample, and out-of-sample partitions. This pattern is consistent with temporal aggregation, since longer intervals incorporate a larger amount of price variation. Second, return distributions are markedly non-Gaussian at all frequencies: skewness is uniformly negative, indicating a heavier left tail, while kurtosis remains far above the Gaussian benchmark, particularly at higher-frequency horizons, which points to substantial tail risk and intermittent extreme movements. Third, the comparison between in-sample and out-of-sample periods suggests meaningful regime heterogeneity, with out-of-sample moments generally showing higher average returns and, in several cases, more pronounced downside realizations. Taken together, these results characterize the BRL/USD intraday environment as volatile, asymmetric, and heavy-tailed, thereby supporting the use of robust multiple-testing procedures and reinforcing the expectation that trading-rule performance is likely to depend on both market regime and aggregation horizon.

4.1 Methodology

We evaluate whether technical trading rules can generate statistically and economically meaningful performance in high-frequency U.S. dollar futures (`DOL`) traded on B3. The dataset contains transaction-level records from the continuous session, and each trading day is represented by the most actively traded contract. We split the sample into a 30-day in-sample window and an out-of-sample window covering the remaining days. The full empirical pipeline is executed separately in each window, and in each case all candidate rules are tested at five intraday horizons (5, 15, 30, 60, and 120 seconds). Statistical predictability is evaluated with the Stepwise SPA test (White, 2000; Hansen, 2005; Romano and Wolf, 2005; Hsu et al., 2016), implemented separately by horizon using daily net returns after transaction costs; in practice, this helps prevent false positives that can appear when thousands of candidate rules are tested. Rejection of the null indicates that at least one rule achieves positive expected net performance after accounting for the joint search across a large rule universe.

The empirical procedure has two distinct stages. In the in-sample stage, the full rule universe is screened statistically and only rules that survive Stepwise SPA are retained; this

Table 1. Descriptive statistics of intraday log returns by aggregation frequency

Stats	5s	15s	30s	60s	120s
Full sample					
Min	-6.06e-03	-7.73e-03	-7.81e-03	-8.50e-03	-8.97e-03
Average	1.99e-07	5.88e-07	1.14e-06	2.25e-06	4.24e-06
Max	5.53e-03	5.45e-03	4.80e-03	4.32e-03	5.29e-03
Std. dev.	1.13e-04	1.92e-04	2.67e-04	3.74e-04	5.14e-04
Skewness	-6.36e-01	-9.22e-01	-9.13e-01	-8.48e-01	-6.32e-01
Kurtosis	7.01e+01	4.46e+01	3.05e+01	2.27e+01	1.40e+01
In-sample					
Min	-2.96e-03	-4.81e-03	-5.65e-03	-6.40e-03	-8.97e-03
Average	1.29e-08	1.23e-08	7.23e-08	1.45e-07	2.20e-07
Max	3.03e-03	3.03e-03	3.03e-03	3.19e-03	3.88e-03
Std. dev.	1.02e-04	1.74e-04	2.41e-04	3.41e-04	4.78e-04
Skewness	-1.23e-01	-3.73e-01	-6.94e-01	-6.82e-01	-1.16e+00
Kurtosis	2.48e+01	2.06e+01	1.88e+01	1.55e+01	2.33e+01
Out-of-sample					
Min	-6.06e-03	-7.73e-03	-7.81e-03	-8.50e-03	-7.29e-03
Average	2.67e-07	7.96e-07	1.53e-06	3.01e-06	5.69e-06
Max	5.53e-03	5.45e-03	4.80e-03	4.32e-03	5.29e-03
Std. dev.	1.17e-04	1.98e-04	2.75e-04	3.86e-04	5.27e-04
Skewness	-7.56e-01	-1.05e+00	-9.62e-01	-8.87e-01	-4.89e-01
Kurtosis	7.86e+01	4.92e+01	3.25e+01	2.40e+01	1.16e+01

Source: Authors' calculations based on B3 transaction-level data. Notes: returns are computed as log differences of consecutive bar closes at each aggregation frequency. Statistics are reported separately for the full sample, the in-sample period, and the out-of-sample period.

test controls for multiple testing within the candidate set, so surviving significance is not an artifact of trying many strategies. In the out-of-sample stage, we evaluate the economic performance of those preselected rules. This separation is meant to reduce data-snooping bias by avoiding economic evaluation on the same data used for statistical selection.

Intraday data are pre-processed by converting transaction timestamps to B3 local time and then aggregating trades into regular open-high-low-close-volume (OHLCV) bars. We aggregate tick-by-tick data because raw ticks are irregularly spaced and heavily affected by bid-ask bounce and microstructure noise, which makes direct rule comparison difficult across time and days. Fixed-time bars create a comparable sampling grid and allow us to study how signal quality changes as noise is smoothed. For each bar, the open, high, low, and close correspond to the first, maximum, minimum, and last transaction prices observed within the interval, and volume is the sum of traded quantities. When no transaction occurs in a bar, the close is carried forward from the last observed trade, while the open, high, and low are all set equal to that carried-forward close and volume is set to zero. The upper bound of 120 seconds is chosen to remain in a clearly intraday high-frequency domain while being long enough to reduce very short-horizon noise and provide a meaningful contrast with 5-second bars. Although OHLCV bars are constructed, the trading rules themselves are applied only to the bar close series. Let P_t denote the aggregated close price at bar t. The return of the underlying futures price used in the performance decomposition is the close-to-close log return,

$$r_t^a = \ln(P_t) - \ln(P_{t-1}) = \ln\left(\frac{P_t}{P_{t-1}}\right). \tag{15}$$

To keep notation aligned with the simple model in Section 3, we use $r_t^a \equiv \Delta f_t$ at the bar level, where $f_t = \ln P_t$.

The rule universe includes Relative Strength Index rules (O1 and O2), filter rules (F1, F2, F3), moving-average rules (MA1, ..., MA5), channel breakout rules (CB1, CB2), support-and-resistance rules (SR1, SR2), Bollinger Bands (BB), and MACD. Each family is evaluated over a large grid of hyperparameters. For each rule, the generated signal is mapped into a discrete position sequence, with values in $\{-1, 0, 1\}$ corresponding to short, neutral, and long positions. Positions are always based only on information available up to the previous bar. In addition, the implementation imposes a day-trading environment, so any open position is closed in the last bar of the day and no overnight exposure is carried.

