

A Multivariate Covariate-Enhanced False Discovery Rate Method: Evidence from Technical Trading in Currency Markets

Ilias Filippou^a, Po-Hsuan Hsu^b, Tren Ma^c, Georgios Sermpinis^d, Mark P. Taylor^{e,f,*}

^a*Herbert Wertheim College of Business, Florida State University, FL 32306, USA*

^b*College of Technology Management, National Tsing Hua University, Taiwan*

^c*Nottingham University Business School, University of Nottingham, Nottingham NG8 1BB, UK*

^d*Adam Smith Business School, University of Glasgow, Glasgow G12 8QQ, UK*

^e*John M. Olin School of Business, Washington University in St. Louis, MO 63130-4899, USA*

^f*Centre for Economic Policy Research, Paris, France*

Abstract

We develop a multivariate functional false discovery rate (*mfFDR*) method for large-scale multiple testing with multiple informative covariates. Relative to existing *FDR* procedures, *mfFDR* delivers higher power while maintaining false-discovery control in the presence of dependence, correlated covariates, and covariate noise. We apply the method to technical trading in foreign exchange markets, using 635,850 rule-currency combinations over a 50-year period. Portfolios formed from *mfFDR*-selected strategies achieve Sharpe ratios close to one and remain profitable out of sample after transaction costs. We further show that this profitability declines with computational power, financial openness, market development and liquidity, linking the returns to technical trading to evolving market efficiency.

Keywords: data snooping, multiple testing, multivariate functional false discovery rate, informative covariates, technical trading, foreign exchange markets.

JEL classification: C12, C15, F31, G12, G15.

*Corresponding author. We would like to thank Gery Geenens for discussions on density estimations and conference participants at the FMA Annual Meeting 2022, the Paris Financial Management Conference 2022, the World Finance Banking Symposium 2023, the INQUIRE Conference 2024, and the Finance and Business Analytics Conference 2025 for their useful comments. All errors are our own.

Email addresses: ifilippou@fsu.edu (Ilias Filippou), pohsuanhsu@mx.nthu.edu.tw (Po-Hsuan Hsu), tren.ma@nottingham.ac.uk (Tren Ma), georgios.sermpinis@glasgow.ac.uk (Georgios Sermpinis), mark.p.taylor@wustl.edu (Mark P. Taylor)

1. Introduction

Empirical finance routinely evaluates large universes of trading strategies, return predictors, and forecasting models, creating acute data-snooping and false-discovery concerns. When researchers test hundreds or thousands of candidate signals, some will appear significant purely by chance, even when they contain no genuine predictive content. In such settings, controlling the false discovery rate (FDR)—the expected proportion of false discoveries among rejected hypotheses—is essential if statistical evidence is to yield economically meaningful inference. Since its introduction by [Benjamini and Hochberg \(1995\)](#) and subsequent development by [Storey \(2002\)](#), the FDR framework has become central to large-scale multiple testing and increasingly important in finance, where it has been applied to fund-performance evaluation and to broader concerns about false discoveries in asset pricing, trading-strategy analysis, and accounting research ([Barras, Scaillet and Wermers, 2010](#); [Harvey and Liu, 2020](#); [Harvey, Liu and Zhu, 2016](#); [Chordia, Goyal and Saretto, 2020](#); [Giglio, Liao and Xiu, 2021](#); [Jennings, Kim, Lee and Taylor, 2024](#)).

We develop a *multivariate functional false discovery rate* ($mfFDR$) procedure that conditions multiple testing on a vector of informative covariates. Standard FDR procedures impose a common rejection threshold across hypotheses and therefore ignore auxiliary information that traders and researchers naturally use when assessing whether a strategy is likely to be genuinely profitable. By contrast, our $mfFDR$ framework allows rejection thresholds to vary flexibly with the joint distribution of multiple observed covariates, improving statistical power while preserving information that would be lost under univariate screening or index-based dimension reduction.

A natural question in any covariate-adaptive multiple-testing procedure is how the covariates should be chosen. The *mfFDR* framework does not require a universal set of covariates; rather, it requires covariates that are available at the time of selection and informative about whether a hypothesis is likely to be truly non-null. Suitable covariates should therefore be economically interpretable and plausibly related either to the prior probability that the null is false or to the behavior of the test statistic under the alternative. In practice, they correspond to the sorts of characteristics that traders or portfolio managers might examine *ex ante* when deciding whether to employ a strategy, including persistence, abnormal return, benchmark exposure, liquidity, turnover, size, expenses, or measures of model fit, depending on the application.

Our paper builds on a natural progression in the literature. The “reality check” of [White \(2000\)](#), along with related procedures, asks whether any model outperforms a benchmark under resampling, but this class of global tests can be conservative in large model universes. The *FDR* methods of [Benjamini and Hochberg \(1995\)](#) and [Storey \(2002\)](#) instead control the expected proportion of false positives among rejected hypotheses, offering greater power in large-scale settings. The *functional FDR* (*fFDR*) framework of [Chen, Robinson and Storey \(2021\)](#) takes a further step by allowing rejection thresholds to vary with a single informative covariate.¹ Many finance applications, however, involve several informative characteristics that capture distinct dimensions of performance. In mutual fund evaluation, for example, these can include fund size, expense ratios, turnover, and past alpha persistence; in asset-

¹In particular, the covariate conveys information about both the prior probability that the null is true and the distribution of p -values under the alternative. [Hsu et al. \(2025\)](#) apply this idea to the conditional performance of mutual funds.

pricing or trading applications, they may include persistence, abnormal return, benchmark exposure, liquidity, and measures of model fit. In the technical trading application we study, persistence, abnormal return, benchmark exposure, and benchmark fit need not move together and cannot always be summarized adequately by a single index. Our $mfFDR$ procedure addresses this limitation by allowing rejection thresholds to depend on the joint distribution of multiple covariates rather than compressing them into a univariate summary. We show empirically that this distinction matters: collapsing the covariate vector into a principal component does not replicate the gains from genuinely multivariate conditioning.

A related strand of the literature addresses dependence across tests through different mechanisms. [Hansen \(2005\)](#) develops the Superior Predictive Ability test, which delivers a global decision on whether any model outperforms a benchmark under time-series dependence. [Lan and Du \(2019\)](#) address cross-test dependence by extracting latent common factors and applying factor-adjusted p -values. [Wang, Han and Tong \(2023\)](#) estimate the dependence structure directly and embed it in optimal decision rules that control false discoveries. Our approach complements these methods by using local, covariate-adaptive thresholds based on multiple *observed* covariates, thereby preserving their joint distribution and maintaining FDR control under realistic dependence without factor extraction or explicit dependence modeling.

We apply $mfFDR$ to a setting in which these issues are especially acute: technical trading in foreign exchange markets. The FX market is a natural laboratory because it combines enormous scale—daily turnover averaged \$9.6 trillion in April 2025 according to the BIS Triennial Survey—with an exceptionally large strategy space, substantial dependence across candidate trading rules, and a long-standing debate over whether technical analysis gener-

ates genuine abnormal performance (Bank for International Settlements, 2025; Allen and Taylor, 1990; Taylor and Allen, 1992; Neely, Weller and Dittmar, 1997; Qi and Wu, 2006; Menkhoff and Taylor, 2007; Hsu, Taylor and Wang, 2016). It therefore provides a demanding environment in which to assess whether a multivariate, covariate-adaptive multiple-testing procedure can separate genuine signals from noise more effectively than existing approaches.

We evaluate $mfFDR$ in a universe of more than 21,000 technical trading rules applied to 30 U.S.-dollar exchange rates, yielding as many as 635,850 rule-currency combinations over a nearly fifty-year period. Trading rules are assessed in rolling 12-month windows, and the selected-rule portfolios are then evaluated out of sample, both before and after transaction costs. We focus on the Sharpe ratio as our primary performance metric, while also considering mean returns and three-factor alphas in robustness checks. Across Monte Carlo designs, $mfFDR$ controls FDR well and delivers substantial power gains relative to standard FDR and single-covariate $fFDR$, with the gains especially large in weak-signal environments. The method also remains robust when test statistics are cross-dependent and when informative covariates are correlated or measured with error.

Our empirical implementation uses four covariates that capture complementary dimensions of trading-rule quality: ρ , the autocorrelation of rule returns; α , abnormal performance relative to a benchmark; β , benchmark exposure; and R^2 , benchmark fit.² Using this richer information set, $mfFDR$ produces materially stronger out-of-sample results than competing procedures. Portfolios formed from $mfFDR$ -selected strategies achieve annualized Sharpe

²These covariates are not structural regressors in the spirit of traditional factor models (Fama and French, 1993). Rather, they improve the efficiency of the multiple-testing procedure by capturing conditional information about whether a trading rule is likely to deliver genuinely abnormal performance.

ratios of 1.06 before transaction costs and 0.95 after transaction costs. Moreover, principal-component compression of the covariates fails to reproduce these gains, underscoring the value of preserving their joint information.³

The FX application also yields broader economic insights. We show that the profitability of *mfFDR*-portfolios selected on the basis of technical strategies declines with computational power, financial openness, market development, and FX market liquidity, but rises with volatility (Menkhoff, Sarno, Schmeling and Schrimpf, 2012a), consistent with theories of adaptive efficiency and evolving information processing in currency markets (Neely, Weller and Ulrich, 2009; Neely and Weller, 2013; Hsu, Taylor and Wang, 2016). At the same time, *mfFDR* adapts to changing opportunity sets by reallocating capital toward those rules and currencies in which technical signals remain effective. Although simple technical predictability has eroded in many developed currencies, significant rules remain detectable even in 2011–2020, especially in more volatile or less developed markets.

Our paper therefore makes three contributions. First, we introduce a multivariate covariate-adaptive *FDR* procedure suited to finance applications involving many competing signals and several informative characteristics. Second, we provide evidence from a very large FX trading-rule universe that multivariate conditioning yields economically meaningful gains in out-of-sample strategy selection and portfolio performance. Third, we show that the profitability of technical trading is systematically related to market conditions, linking conditional trading-rule performance to computational capacity, openness, volatility, liquidity,

³Menkhoff, Sarno, Schmeling and Schrimpf (2012b) similarly study technical trading across currencies and relate it to cross-sectional momentum. They show that portfolios formed from technical rules are profitable, but that currency momentum is not simply a reflection of standard technical trading strategies.

and market development.

The rest of the paper is organized as follows. Section 2 develops the *mfFDR* framework and presents simulation evidence. Sections 3 and 4 describe the data and performance measures. Section 5 presents the main empirical results, and Section 6 examines the market environment and economic drivers of technical profitability. Section 7 concludes.

2. The Use of Covariates in the *FDR* Framework

This section develops the *multivariate functional false discovery rate* (*mfFDR*) framework. We first present the method and its intuition in the context of large-scale strategy evaluation. We then use simulation evidence to compare its *FDR* control and power with those of existing approaches. Finally, we examine its robustness under features that are common in financial data, including dependence across tests, correlated covariates, and noisy covariate measurement.

2.1. Method Overview

Suppose we have n hypotheses, each associated with a performance metric ϕ . In our application, ϕ is the Sharpe ratio, and the object of interest is whether a given trading rule delivers genuinely abnormal performance. For each hypothesis, we observe a vector of covariates that may contain information about the likelihood that the null hypothesis is true. We therefore consider the testing problem

$$H_0 : \phi = 0, \quad H_1 : \phi \neq 0. \tag{1}$$

Across n hypotheses, this yields a large-scale multiple-testing problem in which the objective is to detect as many true nonzero ϕ as possible while controlling the FDR , that is, the expected proportion of false discoveries among rejected hypotheses (Benjamini and Hochberg, 1995).

For notational convenience, consider one such test. Let P denote the corresponding p -value, let $\mathbf{Z} = (Z^1, \dots, Z^d)$ denote a d -dimensional covariate vector, and let $h \in \{0, 1\}$ indicate whether the null is false. Subscript i indexes hypotheses, so that $(P_i, h_i, \mathbf{Z}_i)$ denotes the observed tuple for hypothesis i .

In standard FDR procedures (e.g., Benjamini and Hochberg, 1995), all hypotheses face a common rejection threshold. By contrast, $mfFDR$ allows the threshold to vary with the realized covariates. Intuitively, hypotheses with more favorable covariate profiles can face more lenient thresholds, whereas hypotheses with less favorable profiles face stricter ones. This covariate-adaptive screening improves power while preserving FDR control.

A practical consideration in implementing $mfFDR$ is the choice of covariates. The framework does not require a universal set of covariates; rather, it requires covariates that can be constructed using information available at the time of selection and that are informative about whether a hypothesis is likely to be non-null. Suitable covariates should therefore be economically interpretable and plausibly related either to the prior probability that the null is false or to the distribution of the test statistic under the alternative. In practice, such covariates correspond to the kinds of characteristics that traders or portfolio managers might examine ex ante when deciding whether to deploy a strategy, including persistence, abnormal return, benchmark exposure, liquidity, turnover, size, expenses, or measures of model fit. Our empirical implementation follows this principle.

To implement this idea, each covariate is rank-transformed to lie in $[0, 1]$, so that $z_i^j = r_i^j/n$, where r_i^j is the rank of Z_i^j among $\{Z_1^j, \dots, Z_n^j\}$. For exposition, we assume that: (i) the covariates are mutually independent; (ii) conditional on $\mathbf{Z} = \mathbf{z}$, the null implies $P \sim U(0, 1)$, while under the alternative P has conditional density $f_{\text{alt}}(p \mid \mathbf{z})$; and (iii) the tuples $(h_i, P_i, \mathbf{Z}_i)$ are i.i.d. across i . As we show below and in the Internet Appendix, the method remains effective when these conditions are relaxed, including when test statistics are dependent and covariates are correlated or noisy.

The key object is the conditional probability that a hypothesis is null, given the observed $(p, \mathbf{z})$:

$$r(p, \mathbf{z}) = \mathbb{P}(h = 0 \mid P = p, \mathbf{Z} = \mathbf{z}) = \frac{\pi_0(\mathbf{z})}{f(p, \mathbf{z})}, \quad (2)$$

where $\pi_0(\mathbf{z})$ is the prior null probability and $f(p, \mathbf{z})$ is the joint density of $(P, \mathbf{Z})$.⁴ Intuitively, $r(p, \mathbf{z})$ is the posterior probability that a hypothesis is null after combining its p -value with its covariate profile.