Net returns per bar incorporate proportional transaction costs linked to position turnover. Let $S_{j,t-1}$ be the position of rule j held from the end of bar $t-1$ to the end of bar t, and let $\Delta S_{j,t} = S_{j,t} - S_{j,t-1}$ denote the position change. In practice, the net return is calculated as

$$R_{j,t} \;=\; S_{j,t-1}\, r_t^a \;-\; \kappa\, |\Delta S_{j,t}|, \tag{16}$$

where κ is the one-way transaction cost per entry/exit, calibrated to the DOL market. Under this specification, entries ($0 \to \pm 1$), exits ($\pm 1 \to 0$), and reversals ($-1 \leftrightarrow +1$) receive different penalties through $|\Delta S_{j,t}|$. This formulation is well suited to second-by-second horizons, where microstructure frictions (bid and ask spreads, fees, and price impact) primarily operate through costs associated with changes in trading state.

Transaction costs are incorporated directly into each rule's bar-level net profit-and-loss series. For each trading day, the one-way fee is set equal to US\$0.11 converted into BRL using the most recent PTAX selling rate (the Brazilian Central Bank reference exchange rate) available before that date. Let $\Pi_{j,b}$ denote the net monetary P&L of rule j in bar b. The code converts this into a normalized return by dividing by the lagged monetary value of one DOL contract, computed as $50P_{b-1}$, where 50 is the contract point value. Thus,

$$\widetilde{R}_{j,b} = \frac{\Pi_{j,b}}{50\, P_{b-1}}. \tag{17}$$

Daily returns used in the Stepwise SPA are obtained by summing these normalized bar returns within each day:

$$R^{\text{day}}_{j,d} = \sum_{b \in d} \widetilde{R}_{j,b}. \tag{18}$$

This section is organized as follows. Subsection 4.1 presents the technical trading rules and rule-universe construction. Subsection 4.1 describes the performance metrics. Subsection 4.1 presents the Stepwise SPA procedure.

Technical trading rules.upe se atrasar Technical trading can convert price dynamics into explicit trading signals and examine the performance of systematically implemented technical trading rules in high-frequency. Our implementation follows the general taxonomy used in the empirical technical-analysis research from Hsu et al. (2016), while adapting the rules to an intraday futures environment. Because the available dataset contains trade prints but not the full order book or message-level order-flow records, we do not include tape-reading/order-flow rules in the candidate set.

The implemented rule universe contains seven families. *Oscillator rules* are represented by two RSI variants (O1 and O2), which seek to detect overbought and oversold conditions from recent price movements. *Filter rules* (F1–F3) generate signals only when cumulative price changes exceed pre-specified thresholds, thereby filtering out minor fluctuations. *Moving-average rules* (MA1–MA5) include single-, double-, and triple-average specifications, with variants that also impose holding periods after signal generation. *Channel-breakout rules* (CB1–CB2) initiate positions when prices cross rolling upper or lower bounds constructed from recent extrema. *Support-and-resistance rules* (SR1–SR2) exploit breaches of rolling price barriers interpreted as directional signals. In addition, the code evaluates

Bollinger Band rules (BB), which compare prices with rolling bands centered on a moving average and scaled by recent volatility, and *MACD-type rules* (MACD and MACDZero), which rely on differences between fast and slow exponential moving averages and their associated signal filters.

The parameterization is adapted to the intraday `DOL` setting rather than imported mechanically from lower-frequency spot-FX studies. Lookback windows, holding periods, and threshold parameters are calibrated in bar units so that the rules remain active and empirically relevant at second-based frequencies. In particular, rules based on percentage moves or breakout distances use intraday-scaled trigger values, avoiding the problem that thresholds designed for daily data become too coarse at high frequency. The exact parameter grids for each family are those implemented in the code and are reported in Appendix A.

The rule universe is generated mechanically by combining the taxonomy above with the parameter grids reported in Appendix A; no rules are selected *ex post* at this stage. Each family is evaluated over all parameter combinations in its corresponding grid, and the resulting Cartesian expansion defines the full candidate set. Enumerating these combinations yields 31,986 candidate rules per horizon (159,930 across the five horizons considered). In other words, the total is purely combinatorial: for each variant, we count all admissible grid points after applying the parameter restrictions, and then add those counts across variants (see Appendix B 6).

Performance Metrics. In this subsection, $d = 1, \ldots, D$ indexes trading days and b indexes bars within each day. The daily net return of rule j is $R_{j,d} \equiv R^{\text{day}}_{j,d}$. Let $r_d \equiv \sum_{b \in d} r^a_b$ denote the daily underlying close-to-close log return obtained by aggregating bar-level returns.

The first metric is the mean excess return of the j-th technical trading rule, already net of transaction costs. It evaluates whether, on average, the strategy is profitable over time. It is defined as:

$$\bar{R}_j = \frac{1}{D} \sum_{d=1}^{D} R_{j,d}. \tag{19}$$

The Sharpe Ratio measures excess return per unit of risk, where risk is proxied by the standard deviation of excess returns. It allows comparisons across strategies while accounting for the volatility of outcomes. A higher Sharpe Ratio indicates a better risk–return trade-off:

$$SR_j = \frac{\bar{R}_j}{\sigma_j}, \tag{20}$$

where σ_j is the standard deviation of excess returns for rule j, estimated using HAC methods.

Total net returns $R_{j,d}$ can be decomposed into two components:

$$R_{j,d} = R(\text{Tilt})_{j,d} + R(\text{Tim})_{j,d}. \tag{21}$$

The tilt component is given by:

$$R(\text{Tilt})_{j,d} = \left(\frac{1}{D} \sum_{d=1}^{D} S_{j,d-1} \right) r_d. \tag{22}$$

The market-timing component is:

$$R(\text{Tim})_{j,d} = R_{j,d} - R(\text{Tilt})_{j,d}. \tag{23}$$

The mean timing return helps determine whether performance arises from active entry/exit decisions or simply from maintaining a near-constant exposure. It is defined as:

$$\bar{R}(\text{Tim})_j = \bar{R}_j - \left(\frac{1}{D} \sum_{d=1}^{D} S_{j,d-1} \right) \left(\frac{1}{D} \sum_{d=1}^{D} r_d \right). \tag{24}$$

Finally, the Sharpe Ratio based solely on the market-timing component assesses whether timing decisions are consistently profitable after accounting for risk, isolating this effect from buy-and-hold behavior

$$SR(\text{Tim})_j = \frac{\bar{R}(\text{Tim})_j}{\sigma(\text{Tim})_j}, \tag{25}$$

where $\sigma(\text{Tim})_j$ is the HAC standard deviation of the market-timing returns.

Stepwise SPA (Stepwise Superior Predictive Ability) Formal statistical inference in the current implementation is based exclusively on the Stepwise SPA procedure. For each horizon and sample split, let $G = \{G_{d,j}\}$ denote the $D \times J$ matrix of daily net returns, where $d = 1, \ldots, D$ indexes trading days and $j = 1, \ldots, J$ indexes technical trading rules. The null hypothesis for each rule is one-sided,

$$H_0^{(j)} : \mathbb{E}[G_{d,j}] \leq 0 \qquad \text{against} \qquad H_1^{(j)} : \mathbb{E}[G_{d,j}] > 0,$$

so the test asks whether a given rule generates positive mean daily net returns after transaction costs. In the code, the observed statistic is based on the sample mean daily return,

$$\theta_j = \frac{1}{D} \sum_{d=1}^{D} G_{d,j}, \qquad S_j = \sqrt{D}\, \theta_j.$$

To approximate the joint null distribution while preserving serial dependence across trading days, the procedure uses the stationary bootstrap of Politis and Romano (1994), implemented with $B = 1000$ bootstrap replications. The bootstrap is applied to the day index, so that entire daily return vectors across rules are resampled jointly. Following the Stepwise SPA logic, the bootstrap statistics are recentered in order to reduce the influence of sufficiently poor-performing rules on the critical values. The stepwise component is then implemented through an iterative elimination scheme: rules that reject the null at significance level $\alpha = 0.10$ are recorded as significant, removed from the active set, and the procedure is repeated on the remaining rules until no further rejections occur.

The final inferential output for each rule consists of its sample mean daily return, the associated observed statistic, a bootstrap-adjusted p-value, an inclusion indicator, and the step at which the rule is selected. These outputs are merged back into the full performance database. Importantly, statistical significance is defined only with respect to mean daily net returns. Other reported measures, such as the sample Sharpe ratio and the tilt-timing decomposition, are descriptive performance summaries computed after the return-based Stepwise SPA classification and are not subjected to separate hypothesis tests.

5 Results

The empirical evaluation starts from the full rule universe in Section 4.1. Each candidate rule is tested with Stepwise SPA on mean excess returns at $\alpha = 0.10$, separately by horizon. Table 2 reports significant rules by family and horizon. The total row shows a clear hump: 163 (5s), 378 (15s), 621 (30s), 774 (60s), and 505 (120s). Predictability is therefore strongest at intermediate intraday horizons.

Table 2. Significant rules by family and horizon

Family	5	Min. p-value	15	Min. p-value	30	Min. p-value	60	Min. p-value	120	Min. p-value
BB	35	3.4e-2	155	2.0e-2	330	1.3e-2	283	9e-3	103	5e-3
CB	52	3e-3	65	1e-3	46	1.3e-2	76	$< 1e{-}2$	86	1e-3
Filter	16	8.6e-2	40	1.2e-2	62	$< 1e{-}2$	107	1e-3	54	1.5e-2
MACD	0	2.37e-1	63	2.0e-2	84	1.1e-2	91	$< 1e{-}2$	76	3e-3
MA	15	9e-3	21	9e-3	38	6e-3	88	3e-3	84	$< 1e{-}2$
RSI	6	4.1e-2	16	2e-3	28	8e-3	35	8e-3	16	9e-3
SR	39	1.0e-2	18	1.7e-2	33	2.9e-2	94	$< 1e{-}2$	86	3e-3
Total	163		378		621		774		505	

Source: Authors' calculations using Stepwise SPA outputs from the rule universe described in Section A.
Notes: Entries report significant rules and minimum p-values by family and horizon.

Table 2 also shows that significance is broad but horizon-dependent. At very short horizons, microstructure noise is more relevant; at intermediate horizons, directional moves are clearer; at longer horizons, aggregation dilutes short-lived episodes.

At the family level, the BB row is illustrative: the number of significant Bollinger Band rules rises sharply from 35 (5s) to 155 (15s) and 330 (30s), remains very high at 60s (283), and then falls to 103 at 120s. This hump-shaped profile suggests that BB-type mean-reversion signals are most detectable at intermediate frequencies, where very high-frequency noise is partly filtered but short-lived intraday reversals are still preserved.

Across families at a fixed horizon, the 60-second column (the peak, with 774 significant rules) is the key benchmark: significance is spread across all families but is concentrated in BB (283), Filter (107), SR (94), MACD (91), and MA (88), with smaller contributions from CB (76) and RSI (35). In proportional terms, BB alone represents about 36.6% of all significant rules in this column (283/774), indicating that the aggregate peak at 60 seconds is not only a level effect but also a composition effect driven strongly by specific families.

The family composition also varies meaningfully across horizons. Breakout-style families—Bollinger Bands, Channel Breakouts, and Support–Resistance—are relatively more prominent at shorter horizons, where abrupt directional moves are more relevant. Momentum-type families—MACD, Moving Average, and RSI—become comparatively more important at slightly longer horizons, where trend persistence has more time to develop. This shift suggests that different families are capturing different manifestations of the same underlying price dynamics across bar sizes.

These findings are consistent with the theoretical framework developed earlier. In the model, episodic predictability comes from a latent directional component in returns. When that latent component is large relative to noise, multiple technically distinct rules can detect the same directional signal at the same horizon, which explains why many different rules become significant simultaneously.

These results are consistent with Hsu et al. (2016) in showing predictive content in technical rules, but they also indicate that in Brazilian intraday futures this predictability is strongly regime- and horizon-specific. This interpretation aligns with Menkhoff and Taylor (2007), which stresses the practical relevance of technical rules at short decision horizons.

Having identified which rules are statistically significant in Table 2, the next subsection evaluates their economic magnitude through performance metrics such as mean excess returns and Sharpe ratios.

Day trading performance metrics. We evaluate performance using two metrics: Mean Excess Return (R) and Sharpe Ratio (SR). Table 3 reports annualized daily medians. Results are strong: median excess returns reach 49.74%, and Sharpe ratios exceed 3, a level commonly viewed as outstanding risk-adjusted performance.

Bollinger Bands perform best at 60s and 120s, especially when combined with their p-values and rule counts from Table 2. SR and MA are stronger at higher frequencies, though returns generally decline as the horizon lengthens. Filter and RSI are weaker at 5s but improve at longer horizons. MACD has no robust rules at 5s, yet performs strongly at other horizons. CB remains positive but is the least prominent family.

Table 3. Median annualized R and SR by family for robust rules

	5s		15s		30s		60s		120s	
Family	R (%)	SR	R (%)	SR	R (%)	SR	R (%)	SR	R (%)	SR
BB	49.27	3.79	9.98	1.86	19.89	1.89	41.75	3.22	45.51	3.52
Filter	5.26	1.56	40.77	3.57	38.34	4.43	28.19	3.80	37.59	3.96
MA	42.85	3.43	41.61	3.59	39.57	3.91	33.50	4.23	33.15	4.50
MACD	–	–	37.51	3.26	47.17	3.83	49.74	4.69	33.69	3.43
CB	4.58	1.49	10.99	1.63	14.63	2.78	8.72	2.69	14.47	3.12
SR	35.39	2.75	33.39	3.19	10.86	3.09	13.68	3.64	18.69	3.24
RSI	1.34	-0.20	2.21	2.95	20.55	3.89	33.78	4.27	22.76	3.60

Notes: Entries report only significant rules. R denotes the median annualized mean excess return, expressed in percentage terms, using the median number of bars per day observed for each aggregation level and a 252-trading-day convention. Sharpe Ratio (SR) is annualized using the square-root-of-time rule.