Empirically, we estimate

$$\hat{r}(p, \mathbf{z}) = \frac{\hat{\pi}_0(\mathbf{z})}{\hat{f}(p, \mathbf{z})},$$

where $\hat{\pi}_0(\mathbf{z})$ and $\hat{f}(p, \mathbf{z})$ are estimated nonparametrically. Appendix A provides implementation details. In practice, the procedure is straightforward: estimate the local null probability, estimate the joint density of the p -value and the covariates, construct $\hat{r}(p, \mathbf{z})$ for each hypothesis, and reject hypotheses whose associated *mfFDR* q -values fall below the desired target level.

⁴Under the maintained assumptions, the conditional density of P under the null, given $\mathbf{Z} = \mathbf{z}$, is equal to one on $[0, 1]$. In addition, the rank transformation implies a uniform covariate density. Bayes' rule therefore yields (2).

To control FDR at target $\tau \in (0, 1)$, hypotheses are ranked by increasing $\hat{r}(p, \mathbf{z})$ and sequentially rejected until the estimated FDR reaches τ . The rejection region is

$$\Gamma(\theta) = \{(p, \mathbf{z}) : \hat{r}(p, \mathbf{z}) \leq \theta\}, \quad (3)$$

where θ is chosen such that

$$\widehat{FDR}_\theta = \frac{1}{\#S_\theta} \sum_{k \in S_\theta} \hat{r}(p_k, \mathbf{z}_k) \leq \tau,$$

with $S_\theta = \{k : \hat{r}(p_k, \mathbf{z}_k) \leq \theta\}$.⁵

If hypothesis i is rejected, then all hypotheses with $\hat{r}(p, \mathbf{z}) \leq \hat{r}(p_i, \mathbf{z}_i)$ are also rejected. Following [Chen, Robinson and Storey \(2021\)](#) and [Storey \(2003\)](#), we therefore compute for each hypothesis a *functional q-value*: the estimated FDR that would result from rejecting all hypotheses with posterior-null estimates no larger than its own. A hypothesis is rejected if and only if its functional q -value is no greater than τ .⁶

We refer to this procedure as the *multivariate functional false discovery rate* ($mfFDR$) method. When $d = 1$, it reduces to the $fFDR$ procedure of [Chen, Robinson and Storey \(2021\)](#). The key distinction is that $mfFDR$ conditions on the joint distribution of multiple informative covariates rather than collapsing them into a single dimension. We now quantify the gains from doing so.

⁵Here $\#S_\theta$ is the number of discoveries, while the numerator is the estimated number of false discoveries. Related estimators are used by [Newton, Noueiry, Sarkar and Ahlquist \(2004\)](#), [Storey, Akey and Kruglyak \(2005\)](#), and [Chen, Robinson and Storey \(2021\)](#).

⁶If $f(p, \mathbf{z})$ is decreasing in p for fixed $\mathbf{z}$, then $r(p, \mathbf{z})$ is increasing in p , and the rule $\hat{r}(p_i, \mathbf{z}_i) \leq \theta^*$ is equivalent to a covariate-specific threshold rule $p_i \leq \Theta(\mathbf{z}_i)$.

2.2. Simulation Studies

We begin with a simple setting involving two informative covariates, $\mathbf{Z} = (U, V)$. We simulate $n = 10,000$ hypotheses, with approximately 80% of the null hypotheses true.⁷ The purpose of the exercise is to compare *mfFDR*, which uses both covariates jointly, with *fFDR*, which uses only one covariate, and with standard *FDR* procedures.

In each simulation, the covariates $U = (u_1, \dots, u_n)$ and $V = (v_1, \dots, v_n)$ are drawn independently from $U(0, 1)$. For each hypothesis i , the null status is then drawn from

$$h_i \sim \text{Bernoulli}(1 - \pi_0(u_i, v_i)),$$

where the null probability function takes one of two forms:

- $\pi_0(u, v) = \sin[\pi(u + v)/2]$ (non-monotonic sine form);
- $\pi_0(u, v) = 1 - (u^4 + v^4)/2$ (monotonic form).

These specifications allow us to assess performance under both nonlinear and monotonic dependence of the null probability on the covariates.

Conditional on (U, V, H) , p -values are generated as follows. If $h_i = 0$, then $p_i \sim U(0, 1)$. If $h_i = 1$, then p_i is drawn from a Beta distribution,

$$f_{\text{alt}}(p \mid u, v) \propto p^{\alpha-1}(1-p)^{\beta-1},$$

with $\beta = 3 + 1.5(u + v)$, following [Chen, Robinson and Storey \(2021\)](#). We consider

⁷Section [ID](#) of the Internet Appendix shows that the main conclusions are robust to alternative values of n .

$\alpha \in \{0.5, 1, 1.5\}$, corresponding to strong, moderate, and weak signal-to-noise settings, respectively.⁸

For each design, we implement *mfFDR* (using both U and V), *fFDR* (using only U), the standard *FDR* procedure of Storey (2002) (*stdFDR*), and the original Benjamini and Hochberg (1995) procedure. Using the known null-status vector, we compute the false discovery proportion (*FDP*) and the correct detection proportion (power). We consider target levels $\tau \in \{0.05, 0.10, \dots, 0.95\}$ and repeat the simulation 1,000 times. This yields six baseline cases: two specifications for $\pi_0(u, v)$ crossed with three signal strengths. We then extend the analysis to settings with covariate correlation, measurement error, and dependence across tests.

Figure 1 reports the results. The upper panels plot actual *FDR*, and the lower panels plot power, under both specifications of $\pi_0(u, v)$ and all three signal strengths. All procedures control *FDR* reasonably well across target levels, with only modest finite-sample variation.


The main differences arise in power. *mfFDR* consistently outperforms *fFDR*, *stdFDR*, and BH across all cases, with the gains especially strong when signals are weak. Under strong signals, *mfFDR* improves power by as much as 21%, 29%, and 37% relative to *fFDR*, *stdFDR*, and BH, respectively. Under moderate and weak signals, the gains are larger still: up to 44%, 67%, and 76% in the moderate-signal case, and 57%, 62%, and 62% in the weak-signal case. At the 20% target level used in our main empirical analysis, *mfFDR* improves on *fFDR* by roughly 10–20%.

These findings show that preserving the joint structure of multiple informative covariates


⁸Section IA of the Internet Appendix provides additional illustrations of these signal strengths.

Figure 1: Performance comparison of $mfFDR$, $fFDR$, standard FDR of Storey (2002) ($stdFDR$), and the BH procedure. Panel A: sine form of $\pi_0(u, v)$. Panel B: monotonic form. Top sub-figures show actual FDR estimates; bottom sub-figures show power.

Panel A: $\pi_0(u, v)$ is sine-shaped.



Panel B: $\pi_0(u, v)$ is monotonic.



yields materially higher detection power than either univariate conditioning or unconditional screening. This advantage is especially valuable in low signal-to-noise settings, where distinguishing genuine signals from noise is most difficult and where many finance applications operate in practice.

2.3. Robustness Checks

The baseline simulations assume independent, uniformly distributed covariates and independent tests. These assumptions are convenient for exposition, but they are restrictive in finance applications. We therefore conduct additional simulations under three departures from the baseline design.

First, we introduce correlation across covariates (Section [IE](#) of the Internet Appendix). We find that $mfFDR$ continues to control FDR well for moderate covariate correlation and asymptotically even at higher levels of correlation. The power gains relative to alternative procedures remain intact.

Second, we allow the covariates to be measured with error (Section [IF](#) of the Internet Appendix). This is important because, in empirical finance, covariates are typically estimated from return histories and are therefore noisy. In these simulations, $mfFDR$ continues to control FDR well. Although power falls somewhat relative to the case of perfectly observed covariates, the method still substantially outperforms competing procedures.

Third, we introduce weak dependence across tests (Section [IG](#) of the Internet Appendix). This setting is especially relevant for technical trading, where rules with nearby parameter values often generate highly correlated returns and therefore dependent test statistics. Even in this non-i.i.d. environment, $mfFDR$ remains robust, maintaining FDR control while con-

tinuing to deliver meaningful power gains.

Taken together, these robustness checks show that $mfFDR$ performs well under conditions that closely resemble empirical finance applications, including correlated covariates, noisy covariate measurement, and dependence across tests.

3. Data Sources and the Trading Strategy Universe

3.1. Data

Our empirical analysis uses daily spot and one-month forward exchange-rate data against the U.S. dollar for 30 countries, comprising nine developed economies and 21 emerging markets.⁹ For the developed currencies, the sample runs from January 4, 1971, to December 31, 2020. For the emerging-market currencies, the start date varies with data availability for exchange rates and short-term interest rates. In particular, the sample begins in January 1978 for Israel, January 1981 for South Africa, January 1982 for Singapore, and October 1983 for Taiwan. All emerging-market series end in December 2020. The exchange-rate and short-term interest-rate data were provided by the London branch of BlackRock and are based on London-market quotations.

Table 1 reports summary statistics for gross currency returns and short-term interest rates in developed and emerging markets. We define the gross return from holding one unit

⁹The developed currencies are the Australian dollar, Canadian dollar, German mark/euro, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, Swiss franc, and U.K. pound. The emerging-market currencies are the Argentine peso, Brazilian real, Chilean peso, Colombian peso, Czech koruna, Hungarian forint, Indian rupee, Indonesian rupiah, Israeli shekel, Korean won, Mexican peso, Philippine peso, Polish zloty, Romanian new leu, Russian ruble, Singaporean dollar, Slovak koruna, South African rand, Taiwanese dollar, Thai baht, and Turkish lira.

of foreign currency for one day as

$$r_t = \ln(s_t/s_{t-1}),$$

where s_t denotes the spot exchange rate at time t . The exchange rate is measured as U.S. dollars per unit of foreign currency, so an increase in s_t corresponds to an appreciation of the foreign currency. Table 1 reports the mean, standard deviation, minimum, maximum, and sample period for each currency. As expected, gross returns are generally more volatile in emerging markets, with wider tails and more extreme observations.

Short-term interest rates are also important because they enter total currency returns, even though technical rules primarily respond to exchange-rate movements. We convert annualized short-term interest rates, i_t^a , into daily rates via

$$i_t = \ln(1 + i_t^a)/360.$$

Table 1 shows that short-term rates are relatively low and stable in developed markets, with average daily rates ranging from 0.45 to 1.94 basis points. In emerging markets, by contrast, short-term rates are typically both higher and more variable. For example, Turkey has the highest average daily rate in the sample, at 8.33 basis points, whereas Chile has the lowest, at 0.14 basis points.

Table 1: Summary Statistics. The table presents summary statistics for gross returns on foreign currencies and short-term interest rates. We report the mean, volatility, minimum, and maximum, together with the sample period for each currency. Results are shown separately for developed and emerging economies.

Country	Gross returns on foreign currencies				Short-term interest rates				Sample Period
	mean(%)	vol	min	max	mean(%)	vol(%)	min(%)	max(%)	
Developed									
Australia	-0.0028	0.0068	-0.1925	0.1073	0.0191	0.0119	0.0001	0.1709	1/4/1971–12/31/2020
Canada	-0.0018	0.0041	-0.0434	0.0505	0.0172	0.0125	0.0005	0.0596	1/4/1971–12/31/2020
Germany/E.U.	-0.0017	0.0058	-0.0421	0.0462	0.0092	0.0078	-0.0020	0.0344	1/4/1971–12/31/2020
Japan	0.0095	0.0063	-0.0626	0.0950	0.0108	0.0124	-0.0006	0.0406	1/4/1971–12/31/2020
New Zealand	-0.0033	0.0073	-0.2050	0.0995	0.0194	0.0111	-0.0007	0.1286	1/4/1971–12/31/2020
Norway	-0.0014	0.0066	-0.0814	0.0646	0.0163	0.0106	0.0023	0.1372	1/4/1971–12/31/2020
Sweden	-0.0036	0.0065	-0.1507	0.0555	0.0171	0.0139	-0.0011	0.2090	1/4/1971–12/31/2020
Switzerland	0.0122	0.0073	-0.0892	0.1267	0.0045	0.0098	-0.0086	0.2276	1/4/1971–12/31/2020
U.K.	-0.0044	0.0058	-0.0848	0.0467	0.0160	0.0112	0.0004	0.0460	1/4/1971–12/31/2020
U.S.	–	–	–	–	0.0128	0.0097	-0.0000	0.0466	1/4/1971–12/31/2020
Emerging									
Argentina	-0.0579	0.0098	-0.3418	0.1712	0.0343	0.0368	0.0035	0.4144	4/1/1991–12/31/2020
Brazil	-0.0238	0.0097	-0.1080	0.1178	0.0428	0.0283	0.0051	0.2792	7/4/1994–12/31/2020
Chile	-0.0070	0.0064	-0.1160	0.1114	0.0014	0.0016	0.0000	0.0214	1/3/1994–12/31/2020
Colombia	-0.0224	0.0062	-0.0593	0.0562	0.0315	0.0240	0.0052	0.0908	1/6/1992–12/31/2020
Czech	0.0040	0.0073	-0.0707	0.0522	0.0104	0.0133	0.0002	0.2628	4/22/1992–12/31/2020
Hungary	-0.0177	0.0080	-0.0842	0.0520	0.0297	0.0227	0.0021	0.0834	6/4/1991–12/31/2020
India	-0.0178	0.0046	-0.1281	0.0376	0.0254	0.0149	0.0002	0.1944	1/3/1986–12/31/2020
Indonesia	-0.0276	0.0130	-0.3576	0.2361	0.0296	0.0242	0.0000	0.2054	1/1/1991–12/31/2020
Israel	-0.0683	0.0059	-0.1725	0.0645	0.0555	0.0962	0.0002	0.6309	1/4/1978–12/31/2020
South Korea	-0.0047	0.0078	-0.1809	0.2012	0.0154	0.0138	0.0012	0.0676	1/3/1992–12/31/2020
Mexico	-0.0349	0.0114	-0.2231	0.2231	0.0403	0.0416	0.0077	0.3387	1/1/1987–12/31/2020
Philippines	-0.0096	0.0044	-0.0860	0.1015	0.0198	0.0165	0.0014	0.1962	1/2/1987–12/31/2020
Poland	-0.0117	0.0079	-0.0715	0.0670	0.0212	0.0198	-0.0004	0.0792	6/4/1993–12/31/2020
Romania	-0.0368	0.0091	-0.3887	0.0953	0.0407	0.0518	0.0006	0.3667	1/2/1997–12/31/2020
Russia	-0.0515	0.0132	-0.3863	0.2779	0.0288	0.0391	0.0022	0.3208	9/1/1994–12/31/2020
South Africa	-0.0286	0.0097	-0.1030	0.1440	0.0289	0.0115	0.0000	0.0588	1/2/1981–12/31/2020
Singapore	0.0043	0.0033	-0.0276	0.0414	0.0053	0.0050	0.0000	0.0181	4/27/1993–12/31/2020
Slovakia	0.0021	0.0063	-0.1097	0.0462	0.0130	0.0164	0.0001	0.2068	1/4/1982–12/31/2020
Taiwan	0.0035	0.0029	-0.0420	0.0430	0.0082	0.0067	0.0004	0.0483	10/4/1983–12/31/2020
Thailand	-0.0022	0.0055	-0.2077	0.0635	0.0102	0.0107	0.0004	0.0664	1/2/1991–12/31/2020
Turkey	-0.1002	0.0109	-0.3348	0.1256	0.0833	0.0643	0.0130	1.0328	2/1/1990–12/31/2020

3.2. The Five Families of Technical Trading Rules

Foreign exchange markets provide a demanding environment in which to evaluate $mfFDR$ because they combine a large candidate strategy space, acute data-snooping concerns, and an active, long-standing debate over the profitability of technical analysis (Allen and Taylor, 1990; Cheung and Chinn, 2001; Qi and Wu, 2006; Menkhoff and Taylor, 2007; Hsu, Taylor and Wang, 2016). To operationalize this setting, we focus on five widely used families of technical trading rules: filter rules, moving average rules, relative strength indicator rules, support-resistance rules, and channel breakout rules.¹⁰ These families span the main types of technical signals used in practice—trend-following, momentum, reversal, and breakout signals—while keeping the baseline strategy universe transparent and economically interpretable. The parameter ranges for these families are chosen with reference to both the academic literature and market practice. Full rule definitions and parameter grids are reported in Section IC of the Internet Appendix.