These results can be compared with Hsu et al. (2016) after excluding MACD and Bollinger Bands. Under that comparison, the HFT setting delivers higher returns: median returns exceed the annualized returns of the best rules reported there. Overall, our findings are closer to (Coakley et al., 2016), but with larger magnitudes.

A compact cross-check by row and column is also informative. In the MA row, performance is consistently strong across all horizons: median annualized returns remain between 33.15% and 42.85%, while median Sharpe ratios range from 3.43 to 4.50 and increase with aggregation, peaking at 120s. This pattern suggests relatively stable risk-adjusted performance rather than isolated outliers at a single horizon. In the 60-second column, most families show high risk-adjusted performance simultaneously (SR from 2.69 to 4.69, with MACD, RSI, and MA above 4), indicating that this horizon is not driven by one strategy class alone but reflects a broad improvement in rule quality across families.

To evaluate timing robustness, we report R_{Tim} and SR_{Tim}, which capture predictive timing power among selected significant rules. Unless noted otherwise, these statistics are computed out-of-sample for rules first selected in-sample by Stepwise SPA. Table 4 shows a heterogeneous pattern: BB, MACD, and RSI are strongly positive in several horizons, while MA and SR are more horizon-sensitive and become negative at longer horizons.

At the row level, MA is the clearest case: timing performance is positive at 5s–30s (R_{Tim} between 20.94 and 25.06; SR_{Tim} around 1.65–2.52), but turns sharply negative at 60s and 120s ($R_{Tim} = -51.69$ and -48.08; $SR_{Tim} = -5.88$ and -5.47). This sign reversal indicates that MA rules that are effective for short-horizon timing can become systematically mistimed at longer horizons.

At the column level, 120 seconds clarifies cross-family composition: BB ($R_{Tim} = 84.84$, $SR_{Tim} = 7.43$), RSI (34.21, 5.70), and MACD (31.02, 3.24) remain strongly positive, while MA and SR are negative. Thus, long-aggregation timing performance is concentrated in a subset of families and is not a broad market-wide effect. The key implication is that timing ability is both family-specific and horizon-dependent, reinforcing the need to combine statistical selection with family-level robustness checks.

Table 4. Median annualized R_{Tim} and SR_{Tim} by family for significant rules

	5s		15s		30s		60s		120s	
Family	R_{Tim}	SR_{Tim}	R_{Tim}	SR_{Tim}	R_{Tim}	SR_{Tim}	R_{Tim}	SR_{Tim}	R_{Tim}	SR_{Tim}
BB	59.39	4.61	4.70	2.25	12.30	2.16	87.20	7.63	84.84	7.43
Filter	5.22	1.56	25.73	2.54	23.69	3.31	17.27	2.23	12.17	1.54
MA	20.94	1.65	25.06	2.50	21.33	2.52	-51.69	-5.88	-48.08	-5.47
MACD	–	–	26.64	2.36	48.98	3.92	36.05	3.75	31.02	3.24
CB	2.35	0.89	2.36	1.40	7.75	1.77	4.17	1.83	3.98	1.27
SR	4.78	0.54	16.71	1.68	6.11	1.96	3.16	1.40	-5.12	-0.41
RSI	1.39	0.75	2.75	3.43	29.33	5.22	43.33	5.51	34.21	5.70

Notes: entries report only significant rules. R_{Tim} is annualized as a percentage using the median number of bars per day observed for each aggregation level and a 252-trading-day convention. SR_{Tim} is annualized with the square-root-of-time rule.

5.1 Robustness Checks

5.1.1 Out-of-Sample Portfolio

Finally, this part of the paper evaluates out-of-sample implementability using day-trading portfolios formed from the top 50 in-sample rules at each aggregation horizon. Figure 2 shows substantial cross-horizon dispersion in cumulative P&L. The 5-second portfolio remains mostly above the zero line and ends as the strongest performer, while the 60-second portfolio experiences a pronounced mid-sample drawdown before recovering and closing in positive territory. By contrast, the 120-second portfolio is profitable for much of the window but reverses sharply near the end and finishes slightly negative. The 30-second and 15-second portfolios display persistent downward trajectories throughout the evaluation window, with the 15-second specification recording the deepest losses. Overall, the figure indicates that out-of-sample profitability is highly horizon-dependent and unstable over time, supporting a conditional-viability interpretation rather than evidence of uniformly persistent alpha.

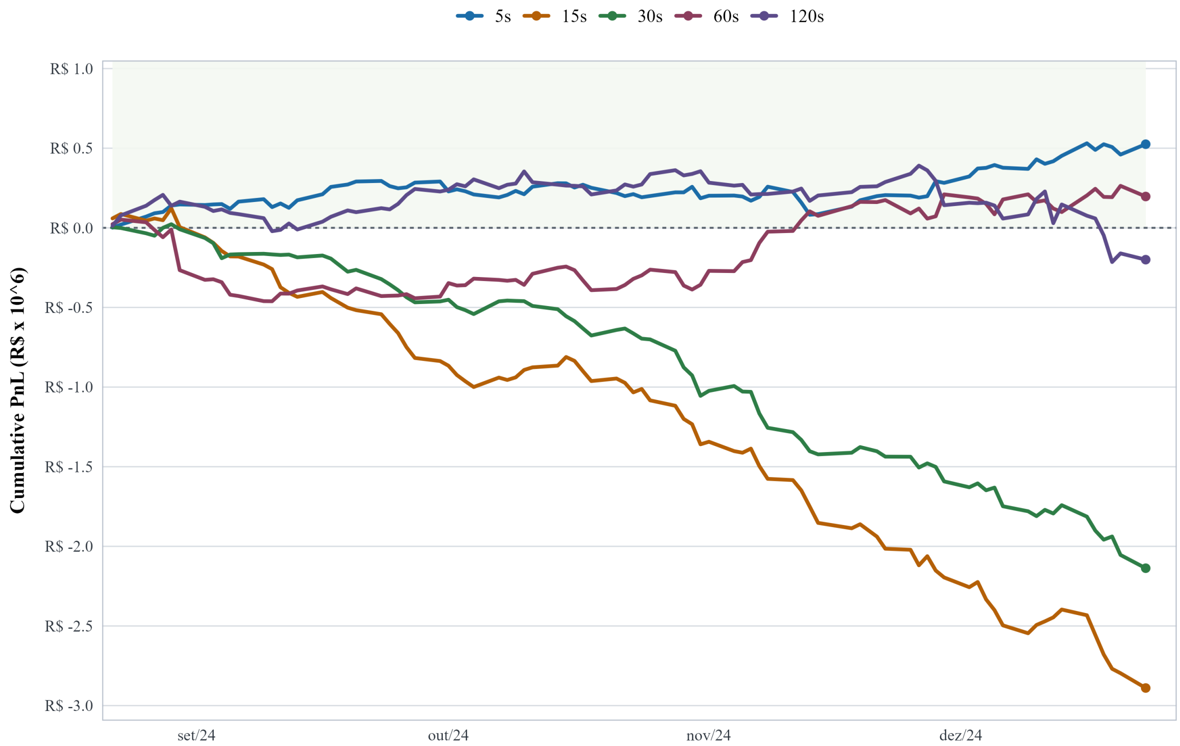

Figure 2. Day Trading portfolio with best 50 rules

Source: Authors' calculations based on B3 transaction-level data.