Filter rules. Filter rules are among the earliest and simplest technical trading strategies, dating back to Alexander (1961) and widely applied to foreign exchange markets following the move to floating rates (e.g., Dooley and Shafer, 1976; Logue, Sweeney and Willett, 1978; Cornell and Dietrich, 1978; Sweeney, 1986; Levich and Thomas, 1993). A filter rule generates a long (short) position when the spot exchange rate rises (falls) by a specified percentage above (below) its most recent low (high). The aim is to capture sustained directional moves while filtering out small, potentially noisy fluctuations.

Moving average rules. Moving average (MA) rules are among the most widely used trend-

¹⁰See, for example, Taylor and Allen (1992); Neely and Weller (1999); Menkhoff and Taylor (2007); Neely and Weller (2012); Hsu, Taylor and Wang (2016) for discussions of their widespread use in FX markets.

following strategies in FX markets (Taylor and Allen, 1992; Henderson, 2006; Menkhoff and Taylor, 2007; Neely and Weller, 2012). These rules generate signals from comparisons of the spot exchange rate with moving averages of different lengths, or from comparisons across moving averages themselves. A simple MA rule, for example, takes a long (short) position when the spot rate rises (falls) sufficiently above (below) a moving average of length j . A dual MA rule generates a signal when a short-horizon moving average crosses above (below) a longer-horizon moving average.

Relative strength indicator rules. Relative strength indicator (RSI) rules form a major class of oscillator strategies designed to identify overbought and oversold conditions (Levy, 1967; Wilder, 1978; Henderson, 2006). The idea is to detect potential reversals by measuring the balance of recent gains and losses. For a window length h , the RSI takes values between 0 and 100, with buy (sell) signals typically triggered when the index crosses a lower (upper) threshold. The standard RSI is given by

$$RSI_t = 100 \times \frac{U_t}{U_t + D_t}, \quad (4)$$

where U_t and D_t denote cumulative up-movements and cumulative absolute down-movements over the preceding h days. Standard threshold values are $m = 70$ for overbought and $n = 30$ for oversold conditions.

Support-resistance rules. Support-resistance rules exploit historically important price levels at which the exchange rate has repeatedly stalled or reversed (Murphy, 1999). When the spot rate breaks through such support or resistance levels, these rules interpret the move as evidence of continued momentum, potentially driven by order clustering and stop-loss

behavior (Osler, 2003, 2005), and generate a corresponding long or short position.

Channel breakout rules. Channel breakout rules extend the support-resistance logic by constructing upper and lower trading bands from rolling windows of past prices. A break above the upper band or below the lower band is interpreted as the start of a new trend and triggers a long or short position. These strategies are particularly relevant in FX markets, where periods of consolidation are often followed by abrupt directional breakouts (Murphy, 1999; Osler, 2003).

Across these five rule families, we generate a large universe of strategy variants by varying signal thresholds, delay parameters, and holding periods. The resulting universe contains 2,835 filter rules, 12,870 moving average rules, 600 RSI rules, 1,890 support-resistance rules, and 3,000 channel breakout rules, for a total of 21,195 distinct trading rules. Additional details on the construction and parameterization of the rule universe are reported in Section IC of the Internet Appendix.

The next section describes the performance measures used to evaluate these rules, including the treatment of transaction costs and the construction of out-of-sample portfolios.

4. Measures of Predictive Ability and Profitability

We distinguish between the gross excess return generated by a trading strategy and the corresponding net excess return after transaction costs. The one-day excess return from holding one unit of foreign currency against the U.S. dollar is

$$r_t = \ln(s_t/s_{t-1}) + \ln(1 + i_{t-1}^*) - \ln(1 + i_{t-1}), \quad (5)$$

where s_t and s_{t-1} are spot exchange rates on days t and $t - 1$, and i_{t-1} and i_{t-1}^* are the corresponding daily U.S. and foreign short-term interest rates observed at $t - 1$.

The excess return generated by trading rule j from day $t - 1$ to day t is then

$$R_{j,t} = S_{j,t-1} r_t, \quad (6)$$

where $S_{j,t-1}$ denotes the position prescribed by the rule using information available at time $t - 1$. The signal takes values in $[-1, 1]$, with positive values corresponding to long foreign-currency positions, negative values to short positions, and zero to neutrality. For most rules, $S_{j,t-1} \in \{+1, -1, 0\}$, although some moving-average rules generate fractional positions such as $+1/3$ or $-1/3$.¹¹

We evaluate each rule using its Sharpe ratio,

$$SR_j = \frac{\bar{R}_j}{\sigma_j}, \quad (7)$$

where $\bar{R}_j$ is the sample mean of rule returns and σ_j is the corresponding sample standard deviation over the evaluation window $[T_1, T_2]$:

$$\bar{R}_j = \frac{1}{N} \sum_{t=T_1}^{T_2} R_{j,t}, \quad \sigma_j = \sqrt{\frac{1}{N-1} \sum_{t=T_1}^{T_2} (R_{j,t} - \bar{R}_j)^2},$$


with $N = T_2 - T_1 + 1$.

¹¹See Internet Appendix Section [IC](#). With an initial one-dollar portfolio, a signal $S_{t-1} \in [-1, 1]$ implies an excess return of $S_{t-1}r_t$. When $S_{t-1} > 0$, the trader allocates S_{t-1} dollars to the foreign currency and the remainder to U.S. bonds. When $S_{t-1} < 0$, the trader borrows in the foreign currency and invests in U.S. bonds accordingly.

Transaction costs arise whenever a trading rule changes its position. We define net excess return as excess return net of estimated trading costs. Following [Neely and Weller \(2013\)](#), our baseline daily transaction-cost measure is one-third of the quoted one-month forward bid-ask spread. For earlier periods in which forward bid-ask data are unavailable, we follow [Neely and Weller \(2013\)](#) in assigning fixed costs of 5 basis points per trade in the 1970s, 4 basis points in the 1980s, and 3 basis points in the 1990s for developed currencies. For emerging-market currencies, we use one-third of the average of the first 500 available bid-ask spreads.

As a preliminary overview, we examine the distribution of Sharpe ratios and average returns for all trading rules across rolling 10-year sub-periods. In each sub-period, this yields as many as 635,850 rule-currency observations. Figure 2 plots the corresponding quantiles and means. Panel A reports gross and net Sharpe ratios, while Panel B reports annualized average returns. Two features stand out. First, the figure shows a marked decline in the profitability of technical trading over time, consistent with the evidence in [Menkhoff and Taylor \(2007\)](#). This decline is visible not only in the center of the distribution but also across the upper quantiles, indicating that even the better-performing rules became less profitable over time. Second, despite this broad erosion, a subset of rules continues to exhibit comparatively strong performance, especially in the upper tail and in the later part of the sample. The key question, therefore, is whether such apparently profitable rules could have been identified *ex ante* rather than merely recognized *ex post*. The next section addresses that question using the *mfFDR* framework.

Figure 2: Evolution of Sharpe ratios and average returns (annualized) of the trading rules across 10-year sub-periods. Dashed (solid) lines in Panel A present the 5%, 25%, 50%, 75%, and 95% quantiles and the mean Sharpe ratio before (after) transaction costs. Panel B shows the corresponding statistics for annualized average returns.



5. Empirical Results

For each trading rule i , we assess performance using the hypothesis test

$$H_0 : \phi_i = 0, \quad H_1 : \phi_i \neq 0, \quad (8)$$

where ϕ_i is the Sharpe ratio of trading rule i before transaction costs. We focus on the Sharpe ratio as our primary performance metric because standard factor structures are less well established in FX trading (Ivanova, Neely, Weller and Famiglietti, 2021), while the Sharpe ratio remains the most widely used summary measure of risk-adjusted trading performance.

For robustness, the Internet Appendix also reports results based on mean returns and three-

factor alphas; see Sections [IL](#) and [IM](#).

Our empirical analysis proceeds in three steps. We first introduce the covariates that inform the *mfFDR* procedure. We then identify outperforming rules both within individual currencies and across the full currency universe. Finally, we form rolling portfolios from the selected rules and evaluate them out of sample. If technical trading rules contain genuine predictive content, portfolios formed from rules identified as profitable in sample should deliver positive out-of-sample performance.

To implement this idea, we apply *mfFDR* to trading rules with positive estimated Sharpe ratios while controlling the *FDR* at a target level. We refer to this procedure as *mfFDR*⁺. For comparison, we also implement the single-covariate analogue, *fFDR*⁺. For brevity, we report results for a 20% *FDR* target in the main text; Section [IH](#) of the Internet Appendix shows that the main conclusions are robust to alternative targets.

5.1. Covariates

Our empirical implementation uses four continuous covariates constructed from the return series of each technical trading rule. These covariates are chosen both for their economic interpretability and for their ability to capture dimensions of performance that practitioners routinely monitor. Taken together, they summarize three broad aspects of signal quality: time-series persistence, abnormal risk-adjusted performance, and deviation from passive benchmarks.

The first covariate is the first-order autocorrelation, ρ , of each rule’s return series. This captures short-run persistence in trading-rule returns and is closely related to the idea that profitable technical signals may reflect local trend continuation or delayed adjustment

(Moskowitz, Ooi and Pedersen, 2012). From a statistical perspective, persistence may enhance predictability; from an economic perspective, it may reflect time-varying risk premia or investor underreaction.

The remaining three covariates are obtained from rolling regressions of rule returns on passive benchmarks. For the within-currency analysis, we estimate

$$r_{i,t} = \alpha_{i,bh} + \beta_{i,bh}r_{bh,t} + \varepsilon_{i,t}, \quad (9)$$

where $r_{bh,t}$ is the excess return on a buy-and-hold strategy in the same currency. The intercept, $\alpha_{i,bh}$, measures abnormal performance relative to the passive benchmark; the slope, $\beta_{i,bh}$, captures benchmark exposure; and R_{bh}^2 measures the extent to which the rule's return variation is explained by the passive benchmark. Lower R_{bh}^2 values indicate more idiosyncratic return profiles and hence greater potential diversification benefits in a multi-rule portfolio (Grinold and Taylor, 2009). Because these regressions are estimated on rolling windows, these characteristics are also allowed to vary over time.

For the pooled cross-currency analysis, we benchmark each rule against the global currency factor RX_t of Lustig, Roussanov and Verdelhan (2011), defined as the excess return on an equal-weighted portfolio of all 30 currencies:

$$r_{i,t} = \alpha_{i,mk} + \beta_{i,mk}RX_t + \varepsilon_{i,t}. \quad (10)$$

This yields the analogous covariates $\alpha_{i,mk}$, $\beta_{i,mk}$, and R_{mk}^2 , which capture abnormal performance, systematic exposure, and benchmark fit relative to common global currency move-

ments.

These four covariates have three useful features. First, they are economically interpretable, mapping naturally into persistence, alpha, benchmark exposure, and diversification. Second, they are grounded in standard portfolio logic: rules with low benchmark correlation can be especially valuable in diversified trading-rule portfolios (Grinold and Taylor, 2009). Third, they are empirically relevant, since persistence, abnormal return, and systematic exposure are all commonly associated with cross-sectional differences in strategy profitability (Jegadeesh and Titman, 1993; Moskowitz, Ooi and Pedersen, 2012; Harvey, Liu and Zhu, 2016).

In what follows, we use ρ , α_{bh} , β_{bh} , and R_{bh}^2 to form portfolios at the individual-currency level, and ρ , α_{mk} , β_{mk} , and R_{mk}^2 to form portfolios in the pooled cross-currency analysis. We then examine whether conditioning on these covariates improves the identification of genuinely profitable technical trading rules.

5.2. Individual Currencies

We begin by evaluating $mfFDR^+$ at the individual-currency level, using R_{bh}^2 , α_{bh} , β_{bh} , and ρ as covariates. Using daily data, we construct portfolios monthly as follows. For each currency, we use the most recent 12 months of observations to compute, for every trading rule, the p -value for the hypothesis in equation (8), together with the associated covariates.¹²

We then apply $mfFDR^+$ to select outperforming rules at a target FDR of 20%. The signals from the selected rules are aggregated by netting out opposing positions. For example,

¹²We also evaluate performance using mean returns and three-factor alphas; see Sections [IL](#) and [IM](#). Following [Bajgrowicz and Scaillet \(2012\)](#), we compute p -values for Sharpe ratios and mean returns using the stationary bootstrap. For alpha-based tests, we use HAC-robust t -statistics ([Newey and West, 1987](#)) for computational efficiency.

suppose 100 rules are selected, of which 30 issue buy signals (+1), 20 issue sell signals (−1), and 50 issue neutral signals (0).¹³ The resulting aggregate signal is +10, which corresponds to a 10% long position in the foreign currency.¹⁴

Each portfolio is then held out of sample for one month. We compare the resulting performance with that of portfolios formed using FDR^+ and $fFDR^+$, where the latter conditions on a single covariate. Selection is based on pre-cost excess returns, whereas out-of-sample performance is evaluated using post-cost net returns.¹⁵

Table 2 reports the pairwise correlations among the four covariates used in $mfFDR^+$. The correlations are generally modest: 98% lie within $[-0.69, 0.59]$, and only 1% exceed $|0.7|$. As shown in Section IE of the Internet Appendix, $mfFDR$ continues to control the false discovery rate well under this degree of dependence.

Table 2: Pairwise correlation of covariates. The table reports quantiles of the correlation coefficients, calculated monthly using one year of in-sample data, for the six pairs formed from the four covariates R_{bh}^2 , α_{bh} , β_{bh} , and ρ . The final row reports the pooled distribution across all covariate pairs.

Covariate pairs	min	1%	5%	25%	50%	75%	95%	99%	max
R_{bh}^2, α_{bh}	-0.73	-0.44	-0.32	-0.12	0.05	0.21	0.40	0.54	0.70
R_{bh}^2, β_{bh}	-0.82	-0.75	-0.68	-0.50	-0.23	0.12	0.50	0.74	0.81
R_{bh}^2, ρ	-0.92	-0.68	-0.40	-0.09	0.02	0.12	0.28	0.47	0.91
α_{bh}, β_{bh}	-0.81	-0.54	-0.39	-0.20	-0.06	0.09	0.30	0.45	0.72
α_{bh}, ρ	-0.86	-0.45	-0.22	-0.10	-0.04	0.02	0.12	0.21	0.72
β_{bh}, ρ	-0.52	-0.32	-0.19	-0.06	0.01	0.08	0.22	0.34	0.58
All pairs	-0.92	-0.69	-0.51	-0.15	-0.02	0.10	0.36	0.59	0.91

¹³As noted above, most trading rules generate discrete sell, neutral, or buy signals, although some moving-average rules generate fractional positions such as $+1/3$ or $-1/3$; these fractional signals are included in the aggregation.

¹⁴This equal-weighted $1/N$ aggregation is motivated by DeMiguel, Garlappi and Uppal (2007), who show that such portfolios are difficult to outperform using more complex weighting schemes. It also preserves the composition of selected rules while reducing transaction costs through the offsetting of opposing signals. Related aggregation approaches are used by Burnside, Eichenbaum and Rebelo (2011) and Filippou, Gozluklu and Taylor (2018).

¹⁵Portfolio formation is based on in-sample information and evaluation is conducted out of sample using a rolling-window procedure, thereby avoiding look-ahead bias.

Table 3 reports an empirical power comparison based on the average number of trading rules selected each month by the FDR^+ and $mfFDR^+$ procedures, both of which control the false discovery rate at 20%. In-sample performance is measured using the Sharpe ratios of technical trading rules computed from portfolio excess returns. Overall, $mfFDR^+$ selects substantially more outperforming rules than FDR^+ while maintaining the same FDR target. This gain in power is evident across currencies and sample periods and is particularly pronounced in the full sample. The table therefore provides a direct comparison between portfolios selected without covariates (FDR^+) and those selected using all four covariates jointly ($mfFDR^+$), highlighting the additional screening value of multivariate conditioning.

Table 3: Empirical power comparison. The table reports the average number of trading rules included in portfolios formed at the beginning of each month using the FDR^+ and $mfFDR^+$ procedures, each controlling the FDR at 20%. The rules are selected using as inputs the p -value from test (8), which tests whether the Sharpe ratio before transaction costs is equal to zero, together with the four covariates α_{bh} , β_{bh} , ρ , and R_{bh}^2 . The first five columns report results for five subperiods, and the final column reports averages over the full sample from the first portfolio-formation month through December 2020.

Country	Period											
	1973–1980		1981–1990		1991–2000		2001–2010		2011–2020		1973–2020	
	FDR^+	$mfFDR^+$	FDR^+	$mfFDR^+$	FDR^+	$mfFDR^+$	FDR^+	$mfFDR^+$	FDR^+	$mfFDR^+$	FDR^+	$mfFDR^+$
Australia	4138	8710	2451	8804	359	6615	1104	8095	114	6191	1530	7636
Canada	4603	8089	1349	6810	335	5826	370	6497	1099	5751	1433	6533
Germany/E.U.	2835	8244	3865	10832	496	7225	655	7707	864	6012	1697	7984
Japan	6121	11179	3453	9309	1173	8791	25	6211	724	6499	2145	8282
New Zealand	3227	10921	4483	10066	1141	8163	2232	8298	137	5047	2195	8388
Norway	2218	7348	3914	10194	127	6278	1243	6592	990	5615	1675	7192
Sweden	2485	9306	3050	10704	2142	8036	688	8146	1288	6079	1906	8410
Switzerland	3255	9296	1531	8960	1053	6513	259	6299	353	4684	1210	7056
U.K.	6303	10826	2889	9646	321	4828	22	6389	349	5413	1803	7280
Argentina					1674	2132	1565	3390	8213	10230	4011	5524
Colombia					3711	9102	3723	9918	1821	6589	3010	8463
India					1305	9112	3399	9262	387	5249	1716	7774
Indonesia			20	3936	988	5291	2801	9641	3084	8593	2090	7493
Israel	15551	3570	9365	5464	1882	8918	1542	7876	620	6223	3658	7021
Philippines			10529	13022	4671	9466	3644	8230	357	6962	3348	8509
Romania					11420	14408	1915	5954	896	6175	2312	6823
Russia					6787	11236	3806	9383	2597	8139	3836	9198
Slovak					345	5616	1643	8526	890	5813	1061	6818
Brazil					4185	5536	1350	10093	650	8451	1583	8580
Chile					5425	10570	3694	8572	210	6692	2635	8214
Czech					16	5919	2646	8221	1670	5788	1619	6726
Hungary					331	6843	473	6934	133	4711	310	6096
Korea					3098	9697	1899	8850	124	5346	1547	7761
Mexico			63	1388	628	5454	905	5326	147	6579	526	5518
Poland					2427	7506	1617	6950	359	5365	1298	6447
Singapore			1060	6313	2422	6740	1014	6655	1068	5400	1415	6268
South Africa			2045	9723	2817	8579	1661	8192	187	6635	1649	8199
Taiwan			8678	9708	4078	9220	2860	8490	116	6551	3278	8317
Thailand					1247	7120	2917	9792	2483	7851	2284	8330
Turkey					13393	11654	1477	8028	1935	9302	5334	9596

Figure 3 reports the out-of-sample (OOS) Sharpe ratios of the portfolios before transaction costs (Panel A) and after transaction costs (Panel B). For each currency, we form portfolios using FDR^+ , $fFDR^+$ with a single covariate (α_{bh} , β_{bh} , R_{bh}^2 , or ρ), $fFDR^+$ using the first principal component (PC1) of the four covariates, and $mfFDR^+$ using all four covariates jointly. Panel A shows that, before transaction costs, the alternative portfolio strategies perform broadly similarly. Once transaction costs are taken into account, however, the $mfFDR^+$ portfolios outperform the alternatives on average and deliver positive Sharpe ratios for almost all currencies, whereas the other procedures often do not. Detailed post-cost results are reported in Table 4.

Table 4 reports the OOS Sharpe ratios of portfolios formed using the FDR^+ , $fFDR^+$, and $mfFDR^+$ procedures, computed from excess returns net of transaction costs. For $fFDR^+$, we also report results based on the first principal component (PC1) of the four covariates as a univariate summary measure. On average, the Sharpe ratio of the FDR^+ portfolios is close to zero. The final row reports t -statistics for paired comparisons of mean Sharpe ratios across currencies, relative to the FDR^+ benchmark. All covariate-augmented procedures outperform FDR^+ at conventional significance levels, with $mfFDR^+$ using all four covariates delivering the strongest performance overall. The $mfFDR^+$ portfolios produce positive Sharpe ratios for all currencies except Hungary and Poland, and 16 currencies exhibit statistically significant post-cost performance. Moreover, the $mfFDR^+$ portfolios achieve higher average net Sharpe ratios than both $fFDR^+$ portfolios, including the PC1 specification, and the FDR^+ portfolios, confirming the gains from multivariate conditioning.

Figure 3: OOS performance of portfolios. The figure reports the OOS Sharpe ratios before (Panel A) and after (Panel B) transaction costs for portfolios formed using FDR^+ , $fFDR^+$, and $mfFDR^+$. For $fFDR^+$, we consider each of the four covariates α_{bh} , β_{bh} , R_{bh}^2 , and ρ , as well as the first principal component (PC1) of these four covariates. For $mfFDR^+$, all four covariates are used jointly. The input p -value is obtained from test (8), which tests whether the pre-transaction-cost Sharpe ratio is equal to zero.




Table 4: Net Sharpe ratios of portfolios. The table presents the annualized Sharpe ratio after transaction cost of seven portfolios based on the FDR^+ , the $fFDR^+$ and $mfFDR^+$ controlling FDR at 20%. The input p -value for the procedures is obtained by the test (8), which tests the Sharpe ratio before transaction cost against zero. For the $fFDR^+$, we first consider four covariates: α_{bh} , β_{bh} , R_{bh}^2 , ρ and the first principal component of the four mentioned covariates. For the $mfFDR^+$ we study $d = 4$ with all four covariates. The second last row is the average of the Sharpe ratio across currencies. The last row is the t -statistic of the test comparing the (paired) means of the portfolios $fFDR^+/mfFDR^+$ to the portfolios FDR^+ . The numbers in parentheses are the corresponding p -values. “*”, “**” and “***” respectively indicate statistical significance at levels of 10%, 5% and 1%.

Country	FDR^+	$fFDR^+$					$mfFDR^+$
		α_{bh}	β_{bh}	R_{bh}^2	ρ	PC1	
Australia	-0.02 (0.42)	-0.01 (0.97)	0.11 (0.42)	0.11 (0.42)	0.11 (0.42)	0.06 (0.63)	0.14 (0.29)
Canada	-0.25 (0.80)	-0.10 (0.50)	0.06 (0.69)	0.05 (0.80)	0.03 (0.89)	0.00 (0.98)	0.11 (0.46)
Germany/E.U.	0.13 (0.02)**	0.20 (0.17)	0.32 (0.02)**	0.31 (0.02)**	0.19 (0.19)	0.27 (0.07)*	0.46 (0.00)***
Japan	0.38 (0.06)*	0.31 (0.02)**	0.24 (0.11)	0.28 (0.06)*	0.11 (0.43)	0.25 (0.09)*	0.41 (0.00)***
New Zealand	0.00 (0.05)**	0.10 (0.43)	0.15 (0.24)	0.24 (0.05)**	0.19 (0.11)	0.24 (0.05)**	0.28 (0.02)**
Norway	-0.08 (0.14)	0.07 (0.64)	0.03 (0.88)	0.21 (0.14)	-0.04 (0.78)	-0.06 (0.64)	0.13 (0.36)
Sweden	0.17 (0.04)**	0.39 (0.01)***	0.28 (0.05)**	0.25 (0.04)**	0.24 (0.07)*	0.26 (0.05)**	0.36 (0.01)***
Switzerland	-0.12 (0.41)	0.16 (0.26)	0.10 (0.44)	0.11 (0.41)	0.01 (0.91)	0.10 (0.47)	0.22 (0.12)
U.K.	0.09 (0.30)	0.02 (0.87)	0.18 (0.22)	0.15 (0.30)	0.15 (0.29)	0.06 (0.64)	0.25 (0.08)*
Argentina	0.27 (0.03)**	0.36 (0.04)**	0.38 (0.02)**	0.37 (0.03)**	0.43 (0.01)***	0.35 (0.04)**	0.31 (0.11)
Colombia	0.71 (0.01)***	0.44 (0.04)**	0.40 (0.07)*	0.51 (0.01)***	0.59 (0.00)***	0.35 (0.10)*	0.53 (0.01)***
India	0.25 (0.12)	0.18 (0.29)	0.13 (0.45)	0.25 (0.12)	0.24 (0.15)	0.30 (0.07)*	0.30 (0.07)*
Indonesia	-0.09 (0.29)	0.23 (0.12)	0.12 (0.40)	0.15 (0.29)	0.11 (0.47)	0.19 (0.18)	0.24 (0.09)*
Israel	0.73 (0.00)***	0.75 (0.00)***	0.68 (0.00)***	0.63 (0.00)***	0.68 (0.00)***	0.65 (0.00)***	0.39 (0.01)***
Philippines	-0.42 (0.05)**	0.36 (0.05)**	0.40 (0.03)**	0.35 (0.05)**	0.39 (0.03)**	0.35 (0.07)*	0.46 (0.01)***
Romania	-0.31 (0.81)	-0.13 (0.53)	0.03 (0.86)	0.04 (0.81)	-0.13 (0.58)	-0.11 (0.60)	0.12 (0.51)
Russia	0.39 (0.07)*	0.38 (0.07)*	0.37 (0.07)*	0.40 (0.07)*	0.40 (0.07)*	0.40 (0.05)**	0.48 (0.01)***
Slovak	-0.12 (0.62)	-0.09 (0.68)	0.17 (0.38)	0.10 (0.62)	0.05 (0.76)	0.08 (0.63)	0.13 (0.48)
Brazil	-0.05 (0.71)	0.20 (0.29)	0.00 (0.98)	0.07 (0.71)	0.02 (0.94)	0.06 (0.78)	0.33 (0.10)*
Chile	0.22 (0.07)*	0.52 (0.01)***	0.35 (0.07)*	0.36 (0.07)*	0.47 (0.03)**	0.43 (0.05)**	0.38 (0.06)*
Czech	-0.11 (0.17)	-0.12 (0.55)	0.01 (0.89)	0.23 (0.17)	-0.14 (0.42)	-0.08 (0.70)	0.06 (0.65)
Hungary	-0.23 (0.97)	-0.09 (0.67)	-0.06 (0.75)	0.00 (0.97)	-0.08 (0.67)	-0.11 (0.55)	-0.17 (0.41)
Korea	0.15 (0.41)	0.18 (0.33)	0.15 (0.39)	0.14 (0.41)	0.09 (0.62)	0.14 (0.42)	0.23 (0.25)
Mexico	0.04 (0.97)	0.10 (0.49)	-0.09 (0.58)	0.00 (0.97)	0.02 (0.86)	0.03 (0.85)	0.04 (0.73)
Poland	-0.23 (0.89)	-0.01 (0.98)	0.03 (0.87)	-0.03 (0.89)	0.06 (0.77)	0.09 (0.64)	-0.03 (0.89)
Singapore	-0.38 (0.67)	-0.19 (0.25)	-0.02 (0.93)	0.08 (0.67)	-0.03 (0.87)	-0.20 (0.21)	0.15 (0.32)
South Africa	-0.27 (0.99)	0.01 (0.93)	-0.02 (0.93)	0.00 (0.99)	0.08 (0.62)	0.10 (0.54)	0.10 (0.49)
Taiwan	0.44 (0.00)***	0.56 (0.00)***	0.45 (0.00)***	0.59 (0.00)***	0.42 (0.01)***	0.57 (0.00)***	0.66 (0.00)***
Thailand	0.11 (0.30)	0.37 (0.05)**	0.18 (0.35)	0.19 (0.30)	0.27 (0.13)	0.07 (0.75)	0.38 (0.04)**
Turkey	0.36 (0.00)***	0.54 (0.00)***	0.53 (0.00)***	0.52 (0.00)***	0.56 (0.00)***	0.54 (0.00)***	0.45 (0.00)***
Average	0.06	0.19	0.19	0.22	0.18	0.18	0.26
t-stats		3.9	3.5	4.8	3.6	3.4	5.1