Notes: Figure produced by the authors from the sample and construction rules described in Section 4.

5.1.2 Trading Rules Persistence

Table 5 shows that statistical significance is highly unstable across sample splits. Of the 2,426 Rule-Sec combinations that are significant in-sample, 2,347 lose significance out of sample, implying that only 79 remain significant in both periods. This low overlap indicates limited persistence of the specific rules selected in-sample and suggests that rule-level profitability is not stable across time.

The same pattern appears within each aggregation frequency. At 5s, only 26 rules remain significant out of the 155 selected in-sample; at 15s, the number falls to only 4 out of 366. Even at the more persistent frequencies, the overlap remains very small: 11 rules at 30s, 21 at 60s, and 17 at 120s. Hence, regardless of aggregation level, the great majority of in-sample significant rules fail to preserve significance in the out-of-sample period.

At the same time, the number of rules that become significant only out of sample is very large, reaching 3,843 Rule-Sec combinations overall. This is especially pronounced at 5s, 15s, and 120s, where out-of-sample significance counts exceed their in-sample counterparts. A natural interpretation is that the rule universe exhibits substantial time variation in performance, so that significance is period-specific rather than persistent. It should also be noted that the out-of-sample period is longer than the in-sample period, which mechanically increases test power and may contribute to the larger number of out-of-sample rejections. Even so, the extremely small intersection between significant in-sample and significant out-of-sample rules points to weak persistence in the identity of the profitable rules.

Table 5. Transition of significant rules from the in-sample to the out-of-sample period

Aggregation	Significant IS	Significant OOS	Lost Significance	Became Significant in OOS
5s	155	1,550	129	1,524
15s	366	753	362	749
30s	629	387	618	376
60s	760	544	739	523
120s	516	688	499	671
ALL	2,426	3,922	2,347	3,843

Notes: counts are computed at the Rule-Sec level. "Lost Significance" corresponds to rules that were significant in-sample but not significant out-of-sample. "Became Significant in OOS" corresponds to rules that were not significant in-sample but became significant out-of-sample.

6 Conclusion

In this paper, we use transaction-level data for U.S. dollar futures traded on B3 to examine the profitability of a large universe of 31,986 technical trading rules per intraday horizon, spanning seven families — RSI, Filter, Moving Average, Channel Breakout, Support–Resistance, Bollinger Bands, and MACD — evaluated at five aggregation levels ranging from 5 to 120 seconds. By applying Stepwise Superior Predictive Ability inference with explicit transaction costs and strict day-trading constraints, we address three related questions: whether high-frequency technical rules can generate net profits after realistic costs and multiple-testing correction; which rule families and horizons drive any predictability; and under which market conditions significance is more likely.

The evidence yields a clear but qualified answer. After correcting for data-snooping bias, significant rules follow a hump-shaped pattern across aggregation levels — peaking at 774 significant rules at the 60-second horizon — and performance varies substantially across families, with Moving Average rules delivering the most consistent risk-adjusted returns and MACD generating no robust rules at the shortest horizon. Profitability is not pervasive: significance is concentrated in specific families and varies with intraday horizon, indicating that design choice and sampling frequency are first-order determinants of performance. Significance is also state-dependent — higher volatility and stronger directional movement increase the likelihood of rule significance, while microstructure-related variables have heterogeneous effects across families — consistent with the theoretical framework in which episodic predictability arises from a latent directional component that different rule families approximate with varying precision across horizons. This horizon-dependence also helps reconcile the more modest profitability documented in daily and swing-trading studies (Hsu et al., 2016; Coakley et al., 2016): at longer horizons, the latent directional component dissipates and fundamentals-driven forces increasingly dominate price dynamics, eroding the short-lived informational advantage that technical rules can exploit at intraday frequencies.

These results reconcile two strands of the FX literature that are often presented as conflicting. Classic exchange-rate predictability results (e.g., Meese and Rogoff, 1983; Engel and West, 2005) emphasize weak unconditional forecastability, while microstructure and technical-trading studies (e.g., Evans and Lyons, 2002; Menkhoff and Taylor, 2007; Hsu et al., 2016) highlight short-run conditional predictability. Our findings support a synthesis: the market is broadly efficient on average, but temporary frictions and order-flow dynamics can create episodic windows in which some rules are profitable.

A more direct test of practical viability confirms the limits of this conditional profitability. Of the 2,426 rule-horizon combinations that are significant in-sample, 2,347 lose significance out of sample — implying a persistence rate below 4%. Out-of-sample portfolio evaluation reinforces this picture: profitability varies sharply across horizons, with the 5-second portfolio ending as the strongest performer while the 15-second specification records the deepest losses, and no horizon delivers uniformly stable positive returns throughout the evaluation window. Importantly, these results do not contradict the evidence in Chague et al. (2020) that most day traders lose money. Extracting returns from the shortest aggregation horizons requires a high risk tolerance, substantial capital, and low-latency infrastructure to avoid execution slippage — barriers that effectively preclude most retail participants from replicating the rule-level profits documented here.

The limited and unstable survival of technical trading rules across sample splits likely reflects time variation in the underlying data generating process. Macroeconomic events, central bank interventions, and shifts in systemic risk can alter the short-run dynamics of the BRL/USD market between the in-sample and out-of-sample periods, rendering previously profitable rules misspecified for the new regime. Consistent with the theoretical model developed in this paper, what determines a rule's viability *ex ante* is the data generating process prevailing at the time of trading. A promising direction for future research is therefore to condition rule selection on pre-market information — such as macroeconomic releases, implied volatility, or order-flow indicators — to identify regimes in which the latent directional component is more likely to be large relative to noise. Such an approach could generate a portfolio of rules with a higher likelihood of out-of-sample robustness. Combining price-based and order-flow-based signals in diversified cross-family portfolios, and incorporating richer implementation frictions such as latency, slippage, and market impact, represent natural next steps for this research agenda.