Overall, these findings have three implications. First, the covariates are clearly informative: they help identify more profitable rules that remain effective after trading costs. Second, conditioning on multiple covariates further improves the OOS performance of technical-trading portfolios, supporting the use of $mfFDR^+$ in this setting. Third, the fact that $mfFDR^+$ outperforms $fFDR^+$ even when the latter uses PC1 suggests that reducing the covariate vector to a single index discards economically relevant information, including possible nonlinearities and interactions across covariates. Taken together, the results underscore the flexibility and empirical value of the multivariate functional FDR approach.

5.3. *Baskets of Currencies*

In the previous section, we selected outperforming trading rules separately for each currency. Because the profitability of technical trading rules varies across both currencies and time, an investor may benefit from evaluating rule performance jointly across currencies rather than in isolation. This cross-currency approach permits diversification and allows capital to be reallocated toward the currencies and rules that appear most promising at each point in time.

We therefore consider an investor who can trade across the full currency universe and select outperforming rule–currency combinations from the entire pool of technical trading opportunities while controlling for false discoveries. Depending on data availability, this pooled universe ranges from 190,755 combinations ($9 \times 21,195$) to 635,850 combinations ($30 \times 21,195$). We construct rolling monthly portfolios exactly as in the previous section, but now apply the procedure to this enlarged cross-currency rule set. The four covariates used in $mfFDR^+$ are ρ , α_{mk} , β_{mk} , and R_{mk}^2 , as described in Section 5.1. Table 5 reports the pairwise

correlations of these covariates. The dependence is generally moderate: approximately 98% of the correlation coefficients lie in the interval $[-0.73, 0.52]$.

Table 5: Summary of correlation coefficients of covariate pairs: pooled currencies. The table reports quantiles of the pairwise correlation coefficients for the six pairs formed from the four covariates R_{mk}^2 , α_{mk} , β_{mk} , and ρ . The final row reports the pooled distribution across all covariate pairs.

Covariate pairs	min	1%	5%	25%	50%	75%	95%	99%	max
R_{mk}^2, α_{mk}	-0.44	-0.38	-0.27	-0.08	0.02	0.10	0.37	0.67	0.71
R_{mk}^2, β_{mk}	-0.81	-0.76	-0.70	-0.52	-0.23	0.12	0.50	0.54	0.55
R_{mk}^2, ρ	-0.26	-0.24	-0.18	-0.08	-0.02	0.07	0.25	0.30	0.36
α_{mk}, β_{mk}	-0.93	-0.90	-0.68	-0.31	-0.16	-0.03	0.31	0.53	0.63
α_{mk}, ρ	-0.20	-0.15	-0.11	-0.03	0.04	0.12	0.26	0.33	0.36
β_{mk}, ρ	-0.40	-0.36	-0.27	-0.11	-0.03	0.05	0.15	0.20	0.21
All pairs	-0.93	-0.73	-0.55	-0.16	-0.03	0.07	0.32	0.52	0.71

At the beginning of each month, the trading rules selected by $mfFDR^+$, controlling the FDR at 20%, are grouped by currency. Suppose, for example, that the procedure selects rules only from currencies A and B , with k_A and k_B selected rules, respectively. Total capital is then allocated equally across the resulting $k = k_A + k_B$ rule-based positions. The portfolio is held for one month, after which performance is evaluated and the procedure is repeated using the next rolling window.

Table 6 reports the performance of these $mfFDR^+$ portfolios in terms of annualized Sharpe ratios and mean returns, both before and after transaction costs. The rightmost column reports the one-way break-even cost, defined as the fixed transaction cost that would reduce the net Sharpe ratio to zero. Over the full sample period, 1973–2020, the results are economically strong: the gross Sharpe ratio is 1.06 and the net Sharpe ratio is 0.95, corresponding to a break-even cost of 60 basis points. These results compare favorably with the single-currency portfolios reported in Section 5.2 and indicate that pooling rule selection across currencies yields substantial diversification benefits.

The subperiod results reveal a pronounced evolution in the profitability of technical trading. In the 1970s, technical rules already delivered strong performance, with a net Sharpe ratio of 1.35. Profitability then rose sharply in the 1980s, when the net Sharpe ratio reached 1.97. The 1990s remained highly profitable, with a net Sharpe ratio of 0.81, although lower than in the preceding decade. A clear decline emerged in the 2000s, when the net Sharpe ratio fell to 0.53, though performance remained economically meaningful after transaction costs. In the most recent decade, 2011–2020, the net Sharpe ratio declined further to 0.21, but profitability remained positive, with a break-even cost of 14 basis points.

Overall, these results point to substantial time variation in the profitability of technical trading. Although increasing market efficiency appears to have eroded returns over time, the systematic application of $mfFDR^+$ continues to identify economically meaningful opportunities even in the most recent decade.

Table 6: Performance of $mfFDR^+$ portfolios. The table reports annualized Sharpe ratios (SR) and mean returns, before and after transaction costs, for $mfFDR^+$ portfolios constructed from the full cross-currency universe. The FDR is controlled at 20%. The final column reports the implied one-way break-even cost.

Period	Excess SR	Net SR	Excess Return (%)	Net Return (%)	Break-even (bps)
Whole Period	1.06	0.95	3.80	3.40	60
1973–1980	1.45	1.35	4.47	4.18	69
1981–1990	2.08	1.97	7.30	6.93	128
1991–2000	0.92	0.81	4.15	3.64	72
2001–2010	0.69	0.53	2.37	1.82	34
2011–2020	0.29	0.21	0.88	0.63	14

To better understand the sources of this time variation, we conduct two complementary analyses. First, we examine the average proportion of rules selected as profitable in each period. Second, we study selection rates by strategy category (Channel breakout, Filter rule,

Moving average, RSI, and Support-resistance).¹⁶

For the first analysis, we compute the average ratio of selected rules to total available rules in each month. Because the universe of candidate rules expands over time from 190,755 to 635,850, this ratio provides a comparable measure of the prevalence of profitable opportunities across subperiods. The results, reported in Table 7, show a clear downward trend. The average selection rate falls from 36% in the 1970s and 1980s to 23% in the 1990s and 2000s, and then to 18% in the 2010s. This pattern closely mirrors the decline in Sharpe ratios reported in Table 6, reinforcing the interpretation that FX markets have gradually become more efficient in eliminating exploitable technical patterns.

Table 7: Average selection rates by rule category under $mfFDR^+$ (FDR controlled at 20%). The table reports the proportion of rules selected each month (in %) by category and subperiod.

Period	Selected strategy rate (%)					
	All	Channel breakout	Filter rule	Moving average	RSI	Support-resistance
Whole Period	27	8	21	34	13	19
1973–1980	36	11	30	45	19	26
1981–1990	36	10	27	46	15	26
1991–2000	23	7	18	29	12	16
2001–2010	23	7	19	29	11	16
2011–2020	18	6	13	23	11	13

Table 7 also shows that, although overall profitability declines over time, the relative ranking of strategy types remains remarkably stable. Moving average rules dominate throughout the sample, with an average selection rate of 34%, followed by filter rules at 21% and support-resistance rules at 19%. Channel breakout and RSI rules consistently lag behind. These patterns suggest that some classes of technical strategy are structurally more robust than others, even as the overall environment becomes less favorable to technical trading.

¹⁶These two analyses reflect the in-sample prevalence of profitable rules based on one-year rolling windows; they do not directly measure individual out-of-sample performance. A third analysis, discussed below, examines the unconditional profitability of all rules.

Taken together, Table 7 yields two main conclusions. First, technical-trading profitability displays substantial time variation, consistent with a gradual increase in market efficiency, although economically meaningful opportunities remain even in the most recent decade. Second, significant cross-sectional variation persists across strategy types, with moving average rules remaining particularly effective, whereas RSI rules appear consistently weak.

Section II of the Internet Appendix further examines performance conditional on rule category. We find that portfolios based on filter rules deliver the highest pre-cost returns, but their higher turnover makes them more sensitive to transaction costs. By contrast, moving average strategies retain strong net profitability and remain especially attractive for practitioners. RSI rules underperform in all periods and therefore appear to be natural candidates for exclusion from mixed-strategy portfolios.

5.4. Evolution of Profitability and Fund Redistribution of $mfFDR^+$

It is well known that the profitability of technical trading tends to decline as foreign exchange markets—especially those for developed currencies—become more efficient. Prior studies document that the profitability of technical trading rules in developed currencies has diminished in recent decades (Qi and Wu, 2006; Menkhoff and Taylor, 2007; Neely, Weller and Ulrich, 2009), although it continues to persist in a number of emerging-market currencies (Hsu, Taylor and Wang, 2016).

Our results in Section 5.3 are consistent with this pattern. Technical trading is highly profitable in the earlier part of the sample, remains strong through the 1980s and 1990s, declines in the 2000s while still delivering economically meaningful returns, and weakens further in the 2010s. The falling selection rates reported in Table 7 likewise indicate that

profitable technical rules become less prevalent over time. This raises a natural question: can the $mfFDR^+$ procedure adapt to these changing conditions by reallocating capital across currencies as profit opportunities shift?

To address this question, we conduct two complementary analyses. First, we study the out-of-sample (OOS) profitability of the rules selected in each subperiod by partitioning the selected basket-level rules into their constituent currencies and evaluating the resulting OOS Sharpe ratios at the currency level. These gross and net Sharpe ratios are reported in Table 8. Second, we examine how the distribution of selected profitable rules evolves across currencies over time, as reported in Table 9.

Because of data availability, the set of tradable currencies expands over time. In the earliest subperiod, 1973–1980, only developed currencies are included. From the 1980s onward, additional developed currencies, such as the Israeli shekel, and emerging-market currencies, such as Mexico, South Africa, Taiwan, and Turkey, enter the sample as reliable data become available. By the 2000s, the full panel of 30 currencies is covered. As a result, each subperiod reflects not only changing market conditions but also a gradually widening cross-sectional opportunity set. When a currency enters only partway through a subperiod, we exclude that first partial subperiod from the summary statistics in order to preserve comparability over time.

Table 8 shows that technical profitability varies sharply across currencies and over time. In the 1970s and 1980s, most developed currencies deliver high net Sharpe ratios, often above 1. In the 1990s, profitability remains strong, with several emerging-market currencies—including Turkey, Taiwan, and the Philippines—joining the set of top performers. In 2001–2010, profitability weakens in many developed markets but remains robust in selected emerging-

market currencies, such as Russia and Brazil. In the most recent decade, 2011–2020, profitability deteriorates sharply in most developed markets, yet pockets of opportunity remain in a number of emerging markets, notably Argentina, Indonesia, Russia, Turkey, and Taiwan.

Table 8: OOS performance of trading rules selected by $mfFDR^+$ and partitioned into currencies (controlling the FDR at 20%) across subperiods. The table reports out-of-sample Sharpe ratios, before and after transaction costs, for the trading rules selected by applying $mfFDR^+$ to the basket of currencies and then partitioning the selected rules by currency. Results are reported for five subperiods: 1973–1980 (P1), 1981–1990 (P2), 1991–2000 (P3), 2001–2010 (P4), and 2011–2020 (P5). For example, in 2011–2020, the outperforming rules selected each month by $mfFDR^+$ from a universe of up to 635,850 rule–currency combinations are partitioned by currency, and the Sharpe ratios of the resulting gross and net return series are then computed for each currency and subperiod.