An important caveat is that the evidence is drawn from a single market (B3 `DOL` futures) and one sample period. At the same time, `DOL` is a highly liquid contract with broad participation and substantial turnover, so there is no strong *a priori* reason to view the qualitative mechanism documented here as unique to a thin or idiosyncratic market. What should not be extrapolated mechanically are effect sizes, which are likely market- and regime-specific. In sum, the evidence supports conditional, state-dependent tradability in a market that remains broadly efficient in the aggregate.

References

Baku, E. (2019). Exchange rate predictability in emerging markets. *International Economics*, 157(C):1–22.

Baron, M., Brogaard, J., and Kirilenko, A. (2012). The trading profits of high frequency traders. *SSRN Electronic Journal*.

Brock, W., Lakonishok, J., and LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47(5):1731–1764.

Carver, R. (2023). *Advanced Futures Trading Strategies: 30 Fully Tested Strategies for Multiple Trading Styles and Time Frames*. Harriman House, Petersfield, Hampshire.

Chague, F., De-Losso, R., and Giovannetti, B. (2020). Day trading for a living? *SSRN Electronic Journal*.

Chu, J., Chan, S., and Zhang, Y. (2020). High frequency momentum trading with cryptocurrencies. *Research in International Business and Finance*, 52:101176.

Coakley, J., Marzano, M., and Nankervis, J. (2016). How profitable are fx technical trading rules? *International Review of Financial Analysis*, 45:273–282.

Engel, C. and West, K. D. (2005). Exchange rates and fundamentals. *Journal of Political Economy*, 113(3):485–517.

Evans, M. D. D. and Lyons, R. K. (2002). Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1):170–180.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417.

Farias Nazário, R. T., e Silva, J. L., Sobreiro, V. A., and Kimura, H. (2017). A literature review of technical analysis on stock markets. *The Quarterly Review of Economics and Finance*, 66:115–126.

Fernandes, A. V. (2008). Microestrutura do mercado cambial brasileiro: comparação do mercado à vista e futuro. Dissertação de mestrado, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Rio de Janeiro, Brasil.

Gaglianone, W. P. and Marins, J. T. M. (2017). Evaluation of exchange rate point and density forecasts: An application to brazil. *International Journal of Forecasting*, 33(3):707–728.

Ghanem, S., Harasheh, M., Sbaih, Q., and Ajmal, T. K. (2024). The predictability of technical analysis in foreign exchange market using forward return: evidence from developed and emerging currencies. *Cogent Business & Management*, 11(1):2428781.

Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4):365–380.

Hsu, P.-H., Taylor, M. P., and Wang, Z. (2016). Technical trading: Is it still beating the foreign exchange market? *Journal of International Economics*, 102:188–208.

Hudson, R., McGroarty, F., and Urquhart, A. (2017). Sampling frequency and the performance of different types of technical trading rules. *Finance Research Letters*, 22:136–139.

Hřebačka, V. (2025). Analyzing market efficiency: The role of business cycles, risk aversion, and occam's razor in the adaptive market hypothesis. *Finance Research Letters*, 75:106840.

InfoMoney (2024). B3 aponta mais de 1 milhão de day traders no brasil em 2024. Financial news report. Accessed: 2026-03-16.

Kohlscheen, E. and Andrade, S. C. (2014). Official fx interventions through derivatives. *Journal of International Money and Finance*, 47:202–216.

Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management, Forthcoming.*

Marçal, E. F. and Junior, E. H. (2016). Is it possible to beat the random walk model in exchange rate forecasting? more evidence for brazilian case. *Brazilian Review of Finance*, 14(1):65–88.

Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1-2):3–24.

Menkhoff, L. and Taylor, M. P. (2007). The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature*, 45(4):936–972.

Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740. High-Frequency Trading.

Narayan, P. K., Mishra, S., Narayan, S., and Thuraisamy, K. (2015). Is exchange rate trading profitable? *Journal of International Financial Markets, Institutions and Money*, 38:217–229.

Osler, C. L. (2003). Currency orders and exchange-rate dynamics: An explanation for the predictive success of technical analysis. *Journal of Finance*, 58(5):1791–1819.

Pincheira, P. and Selaive, J. (2011). External imbalances, valuation adjustments and real exchange rate: Evidence of predictability in an emerging economy. *Revista de Análisis Económico*, 26(1):107–125.

Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.

Romano, J. P. and Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282.

Seyhun, H. (1986). Insiders' profits, costs of trading, and market efficiency. *Journal of Financial Economics*, 16(2):189–212.

Sullivan, R., Timmermann, A., and White, H. (1999). Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, 54(5):1647–1691.

Velu, R. (2020). *Algorithmic Trading and Quantitative Strategies*. Chapman and Hall/CRC, 1 edition.

White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5):1097–1126.

XP Investimentos (2024). Relatórios de análise técnica e estratégias intraday para contratos futuros. Broker research publication. Accessed: 2026-03-16.

A Details on trading rules

This section describes the technical trading rules implemented in this study. Our rule taxonomy follows Hsu et al. (2016), but the implementation is adapted to an intraday futures setting. First, all rules are computed on fixed-width time bars constructed from transaction data at five horizons (5s, 15s, 30s, 60s, and 120s).

RSI oscillator rules (O1–O2) Oscillator rules are based on the Relative Strength Index (RSI), which measures the strength of recent positive price changes relative to recent negative changes and is normalized to the interval $[0, 100]$. Let one-bar price changes be $\Delta P_t = P_t - P_{t-1}$ and define

$$U_t = \max(\Delta P_t, 0), \qquad D_t = \max(-\Delta P_t, 0). \tag{26}$$

For a lookback window of h bars, the RSI is computed from rolling sums (equivalently, rolling averages) of gains and losses:

$$\mathrm{RSI}_t(h) = 100 \cdot \frac{\sum_{i=0}^{h-1} U_{t-i}}{\sum_{i=0}^{h-1} U_{t-i} + \sum_{i=0}^{h-1} D_{t-i}}. \tag{27}$$

We implement two RSI variants:

- **O1 (no fixed holding period).** A long signal is generated when the RSI exits the oversold region, i.e., when it rises above $50 - v$ after having remained below $50 - v$ for at least d consecutive bars. A short signal is generated when the RSI exits the overbought region, i.e., when it falls below $50 + v$ after having remained above $50 + v$ for at least d consecutive bars. In the absence of a new signal, the previous position is maintained; hence positions may flip when the opposite exit condition is triggered.
- **O2 (fixed holding period).** O2 uses the same RSI trigger as O1, but once a position is opened it is locked for k bars, during which intermediate RSI signals are ignored. After the holding counter expires, the rule resumes evaluating new RSI signals; if no new signal is generated, the position returns to the neutral state.

RSI parameter grid (bars/levels).

Parameter	Values
h	$\{5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$
v	$\{10, 15, 20, 25\}$
d	$\{1, 2, 5\}$
k	$\{1, 5, 10, 25\}$

Filter rules (F1–F3) Filter rules are trend-following rules that compare the current price with recent extrema and act only after sufficiently large price movements, thereby filtering out small fluctuations. Let $\mathrm{Low}_{t-1}(j)$ and $\mathrm{High}_{t-1}(j)$ denote the minimum and maximum close over the previous j bars (ending at $t-1$). These serve as reference points for the running most recent low and high used by the filter rules. We implement:

- **F1 (persistent reversal rule).** A long signal is generated when P_t is at least $x\%$ above the most recent low and the condition receives at least d consecutive confirmations. A short signal is generated when P_t is at least $x\%$ below the most recent high under the analogous confirmation rule. Once a position is established, it is maintained until the opposite signal is triggered, subject to end-of-day flattening.
- **F2 (entry/exit with neutrality).** Starting from the neutral state, the rule enters long (short) when the $x\%$ filter condition is satisfied with entry confirmation parameter $d(x)$. After a long entry, the rule tracks the highest price observed since entry and exits to neutral when price falls by at least $y\%$ from that post-entry high for $d(y)$ confirmations. After a short entry, it analogously tracks the lowest price observed since entry and exits when price rises by at least $y\%$ from that post-entry low. The admissible parameter space imposes $y < x$ and $d(y) \le d(x)$.
- **F3 (fixed-holding filter rule).** F3 uses the same entry trigger as F1, but once a position is opened it is locked for k bars, during which intermediate filter signals are ignored. After the holding counter expires, the rule resumes evaluating new filter conditions; if no new signal is generated, it returns to the neutral state.

Filter parameter grid (percent thresholds and bars).

Parameter	Values
x	$\{0.005, 0.01, 0.05, 0.1, 0.5, 1.0, 5.0\}\%$
y	$\{0.005, 0.01, 0.05, 0.1, 0.5, 1.0, 5.0\}\%$
d	$\{0, 1, 2, 3, 4, 5\}$
j	$\{1, 2, 5, 10, 20\}$
k	$\{5, 10, 15, 20, 25\}$
$d(x)$	$\{0, 1, 2, 3, 4, 5\}$
$d(y)$	$\{0, 1, 2, 3, 4\}$

Moving-average rules (MA1–MA5). Let the simple moving average (SMA) over m bars be

$$\mathrm{MA}_t(m) \;=\; \frac{1}{m}\sum_{i=0}^{m-1} P_{t-i}. \tag{28}$$

Moving-average rules are based either on the deviation of price from a single moving average or on the relative ordering of moving averages with different horizons:

- **MA1 (price vs. single MA, persistent directional rule).** A long signal is generated when P_t is at least $x\%$ above $\mathrm{MA}_t(q)$ for at least d consecutive bars; a short signal is generated when P_t is at least $x\%$ below $\mathrm{MA}_t(q)$ for at least d consecutive bars. In the absence of a new signal, the previous position is maintained, subject to end-of-day flattening.

- **MA2 (price vs. single MA, fixed-holding version).** MA2 uses the same trigger as MA1, but once a position is opened it is locked for k bars, during which intermediate signals are ignored. After the holding counter expires, the rule resumes evaluating new conditions; if no new signal is generated, it returns to the neutral state.

- **MA3 (two-MA crossover, persistent directional rule).** With $p < q$, a long signal is generated when $\mathrm{MA}_t(p)$ is at least $x\%$ above $\mathrm{MA}_t(q)$ for at least d consecutive bars, and a short signal is generated when $\mathrm{MA}_t(p)$ is at least $x\%$ below $\mathrm{MA}_t(q)$ for at least d bars. Positions are then maintained until an opposite signal occurs, again subject to end-of-day flattening.

- **MA4 (two-MA crossover, fixed-holding version).** MA4 uses the same crossover trigger as MA3, but once a position is entered it is held for k bars, after which the rule resumes evaluating fresh crossover conditions.

- **MA5 (triple moving-average ordering rule).** With $n < p < q$, the rule compares three moving averages of increasing horizon. A long signal is generated when $\mathrm{MA}_t(n)$ is at least $x\%$ above $\mathrm{MA}_t(p)$ and, simultaneously, $\mathrm{MA}_t(p)$ is at least $x\%$ above $\mathrm{MA}_t(q)$ for at least d consecutive bars. A short signal is generated under the symmetric inequalities. Thus, MA5 is a strict three-average alignment rule rather than a price-versus-average rule.

MA parameter grid (bars and percent thresholds).

Parameter	Values
q	$\{2, 5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$
p	$\{2, 5, 10, 15, 20, 25, 50, 100, 150, 200\}$, with $p < q$
n	$\{2, 5, 10, 15, 20, 25, 50, 100, 150\}$, with $n < p < q$
x	$\{0, 0.005, 0.01, 0.05, 0.1, 0.5\}\%$
d	$\{0, 2, 3, 4, 5\}$
k	$\{5, 10, 25\}$

Channel breakout rules (CB1–CB2). Channel breakout rules first identify a low-variability price channel and then trade on breakouts from that channel. Over the previous

j bars, let $\text{Low}_{t-1}(j)$ and $\text{High}_{t-1}(j)$ denote the rolling minimum and maximum close, respectively. In the implementation, a valid $c\%$ channel exists when

$$\text{High}_{t-1}(j) \;\le\; \text{Low}_{t-1}(j)(1+c), \tag{29}$$

so that the recent high lies within a narrow band above the recent low. Conditional on this channel being present, the implied upper and lower channel bounds are

$$U_t(j,c) = \text{Low}_{t-1}(j)(1+c), \qquad L_t(j,c) = \text{High}_{t-1}(j)(1-c). \tag{30}$$

Breakouts are then defined relative to these bounds:

- **CB1 (persistent breakout rule).** A long signal is generated when $P_t \geq U_t(j,c)(1+x)$ for at least d consecutive bars, while a short signal is generated when $P_t \leq L_t(j,c)(1-x)$ for at least d bars. In the absence of a new signal, the previous position is maintained, subject to end-of-day flattening.
- **CB2 (fixed-holding breakout rule).** CB2 uses the same breakout trigger as CB1, but once a position is opened it is locked for k bars, during which intermediate breakout signals are ignored. After the holding counter expires, the rule resumes evaluating fresh breakout conditions; if no new signal is generated, it returns to the neutral state.

CB parameter grid (percent thresholds and bars).