Country	Gross Sharpe Ratio					Net Sharpe Ratio				
	P1	P2	P3	P4	P5	P1	P2	P3	P4	P5
Australia	0.17	0.44	0.07	0.30	-0.26	0.12	0.38	-0.01	0.28	-0.29
Canada	0.83	0.49	0.12	0.10	-0.11	0.62	0.36	0.01	0.07	-0.15
Germany/E.U.	0.51	1.14	0.16	0.41	-0.17	0.43	1.12	0.13	0.38	-0.19
Japan	1.28	0.87	0.27	-0.13	0.16	1.23	0.81	0.23	-0.15	0.14
New Zealand	0.65	0.65	0.54	0.39	-0.65	0.61	0.54	0.44	0.36	-0.69
Norway	0.36	0.86	-0.14	0.24	-0.27	0.27	0.79	-0.22	0.21	-0.30
Sweden	0.52	1.05	0.33	0.31	-0.28	0.45	1.00	0.25	0.29	-0.31
Switzerland	0.67	0.74	0.15	-0.00	-0.33	0.63	0.70	0.11	-0.03	-0.38
U.K.	1.09	0.88	-0.50	0.00	-0.26	1.02	0.84	-0.57	-0.02	-0.28
Argentina			1.05	0.04	1.09			0.50	-0.01	1.05
Colombia			0.81	0.74	0.30			0.76	0.66	0.20
India			0.47	0.75	-0.19			0.44	0.68	-0.23
Indonesia		-0.01	0.38	0.54	0.67		-0.04	0.30	0.13	0.56
Israel	0.84	2.03	0.61	0.30	-0.10	0.83	1.92	0.27	0.10	-0.15
Philippines		0.78	0.80	0.59	0.09		0.72	0.54	0.41	-0.01
Romania			0.54	0.15	-0.17			0.52	-0.02	-0.24
Russia			0.37	0.81	0.44			0.37	0.76	0.39
Slovak			-0.23	0.61	-0.23			-0.31	0.56	-0.25
Brazil			-0.49	0.63	0.19			-0.52	0.60	0.17
Chile			1.07	0.66	0.00			0.90	0.57	-0.07
Czech			-0.06	0.35	-0.28			-0.11	0.30	-0.37
Hungary			-0.05	0.20	-0.41			-0.09	0.14	-0.47
Korea			0.45	0.33	-0.32			0.45	0.23	-0.52
Mexico		-0.07	0.26	0.04	-0.03		-0.17	0.17	0.01	-0.04
Poland			0.18	0.28	-0.53			0.14	0.22	-0.57
Singapore		0.32	0.42	0.38	-0.04		0.20	0.11	0.19	-0.09
South Africa		0.40	0.36	0.35	-0.13		0.29	0.19	0.17	-0.18
Taiwan		1.91	0.86	0.69	0.34		1.88	0.79	0.52	0.29
Thailand			0.43	0.75	0.17			0.33	0.65	0.08
Turkey			1.01	0.40	0.36			0.99	0.27	0.33

Table 9: Distribution across currencies of trading rules selected by $mfFDR^+$ (controlling the FDR at 20%) across subperiods. The table reports the average proportion (%) of trading rules selected by $mfFDR^+$ across currencies in five subperiods: 1973–1980 (P1), 1981–1990 (P2), 1991–2000 (P3), 2001–2010 (P4), and 2011–2020 (P5). For example, in 2011–2020, an average of 113,160 rules are selected per month (bottom row of the final column), and 3.1% of those selected rules are allocated to AUD (top row of the final column). Bold values indicate the largest proportion in each subperiod.

Country	Period				
	P1	P2	P3	P4	P5
Australia	9.6	6.7	3.0	3.5	3.1
Canada	9.3	4.2	2.5	2.4	2.3
Germany/E.U.	9.0	9.3	3.9	3.3	3.0
Japan	14.8	7.9	4.8	2.2	3.2
New Zealand	12.1	8.4	3.7	3.9	2.3
Norway	7.7	8.1	3.5	2.9	2.8
Sweden	9.8	9.1	4.4	3.5	3.1
Switzerland	11.4	7.9	4.0	2.6	2.3
U.K.	12.9	7.7	2.7	2.4	2.5
Argentina			1.1	2.8	10.1
Colombia			2.8	4.8	3.2
India			3.7	3.4	2.6
Indonesia		1.7	2.8	4.1	3.7
Israel	3.2	10.9	4.4	2.7	2.3
Philippines		2.9	6.1	3.2	3.1
Romania			1.5	3.4	3.2
Russia			2.9	4.0	4.5
Slovak			1.1	4.0	2.9
Brazil			1.2	4.9	5.0
Chile			2.2	3.5	3.0
Czech			1.7	3.9	3.2
Hungary			2.5	3.4	2.6
Korea			3.1	3.7	2.2
Mexico		0.2	3.0	1.9	3.3
Poland			1.6	3.2	2.9
Singapore		2.9	3.2	2.1	2.2
South Africa		6.2	4.3	4.1	4.1
Taiwan		5.9	4.7	2.6	2.5
Thailand			1.9	3.8	3.2
Turkey			11.6	3.8	5.4
Average number of selected rules (per month)	69,028	93,317	122,476	146,532	113,160

Table 9 shows that $mfFDR^+$ responds to these shifts by reallocating selected rules across currencies. In the earlier decades, the largest shares are concentrated in highly profitable developed currencies such as the Japanese yen, the Deutsche mark (and later the euro),

and the Israeli shekel. From the 2000s onward, the distribution shifts increasingly toward emerging markets, which account for a growing share of selected rules in the final subperiod.

The same pattern appears when we focus on profitable destinations of capital. The fraction of currencies delivering positive net Sharpe ratios declines from 100% in P1 to only 30% in P5. Even so, $mfFDR^+$ continues to direct the bulk of selected rules toward those currencies in which profitability survives: 100%, 98%, 85%, 87%, and 41% of selected rules are allocated to profitable currencies in periods P1 through P5, respectively.

Taken together, these results highlight both the gradual erosion of technical-rule profitability as FX markets mature and the ability of $mfFDR^+$ to adapt to changing opportunity sets. Dynamic redistribution across currencies plays an important role in preserving overall portfolio profitability despite the broad decline in technical predictability among developed currencies. More broadly, the success of the $mfFDR^+$ portfolio reflects two complementary forces: first, the continuing profitability of technical signals in a subset of currencies and periods; and second, the ability of $mfFDR^+$ to identify and exploit those opportunities through systematic reallocation.

6. Market Environment and the Profitability of Technical Trading

The preceding results show that the profitability of technical trading varies markedly across periods and currencies. We now examine whether that variation is systematically related to the broader market environment. Specifically, we consider three sets of potential drivers: computational power, financial openness, and a wider set of macroeconomic and market variables. These tests help clarify whether the decline and cross-sectional concentration of technical-trading profits are associated with changes in market efficiency, competition,

volatility, liquidity, and financial development.

6.1. Computational power

Throughout this study, technical trading rules are implemented using daily data and executed at the daily frequency. Over time, however, foreign exchange trading has increasingly shifted toward intraday and high-frequency horizons, facilitated by rapid advances in computational technology and machine learning (Cavalcante, Brasileiro, Souza, Nobrega and Oliveira, 2016). It is therefore natural to conjecture that improvements in computational power have increased market efficiency and, in turn, reduced the profitability of slower-moving technical trading rules.

Because computational power is difficult to observe directly at the country level, we proxy for technological innovation using the volume of related academic publications. Specifically, we collect country-level data on the number of publications whose titles, abstracts, or keywords contain at least one of the following phrases: “algorithmic trading,” “artificial intelligence in trading,” “machine learning in trading,” “neural networks in trading,” or “high-frequency trading.”¹⁷ We scale this count by population to obtain a measure of publications per million inhabitants, denoted *Publication*. Although indirect, this proxy captures cross-country variation in research activity and innovation related to advanced trading technologies.

¹⁷Publication data are sourced from Scopus as of January 2024. Chaboud, Chiquoine, Hjalmarsson and Vega (2014) examine the impact of algorithmic trading in the foreign exchange market.

We then estimate the following annual fixed-effects panel regression:

$$R_{i,t} = a_i + \beta_1 Publication_{i,t} + \sum_{k=1}^4 \beta_{k+1} Z_{k,i,t} + \epsilon_{i,t}, \quad (11)$$

where $R_{i,t}$ is the cumulative annual return of the *mfFDR*-based portfolio for currency i in year t , a_i denotes country fixed effects, $Z_{k,i,t}$ is a vector of macroeconomic controls (GDP growth, unemployment rate, stock market development, and international trade ratio), and $\epsilon_{i,t}$ is the error term. Macroeconomic variables are obtained from the World Bank database.¹⁸ The results are reported in the first column of each panel of Table 10. The regressions have modest but nontrivial explanatory power both before and after transaction costs. In both cases, the coefficient on *Publication* is negative and statistically significant, consistent with the view that technological innovation has increased market efficiency and thereby reduced the profitability of technical trading rules over time.

6.2. Financial openness

In addition to *Publication*, we examine whether financial openness affects the profitability of technical trading rules. More open capital markets facilitate cross-border arbitrage, accelerate information diffusion, and intensify competition from global investors, all of which may reduce the scope for profitable technical trading. We measure openness using the Chinn-Ito (2006) index, denoted *Openness*.

¹⁸World Bank Data, January 2024.

Using the same panel-regression framework, we estimate

$$R_{i,t} = a_i + \beta_1 Openness_{i,t} + \sum_{k=1}^4 \beta_{k+1} Z_{k,i,t} + \epsilon_{i,t}. \quad (12)$$

The results are reported in the second column of each panel of Table 10. The regressions again exhibit modest but nontrivial explanatory power for both gross and net returns. The coefficient on *Openness* is consistently negative and statistically significant, indicating that technical trading tends to be less profitable in markets with more open capital accounts. This pattern is consistent with the view that greater financial openness is associated with faster arbitrage and more efficient incorporation of information into exchange rates.

Finally, we estimate a combined specification that includes both *Publication* and *Openness*,

Table 10: Regressions of *mfFDR*-based portfolio returns on *Publication*, *Openness*, and macro factors. The regressions are estimated with country fixed effects. The dependent variable is the annual return of the *mfFDR*-based portfolios, partitioned by currency as in Table 8. The main explanatory variables are *Publication* (publications per million inhabitants relating to advanced trading technologies), the Chinn-Ito (2006) index of financial openness (*Openness*), and macroeconomic controls. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses.

	<i>Dependent variable</i>					
	Excess return			Net return		
Publication	−0.007** (0.003)		−0.007** (0.003)	−0.007** (0.003)		−0.006** (0.003)
Openness		−0.007** (0.003)	−0.006* (0.003)		−0.007** (0.003)	−0.006* (0.003)
GDP growth	−0.002 (0.001)	−0.002 (0.001)	−0.002 (0.001)	−0.002 (0.001)	−0.002 (0.001)	−0.002 (0.001)
Stock market development	−0.0005*** (0.0001)	−0.0005*** (0.0001)	−0.0005*** (0.0001)	−0.0004*** (0.0001)	−0.0005*** (0.0001)	−0.0005*** (0.0001)
International trade	0.00003 (0.0002)	0.00005 (0.0002)	0.00004 (0.0002)	−0.00001 (0.0002)	0.00001 (0.0002)	0.000002 (0.0002)
Unemployment rate	−0.002 (0.002)	−0.001 (0.002)	−0.002 (0.002)	−0.002 (0.002)	−0.002 (0.001)	−0.002 (0.002)
Observations	790	772	772	790	772	772
Adjusted R^2	0.028	0.027	0.032	0.021	0.020	0.024

together with the full set of macroeconomic controls. The results are presented in the final column of each panel of Table 10. Both variables remain negative and statistically significant, at the 5% level for *Publication* and the 10% level for *Openness*. These findings suggest that advances in trading technology and greater financial openness each independently contribute to the erosion of technical-rule profitability over time. This interpretation is consistent with the adaptive markets hypothesis of Lo (2004), under which market efficiency evolves endogenously as traders learn and adapt.

6.3. Other Macro Factors

We next examine how the profitability of technical trading rules varies with broader macroeconomic and market conditions commonly linked to currency risk premia. Following Filippou, Rapach, Taylor and Zhou (2023), we consider a set of theoretically motivated variables widely used in the FX literature. Specifically, we include local-minus-U.S. differentials in three interest rates (the 3-month T-bill, the 5-year government note, and the 10-year government bond), inflation, the unemployment gap, the dividend-to-price ratio, the price-to-earnings ratio, and stock market capitalization. We also include three market-microstructure variables: idiosyncratic volatility, aggregate volatility, and illiquidity.¹⁹

Because the macroeconomic data are available at the monthly frequency, we aggregate the daily returns of the *mfFDR*-based portfolios to monthly returns. We then estimate the

¹⁹Following Filippou, Gozluklu and Taylor (2018), volatility (VOL) and illiquidity (ILL) are defined as $VOL_{i,t} = (1/T_t) \sum_{d \in T_t} |\Delta s_{i,d}|$ and $ILL_{i,t} = (1/T_t) \sum_{d \in T_t} BAS_{i,d}$, where $|\Delta s_{i,d}|$ is the absolute daily change in the log spot exchange rate of currency i , $BAS_{i,d}$ is the bid-ask spread (in percent) of currency i on day d , and T_t is the number of trading days in month t . Higher values indicate greater volatility or illiquidity.

following panel regression with currency fixed effects:

$$R_{i,t} = a_i + \sum_{k=1}^{11} \beta_k X_{k,i,t} + \epsilon_{i,t}, \quad (13)$$

where $R_{i,t}$ is the monthly return of the *mfFDR*-based portfolio for currency i in month t , a_i is a currency fixed effect, and $X_{k,i,t}$ is the vector of 11 explanatory variables. Data availability limits this analysis to 23 currencies, excluding Argentina, Colombia, Indonesia, Romania, Slovakia, Chile, and Turkey. We estimate separate regressions for gross and net returns. The results are reported in Table 11.

Several findings stand out. Stock market capitalization and aggregate volatility are highly significant at the 1% level in both specifications. The positive association between aggregate volatility and portfolio profitability is consistent with the stylized fact documented by [Menkhoff and Taylor \(2007\)](#) that technical trading performs better in more volatile markets. This interpretation is reinforced by the positive and highly significant coefficient on illiquidity, suggesting that technical rules are more profitable in less liquid trading environments.

By contrast, the negative and highly significant coefficient on stock market capitalization, together with the negative effect of the price-to-earnings ratio, suggests that technical trading tends to be more profitable in less developed and less informationally efficient markets.²⁰

In Sections [II](#) and [IM](#) of the Internet Appendix, we also repeat these analyses using alternative performance measures, namely mean return and three-factor alpha based on pre-cost returns. The conclusions are robust. First, the performance of the *mfFDR*-based

²⁰In Section [IJ](#) of the Internet Appendix, we repeat this analysis using predictive regressions. We find that the 3-month T-bill, the 5-year note, inflation, idiosyncratic volatility, and illiquidity are important predictors of which rules are selected as profitable, although they explain only a modest share of the variation in returns.

Table 11: Regression of $mfFDR$ -based portfolio returns on macro variables. The table reports panel regressions with currency fixed effects. The dependent variable is the monthly return, gross or net of transaction costs, of the $mfFDR$ -based portfolio partitioned by currency. Explanatory variables include interest rates, inflation, valuation ratios, volatility, and illiquidity, as defined in the text. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses.

	Excess return	Net return
3-month T-Bill	0.039 (0.025)	0.040 (0.025)
5-year note	-0.027 (0.056)	-0.025 (0.056)
10-year bond	0.002 (0.038)	-0.004 (0.039)
Inflation	0.061 (0.049)	0.062 (0.049)
Dividend-to-price ratio	-0.045* (0.027)	-0.042 (0.027)
Price-to-earnings ratio	-0.00002** (0.00001)	-0.00002** (0.00001)
Stock market capitalization	-0.021*** (0.005)	-0.021*** (0.005)
Unemployment gap	0.008 (0.013)	0.008 (0.013)
Idiosyncratic volatility	-0.105 (0.067)	-0.118* (0.068)
Illiquidity	0.029*** (0.009)	0.027*** (0.010)
Aggregate volatility	0.010*** (0.003)	0.010*** (0.003)
Observations	6,609	6,609
Adjusted R^2	0.030	0.029

portfolios is qualitatively consistent across metrics. Second, performance is even stronger when measured using three-factor alpha, with *Publication* emerging as the most significant explanatory variable.

Taken together, these findings reinforce the central message of this section: the profitability of technical trading rules is systematically related to the surrounding market environment. Technical trading performs better in more volatile, less liquid, and less financially developed markets, and worse in more open and technologically advanced settings. This pattern is

consistent with adaptive learning and a gradual increase in market efficiency over time, and it helps explain why pockets of technical profitability persist in some currencies and periods even as the broader environment becomes less favorable.

7. Conclusion

This paper develops the *multivariate functional false discovery rate* (*mfFDR*) method, a flexible approach to large-scale multiple testing in which rejection thresholds depend on multiple informative covariates. Relative to existing procedures, *mfFDR* delivers materially greater power while maintaining false-discovery control, and it provides a natural way to incorporate auxiliary information into the evaluation of competing signals. The method is straightforward to implement, robust to moderate covariate noise and dependence, and broadly applicable in empirical settings characterized by many simultaneous tests. Together with recent contributions such as [Harvey and Liu \(2020\)](#) and [Giglio, Liao and Xiu \(2021\)](#), our paper adds to a growing literature in finance that addresses data-snooping bias in large-scale multiple-hypothesis testing.

Our application to foreign exchange markets illustrates the empirical value of the approach. In single-currency tests, *mfFDR*-based portfolios outperform traditional alternatives and deliver stronger, more robust results. In a global application spanning 30 currencies and more than 21,000 trading rules, the method identifies portfolios with Sharpe ratios close to one over a 50-year horizon, with economically meaningful out-of-sample profitability persisting even in recent decades. These findings show that multivariate conditioning is not merely a statistical refinement: it has tangible economic value in the selection of trading strategies.

The paper also sheds light on the broader evolution of technical trading profitability.

We show that profitability declines with advances in computational power, greater capital-account openness, higher levels of financial development, and deeper FX market liquidity. At the same time, the surviving profits become increasingly concentrated in particular currencies and market environments, making adaptive selection especially important. In this sense, *mfFDR* is useful not only for controlling false discoveries, but also for tracing how changing market conditions shape the opportunity set available to traders.

More broadly, *mfFDR* provides a useful framework for empirical research in finance and economics wherever researchers confront large numbers of competing signals, models, or strategies. The central message of the paper is therefore twofold: multivariate covariate information materially improves large-scale inference, and careful control of false discoveries can also generate new insight into the economic structure of predictability. For that reason, we expect the *mfFDR* approach to be useful well beyond foreign exchange markets, in other settings where large-scale testing and adaptive behavior are central empirical challenges.

Appendix A. The multivariate functional false discovery rate

To complement Section 2.1, we present the full procedure for estimating the key components of the *mfFDR* method: the local null proportion $\pi_0(\mathbf{z})$, the joint density $f(p, \mathbf{z})$, and the q -value function used to control the false discovery rate.

Estimating $\pi_0(\mathbf{z})$. We follow and extend the “bin” approach of [Hsu, Kyriakou, Ma and Sermpinis \(2025\)](#), in which tests are partitioned into groups based on covariates and $\pi_0(\mathbf{z})$ is estimated within each group. In addition, we use a kernel-based density estimation procedure to obtain a smoother estimate.

For $\lambda \in [0, 1)$, define:

$$\pi_0(\mathbf{z}, \lambda) = \frac{\mathbb{P}(P > \lambda \mid \mathbf{Z} = \mathbf{z})}{1 - \lambda}. \quad (\text{A.1})$$

Since $\mathbb{P}(P > \lambda \mid \mathbf{Z} = \mathbf{z}) \geq (1 - \lambda)\pi_0(\mathbf{z})$, this provides a conservative estimate: $\pi_0(\mathbf{z}, \lambda) \geq \pi_0(\mathbf{z})$.

We rewrite $\pi_0(\mathbf{z}, \lambda)$ as:

$$\pi_0(\mathbf{z}, \lambda) = \mathbb{P}(\mathbf{Z} = \mathbf{z} \mid P > \lambda) \cdot \frac{\mathbb{P}(P > \lambda)}{1 - \lambda}. \quad (\text{A.2})$$

The second term is estimated by:

$$\hat{\pi}_0(\lambda) = \frac{\#\{p_i > \lambda\}}{n(1 - \lambda)},$$

where $\#\{p_i > \lambda\}$ denotes the number of p -values exceeding λ ([Storey, 2002](#)). The conditional

density $\mathbb{P}(\mathbf{Z} = \mathbf{z} \mid P > \lambda)$ is estimated by $\hat{h}_\lambda(\mathbf{z})$, yielding:

$$\hat{\pi}_0(\mathbf{z}, \lambda) = \hat{h}_\lambda(\mathbf{z}) \cdot \hat{\pi}_0(\lambda). \quad (\text{A.3})$$

Estimating $f(p, \mathbf{z})$ and $\hat{h}_\lambda(\mathbf{z})$. We estimate both $\hat{h}_\lambda(\mathbf{z})$ and $f(p, \mathbf{z})$ using local likelihood kernel density estimation (KDE), following the approach of Loader (1999) and adopting the probit transformation of Geenens (2014).

Specifically, we transform:

$$\tilde{p}_i = \Phi^{-1}(p_i), \quad \tilde{z}_i^k = \Phi^{-1}(z_i^k), \quad k = 1, \dots, d,$$

where Φ^{-1} is the inverse standard normal CDF. The joint density of $(\tilde{p}, \tilde{\mathbf{z}})$ is denoted $\tilde{f}(\tilde{p}, \tilde{\mathbf{z}})$, and the conditional density of $\tilde{\mathbf{z}}$ given $P > \lambda$ is $\tilde{h}_\lambda(\tilde{\mathbf{z}})$.

We estimate these densities using local likelihood KDE via the `locfit` R package. For $d \geq 3$, we omit cross-product terms to mitigate the curse of dimensionality. When $d < 3$, bandwidth selection is performed locally via k -nearest-neighbor cross-validation, following Chen, Robinson and Storey (2021) (hereafter CRS).

The desired densities are recovered as:

$$\hat{h}_\lambda(\mathbf{z}) = \frac{\tilde{h}_\lambda(\tilde{\mathbf{z}})}{\prod_{k=1}^d \phi(\tilde{z}^k)}, \quad \hat{f}(p, \mathbf{z}) = \frac{\tilde{f}(\tilde{p}, \tilde{\mathbf{z}})}{\phi(\tilde{p}) \prod_{k=1}^d \phi(\tilde{z}^k)},$$

where $\phi(x)$ is the standard normal density.

Choice of λ . The tuning parameter λ is selected from $\{0.4, 0.5, \dots, 0.8\}$ to minimize the mean integrated squared error of $\hat{\pi}_0(\mathbf{z}, \lambda)$, as in CRS.

Monotonic adjustment of $\hat{f}(p, \mathbf{z})$. Following footnote 6, if $f(p, \mathbf{z})$ is theoretically non-increasing in p for fixed $\mathbf{z}$, we impose this by adjusting:

$$\hat{f}(p_i, \mathbf{z}_i) := \min \left\{ \hat{f}(p_i, \mathbf{z}_i), \min_{p_l < p_i, \|\mathbf{z}_i - \mathbf{z}_l\| \leq \epsilon} \hat{f}(p_l, \mathbf{z}_l) \right\},$$

where $\epsilon = 0.01$.

Computing q -values and FDR control. We compute:

$$\hat{r}(p, \mathbf{z}) = \frac{\hat{\pi}_0(\mathbf{z})}{\hat{f}(p, \mathbf{z})},$$

and estimate the q -value:

$$\hat{q}(p_i, \mathbf{z}_i) = \frac{1}{\#S_i} \sum_{k \in S_i} \hat{r}(p_k, \mathbf{z}_k), \tag{A.4}$$

where $S_i = \{k : \hat{r}(p_k, \mathbf{z}_k) \leq \hat{r}(p_i, \mathbf{z}_i)\}$.

FDR control. Following Storey (2003), we note that controlling the “positive false discovery rate” (pFDR):

$$pFDR = \mathbb{E} \left(\frac{V}{R} \mid R > 0 \right),$$

is asymptotically equivalent to controlling FDR as n grows large—which holds in our application. For target level $\tau \in [0, 1]$, hypotheses $H_{0,i}$ are rejected if $\hat{q}(p_i, \mathbf{z}_i) \leq \tau$. As shown in CRS, the $mfFDR$ procedure thus controls FDR at level τ .

References

- Alexander, S. S. (1961) Price movements in speculative markets: Trends or random walks? *Industrial Management Review*, **2**, 7–26.
- Allen, H. and Taylor, M. P. (1990) Charts, noise and fundamentals in the london foreign exchange market. *The Economic Journal*, **100**, 49–59.
- Amihud, Y. and Goyenko, R. (2013) Mutual fund’s R^2 as predictor of performance. *Review of Financial Studies*, **26**, 667–694.
- Bajgrowicz, P. and Scaillet, O. (2012) Technical trading revisited: False discoveries, persistence tests, and transaction costs. *Journal of Financial Economics*, **106**, 473–491.
- Bank for International Settlements (2025) Triennial Central Bank Survey of foreign exchange and Over-the-counter (OTC) derivatives markets in 2025. Tech. rep., Bank for International Settlements, Basel. Available at <https://www.bis.org/statistics/rpfx25.htm>.
- Barras, L., Scaillet, O. and Wermers, R. (2010) False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance*, **65**, 179–216.
- Benjamini, Y. and Hochberg, Y. (1995) Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, **57**, 289–300.
- Burnside, C., Eichenbaum, M. S. and Rebelo, S. (2011) Carry trade and momentum in currency markets. Tech. rep., National Bureau of Economic Research.

- Cavalcante, R. C., Brasileiro, R. C., Souza, V. L., Nobrega, J. P. and Oliveira, A. L. (2016) Computational intelligence and financial markets: A survey and future directions. *Expert Systems with Applications*, **55**, 194–211.
- Chaboud, A. P., Chiquoine, B., Hjalmarsson, E. and Vega, C. (2014) Rise of the machines: Algorithmic trading in the foreign exchange market. *Journal of Finance*, **69**, 2045–2084.
- Chen, X., Robinson, D. G. and Storey, J. D. (2021) The functional false discovery rate with applications to genomics. *Biostatistics*, **22**, 68–81.
- Cheung, Y.-W. and Chinn, M. D. (2001) Currency traders and exchange rate dynamics: a survey of the us market. *Journal of international Money and Finance*, **20**, 439–471.
- Chinn, M. D. and Ito, H. (2006) What Matters for Financial Development? Capital Controls, Institutions, and Interactions. *Journal of Development Economics*, **81**, 163–192.
- Chordia, T., Goyal, A. and Saretto, A. (2020) Anomalies and false rejections. *Review of Financial Studies*, **33**, 2134–2179.
- Cornell, B. and Dietrich, J. K. (1978) The efficiency of the market for foreign exchange under floating exchange rates. *Review of Economics and Statistics*, **60**, 111–120.
- DeMiguel, V., Garlappi, L. and Uppal, R. (2007) Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, **22**, 1915–1953.
- Dooley, M. and Shafer, J. (1976) Analysis of short-run exchange rate behavior: March 1973 to july 1976. In *The International Transmission of Inflation* (ed. M. R. Darby), 555–579. University of Chicago Press.

- Fama, E. F. and French, K. R. (1993) Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, **33**, 3–56.
- Filippou, I., Gozluklu, A. E. and Taylor, M. P. (2018) Global political risk and currency momentum. *Journal of Financial and Quantitative Analysis*, **53**, 2227–2259.
- Filippou, I., Rapach, D., Taylor, M. P. and Zhou, G. (2023) Economic fundamentals and short-run exchange rate prediction: A machine learning perspective. *Available at SSRN 3455713*.
- Geenens, G. (2014) Probit transformation for nonparametric kernel estimation on the unit interval. *Journal of the American Statistical Association*, **109**, 346–358.
- Giglio, S., Liao, Y. and Xiu, D. (2021) Thousands of alpha tests. *Review of Financial Studies*, **34**, 3456–3496.
- Grinold, R. C. and Taylor, M. P. (2009) The opportunity set: Market opportunities and the effective breadth of a portfolio. *Journal of Portfolio Management*, **35**, 12–26.
- Hansen, P. R. (2005) A test for superior predictive ability. *Journal of Business & Economic Statistics*, **23**, 365–380.
- Harvey, C. R. and Liu, Y. (2020) False (and missed) discoveries in financial economics. *Journal of Finance*, **75**, 2503–2553.
- Harvey, C. R., Liu, Y. and Zhu, H. (2016) ...and the cross-section of expected returns. *Review of Financial Studies*, **29**, 5–68.