Parameter	Values
x	$\{0.005, 0.01, 0.05, 0.1, 0.5\}\%$
c	$\{0.01, 0.05, 0.1, 0.5, 1.0\}\%$
d	$\{0, 1, 2\}$
j	$\{5, 10, 15, 20, 25, 50, 100, 200\}$
k	$\{1, 5, 10, 25\}$

Support-and-resistance rules (SR1–SR2). Support-and-resistance rules are similar in spirit to filter and breakout rules, but use endogenous resistance and support levels defined from recent price extrema. Let resistance be $\text{High}_{t-1}(j)$ and support be $\text{Low}_{t-1}(j)$, where both are computed over the previous j bars. Signals are generated when price moves sufficiently far beyond these reference levels:

- **SR1 (persistent directional rule).** A long signal is generated when P_t is at least $x\%$ above resistance, that is, when

$$P_t \geq \text{High}_{t-1}(j)(1+x),$$

for at least $d+1$ consecutive bars. A short signal is generated when

$$P_t \leq \text{Low}_{t-1}(j)(1-x),$$

for at least $d+1$ bars. In the absence of a new signal, the previous position is maintained, subject to end-of-day flattening.

- **SR2 (fixed-holding support-and-resistance rule).** SR2 uses the same trigger as SR1, including the $d+1$ confirmation convention, but once a position is opened it is locked for k bars, during which intermediate signals are ignored. After the holding counter expires, the rule resumes evaluating fresh support-and-resistance signals; if no new signal is generated, it returns to the neutral state.

SR parameter grid (percent thresholds and bars).

Parameter	Values
x	$\{0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1.0, 2.5\}\%$
d	$\{0, 1, 2, 3, 4, 5\}$
j	$\{2, 5, 10, 15, 20, 25, 50, 100, 250\}$
k	$\{1, 5, 10, 25\}$

Bollinger Bands rule (BB). The Bollinger Bands rule is implemented as a mean-reversion rule based on a rolling moving-average envelope. For a lookback window of j bars, let

$$\mu_t(j) = \frac{1}{j}\sum_{i=1}^{j} P_{t-i}, \qquad \sigma_t(j) = \sqrt{\frac{1}{j-1}\sum_{i=1}^{j}\Big(P_{t-i} - \mu_t(j)\Big)^2}. \tag{31}$$

The lower and upper Bollinger bands are then defined as

$$L_t(j, n_\sigma, x) = \Big(\mu_t(j) - n_\sigma \sigma_t(j)\Big)(1-x), \qquad U_t(j, n_\sigma, x) = \Big(\mu_t(j) + n_\sigma \sigma_t(j)\Big)(1+x), \tag{32}$$

where n_σ denotes the band-width multiplier and x is an additional percentage filter.

The implemented rule is contrarian. A long signal is generated when price crosses below the lower band and remains below it long enough to satisfy the confirmation requirement. More precisely, if the price was not below the lower band at $t-2$ but is below it at $t-1$, a long-run counter is initiated; if the condition continues to hold, the counter is incremented, and a long position is opened once at least $d+1$ confirmations have been accumulated. Symmetrically, a short signal is generated when price crosses above the upper band and remains above it for at least $d+1$ confirmations. In the absence of a new signal, the previous position is maintained.

The rule also includes an optional holding-period parameter c. Once a new position is opened, it may be locked for c bars, during which intermediate band signals are ignored. After the holding counter expires, the rule resumes evaluating fresh Bollinger-band conditions. All positions are forcibly closed at the end of the trading day.

BB parameter grid (bars, band width, and holding period).

Parameter	Values
x	$\{0, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$
d	$\{0, 1, 2, 3, 4, 5\}$
c	$\{0, 1, 2, 5, 10, 15, 20\}$
n_σ	$\{1, 2, 3, 4\}$
j	$\{14, 20, 25, 30, 35, 40, 45, 50\}$

MACD rules (MACD and MACDZero). The MACD rules are based on exponential moving averages (EMAs) of the price series. For a span of m bars, the EMA is computed recursively as

$$\mathrm{EMA}_t(m) = \alpha_m P_t + (1 - \alpha_m)\mathrm{EMA}_{t-1}(m), \qquad \alpha_m = \frac{2}{m+1}. \tag{33}$$

Using a fast span f, a slow span s, and a signal span g, the MACD line and the signal line are defined as

$$\mathrm{MACD}_t = \mathrm{EMA}_t(f) - \mathrm{EMA}_t(s), \qquad \mathrm{SIG}_t = \mathrm{EMA}_t^{(\mathrm{MACD})}(g). \tag{34}$$

In the implementation, the baseline spans are fixed at $(f, s, g) = (12, 26, 9)$.

We consider two MACD variants:

- **MACD (signal-line rule).** A long signal is generated when the MACD line crosses above the signal line by at least x, and a short signal is generated when it crosses below the signal line by at least x. Operationally, the code detects a transition from a non-signal state at $t-2$ to a signal state at $t-1$, then requires the condition to persist until at least $d+1$ confirmations have been accumulated.

- **MACDZero (zero-line rule).** A long signal is generated when the MACD line crosses above the threshold $+x$, while a short signal is generated when it crosses below $-x$. As in the standard MACD rule, the condition must persist for at least $d+1$ confirmations before the position is opened.

In both variants, once a new position is established it may be locked for c bars, during which intermediate MACD signals are ignored. After the holding counter expires, the rule resumes evaluating new MACD conditions. In the absence of a new signal, the previous position is maintained. All positions are forcibly closed at the end of the trading day.

MACD parameter grid (thresholds and holding period).

Parameter	Values
x	$\{0, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$
d	$\{0, 1, 2, 3, 4, 5\}$
c	$\{0, 1, 2, 5, 10, 15, 20\}$
(f, s, g)	$(12, 26, 9)$

B Number of trading rules tested by rule variant

Table 6. Number of parameterized trading rules tested by rule variant

Rule variant	rules (per horizon)	Total
O1	120	600
O2	480	2,400
F1	210	1,050
F2	2,100	10,500
F3	1,050	5,250
MA1	330	1,650
MA2	990	4,950
MA3	1,650	8,250
MA4	4,950	24,750
MA5	4,950	24,750
CB1	600	3,000
CB2	2,400	12,000
SR1	432	2,160
SR2	1,728	8,640
BB	9,408	47,040
MACD	294	1,470
MACDZero	294	1,470
Total	31,986	159,930

Source: Authors' calculations from parameter grids implemented in this study.
Notes: Counts reflect the number of distinct parameter combinations per rule variant, computed separately for each horizon. The current workflow uses five horizons: 5s, 15s, 30s, 60s, and 120s. For F2, we impose $y < x$ and $d_y \leq d_x$. For MA3 and MA4, we impose $p < q$; for MA5, we impose $n < p < q$. Totals therefore sum to 31,986 rules per horizon and 159,930 rules across all five horizons.