- Henderson, C. (2006) *Currency strategy: the practitioner's guide to currency investing, hedging and forecasting*. John Wiley & Sons.
- Hofert, M., Kojadinovic, I., Maechler, M. and Yan, J. (2025) *copula: Multivariate Dependence with Copulas*. URL <https://CRAN.R-project.org/package=copula>. R package version 1.1-6.
- Hsu, P.-H., Kyriakou, I., Ma, T. and Sermpinis, G. (2025) Mutual funds' conditional performance free of data snooping bias. *Journal of Financial and Quantitative Analysis*, **60**, 1373–1400.
- 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.
- Huang, J., Bai, L., Cui, B. *et al.* (2020) Leveraging biological and statistical covariates improves the detection power in epigenome-wide association testing. *Genome Biology*, **21**, 88.
- Ignatiadis, N. and Huber, W. (2021) Covariate powered cross-weighted multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **83**, 720–751.
- Ignatiadis, N., Klaus, B., Zaugg, J. B. and Huber, W. (2016) Data-driven hypothesis weighting increases detection power in genome-scale multiple testing. *Nature Methods*, **13**, 577–580.
- Ivanova, Y., Neely, C. J., Weller, P. and Famiglietti, M. T. (2021) Can risk explain the profitability of technical trading in currency markets? *Journal of International Money and Finance*, **110**, 102285.

- Jegadeesh, N. and Titman, S. (1993) Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, **48**, 65–91.
- Jennings, J., Kim, J. M., Lee, J. and Taylor, D. (2024) Measurement error, fixed effects, and false positives in accounting research. *Review of Accounting Studies*, **29**, 959–995.
- Joe, H. (1997) *Multivariate Models and Multivariate Dependence Concepts*. New York: Chapman and Hall/CRC, 1 edn.
- Lan, W. and Du, L. (2019) A factor-adjusted multiple testing procedure with application to mutual fund selection. *Journal of Business & Economic Statistics*, **37**, 147–157.
- Levich, R. M. and Thomas, L. R. (1993) The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach. *Journal of International Money and Finance*, **12**, 451–474.
- Levy, R. A. (1967) Relative strength as a criterion for investment selection. *Journal of Finance*, **22**, 595–610.
- Lo, A. W. (2004) The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, **30**, 15–29.
- Loader, C. (1999) *Local Regression and Likelihood*. Springer.
- Logue, D. E., Sweeney, R. J. and Willett, T. D. (1978) The theory of floating exchange rate regimes and the multinational corporation. *Journal of International Business Studies*, **9**, 15–27.

- Lustig, H., Roussanov, N. and Verdelhan, A. (2011) Common Risk Factors in Currency Markets. *Review of Financial Studies*, **24**, 3731–3777.
- Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A. (2012a) Carry trades and global foreign exchange volatility. *Journal of Finance*, **67**, 681–718.
- Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A. (2012b) Currency momentum strategies. *Journal of Financial Economics*, **106**, 660–684.
- Menkhoff, L. and Taylor, M. P. (2007) The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature*, **45**, 936–972.
- Moskowitz, T. J., Ooi, Y. H. and Pedersen, L. H. (2012) Time series momentum. *Journal of Financial Economics*, **104**, 228–250.
- Murphy, J. J. (1999) *Technical Analysis of the Financial Markets*. New York: New York Institute of Finance.
- Neely, C., Weller, P. and Dittmar, R. (1997) Is technical analysis in the foreign exchange market profitable? a genetic programming approach. *Journal of Financial and Quantitative Analysis*, **32**, 405–426.
- Neely, C. J. and Weller, P. A. (1999) Technical trading rules in the european monetary system. *Journal of International Money and Finance*, **18**, 429–458.
- Neely, C. J. and Weller, P. A. (2012) Technical analysis in the foreign exchange market. *Handbook of exchange rates*, 343–373.

- Neely, C. J. and Weller, P. A. (2013) Lessons from the evolution of foreign exchange trading strategies. *Journal of Banking and Finance*, **37**, 3783–3798.
- Neely, C. J., Weller, P. A. and Ulrich, J. M. (2009) The adaptive markets hypothesis: Evidence from the foreign exchange market. *Journal of Financial and Quantitative Analysis*, **44**, 467–488.
- Newey, W. K. and West, K. D. (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, **55**, 703–708.
- Newton, M. A., Noueir, A., Sarkar, D. and Ahlquist, P. (2004) Detecting differential gene expression with a semiparametric hierarchical mixture method. *Biostatistics*, **5**, 155–176.
- Osler, C. (2003) Currency orders and exchange-rate dynamics: Explaining the success of technical analysis. *Journal of Finance*, **58**, 1791–1819.
- Osler, C. (2005) Stop-loss orders and price cascades in currency markets. *Journal of International Money and Finance*, **24**, 219–241.
- Qi, M. and Wu, Y. (2006) Technical trading-rule profitability, data snooping, and reality check: Evidence from the foreign exchange market. *Journal of Money, Credit and Banking*, **38**, 2135–2158.
- Storey, J. D. (2002) A direct approach to false discovery rates. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **64**, 479–498.
- Storey, J. D. (2003) The positive false discovery rate: a Bayesian interpretation and the q -value. *Annals of Statistics*, **31**, 2013–2035.

- Storey, J. D., Akey, J. M. and Kruglyak, L. (2005) Multiple locus linkage analysis of genomewide expression in yeast. *PLOS Biology*, **3**, 1380–1390.
- Sweeney, R. J. (1986) Beating the foreign exchange market. *Journal of Finance*, **41**, 163–182.
- Taylor, M. P. and Allen, H. (1992) The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance*, **11**, 304–314.
- Wang, L., Han, X. and Tong, X. (2023) Skilled mutual fund selection: False discovery control under dependence. *Journal of Business & Economic Statistics*, **41**, 578–592.
- White, H. (2000) A reality check for data snooping. *Econometrica*, **68**, 1097–1126.
- Wilder, J. W. (1978) *New Concepts in Technical Trading Systems*. Greensboro, NC: Trend Research.

Internet Appendix for

“A Multivariate Covariate-Enhanced False Discovery Rate


Method: Evidence from Technical Trading in Currency Markets”

IA. The role of informative covariates in detecting false nulls

We begin by illustrating the differences across the three signal-strength settings. Figure I reports, from left to right, the distributions of p -values under the strong-, weak-, and very weak-signal cases. Panel A focuses on false null hypotheses only, while Panel B shows the corresponding distributions for the full population of hypotheses. In all three cases, the probability that a given hypothesis is null is generated from the same sine form of $\pi_0(u, v)$, and the p -values of the true null hypotheses are drawn from the uniform distribution on $[0, 1]$.

In the strong-signal case ($\alpha = 0.5$), the p -values of false null hypotheses are concentrated near zero. Under the weak-signal case, by contrast, these p -values are less concentrated and spread more substantially over the unit interval, with considerable mass extending well

Figure I: Distribution of p -values for false null hypotheses (Panel A) and for the full population of hypotheses (Panel B) under three signal-strength scenarios: strong, weak, and very weak.




beyond conventional significance cutoffs. In the very weak-signal case, the mode shifts even farther away from zero. As a result, false null hypotheses become progressively more difficult to detect as signal strength declines. For example, if one rejects a null whenever its p -value is below 0.05, a large fraction of false nulls is detected in the strong-signal case, whereas the detected fraction is much smaller in the very weak-signal case.²¹ In the context of technical trading rules, a weak signal corresponds to genuinely out- or under-performing rules having only small absolute Sharpe ratios, which in turn produce larger p -values and make them harder to distinguish from noise.

We next use a simulated sample to illustrate how informative covariates operate within the *FDR* framework. Specifically, we focus on the sample underlying Panel A of Figure I, generated under the strong-signal setting with the sine form of $\pi_0(u, v)$. To show how the covariates convey information about the hypotheses, we partition the sample into nine groups by dividing each of u and v into three equal intervals. For convenience, the groups are labeled 1 through 9, corresponding to the nine regions of the (u, v) space. Thus, group 1 contains hypotheses with $(u, v) \in [0, 1/3] \times [0, 1/3]$, group 2 contains hypotheses with $(u, v) \in [0, 1/3] \times (1/3, 2/3]$, and so on, up to group 9, which contains hypotheses with $(u, v) \in (2/3, 1] \times (2/3, 1]$. By construction, each group contains roughly the same number of hypotheses. The resulting group-specific p -value distributions are shown in Figure II.

To control the *FDR* at a target level τ —here, $\tau = 0.1$ —the *BH* and *StdFDR* procedures reject hypotheses solely on the basis of their p -values, using a common cutoff across all groups. By contrast, Figure II shows that the distributions of p -values vary markedly across groups.

²¹At a fixed threshold, the number of true null hypotheses incorrectly rejected is similar across cases. Hence, the false discovery proportion is substantially higher in the very weak-signal case.

Figure II: Distribution of p -values across the nine groups obtained by partitioning the sample in Panel A of Figure I according to the covariates u and v . In each sub-figure, the blue histogram shows the distribution of p -values for all hypotheses in the group, while the green histogram shows the distribution for false null hypotheses only. The title of each sub-figure reports the corresponding range of (u, v) .




In particular, groups 1, 2, 4, 6, 8, and 9 display much stronger left-skewness, especially in the green histograms, indicating that these groups contain relatively many false null hypotheses. By contrast, the nearly uniform blue histograms in groups such as 3 and 7 suggest that most hypotheses there are true nulls. This is consistent with the fact that the green histograms for groups 3 and 5 contain only a handful of false nulls among more than 1,000 hypotheses. Put differently, the true-null proportion is very high in those groups. The clear heterogeneity in the group-specific p -value distributions confirms that the appropriate rejection threshold should depend on informative covariates.

Accordingly, if the objective is to maximize the number of discoveries while controlling the FDR at a given target, it is natural to allow rejection thresholds to vary across groups rather than imposing a single universal cutoff, such as 0.05. For example, one might use a relatively lenient threshold in group 1 and a much stricter threshold in group 3. The $mfFDR$ implements this idea in a much finer and fully data-driven way, as explained below. The resulting gains in detecting true alternatives are illustrated more clearly in Figure III.

In Panel A1 (Panel B1) of Figure III, each line shows the proportion of hypotheses rejected by the four procedures across groups 1 through 9, for all hypotheses (false null hypotheses only). Here, the rejection proportion is defined as the number of rejected hypotheses in a group divided by the total number of hypotheses in that group. All procedures reject more hypotheses in groups 1 and 9 than in groups 3, 5, and 7. More importantly, the rejection proportions of $mfFDR$ and $fFDR$ are much higher than those of BH and $StdFDR$ in groups 1 and 9, whereas in group 7 they are slightly lower than that of $StdFDR$.

Panels A2 and B2 report the maximum p -value among rejected hypotheses. Because BH and $StdFDR$ use a single rejection threshold, their maximum rejected p -values are

Figure III: Comparison of procedures across the nine groups defined in Figure II. Panel A reports results for all hypotheses, and Panel B reports results for false null hypotheses only. Panels A1 and B1 show the proportion of rejections. Panels A2 and B2 show the maximum p -value among rejected hypotheses. Panels A3 and B3 show the minimum p -value among non-rejected hypotheses.



nearly constant across groups. In this example, those thresholds are approximately 0.011 for BH and 0.016 for $StdFDR$, as reflected in the green and brown lines in Panel A2. By contrast, $mfFDR$ allows thresholds to vary sharply across groups: in group 1 it rejects some hypotheses with p -values as high as 0.491, whereas in group 3 it rejects no hypotheses with p -values above 0.01, a cutoff that is even more stringent than that of $StdFDR$. Hence, some hypotheses rejected by $StdFDR$ are not rejected by $mfFDR$, while others with much larger p -values are rejected in groups rich in false nulls. Indeed, within the same group, it is possible for a hypothesis with a relatively large p -value to be rejected while another with a smaller p -value is not. For example, in group 1, one hypothesis with a p -value of 0.017 is not rejected by $mfFDR$, whereas other hypotheses in the same group with p -values as high as 0.491 are rejected.

Overall, the experiment shows that $mfFDR$ rejects more false null hypotheses in groups where false nulls are more prevalent. Although the grouping exercise is useful for illustrating the role of informative covariates, the $mfFDR$ procedure itself does not rely on partitioning hypotheses into a small number of coarse groups.²² Instead, $mfFDR$ effectively treats each hypothesis individually and assigns a covariate-specific rejection threshold, as discussed in Section 2.1. This is implemented through the q -value construction described in Appendix A.

²²Ignatiadis *et al.* (2016) and Ignatiadis and Huber (2021) propose group-weighting approaches that partition hypotheses into groups based on a single covariate and determine a rejection threshold for each group. However, these approaches are less powerful than the $fFDR$ procedure of Chen *et al.* (2021); see Section 1 of the Supplementary Materials of CRS.

IB. A Test for Covariate- p -Value Association

As illustrated in Section IA, one can assess the informativeness of a covariate by partitioning hypotheses according to the covariate and examining the resulting distribution of p -values. Huang *et al.* (2020) formalize this idea through an omnibus test that evaluates whether a single covariate is associated with the corresponding p -values. The procedure discretizes the covariate into groups, converts the p -value into a binary indicator based on a cutoff, and then applies classical tests of association between the resulting categorical variables. It searches across multiple choices of cutoff values and groupings, retaining the strongest observed association, and uses a permutation procedure to obtain a valid p -value. We briefly describe the test below.

Let $\{(p_i, z_i)\}_{i=1}^n$ denote the observed p -values and their corresponding covariate values. The null hypothesis is that p_i and z_i are independent.

Step 1. *Categorizing the p -value and the covariate.* Consider a set of p -value cutoffs $\{c_1, \dots, c_L\}$ and a set of candidate numbers of covariate groups $\{g_1, \dots, g_G\}$. For each g_j , discretize the covariate z_i into g_j quantile bins, denoted $y_i^{(j)}$, where $y_i^{(j)} = k$ if z_i belongs to bin k , for $k = 1, \dots, g_j$. For each cutoff c_ℓ , define a binary variable $x_i^{(\ell)}$ such that $x_i^{(\ell)} = 1$ if $p_i \leq c_\ell$ and $x_i^{(\ell)} = 0$ otherwise.

Thus, for each pair (c_ℓ, g_j) , we obtain a binary variable $x_i^{(\ell)}$ based on the p -value and a grouped variable $y_i^{(j)}$ based on the covariate.

Step 2. *Constructing association statistics.* For each pair (c_ℓ, g_j) , we compute two measures of association: (i) the p -value of the Pearson chi-square test, denoted $p\text{-value}_{\text{Pearson}}$, which tests for general dependence between $x_i^{(\ell)}$ and $y_i^{(j)}$; and (ii) the p -value of the Cochran-

Armitage trend test, denoted $p\text{-value}_{\text{CA}}$, which tests for a monotonic relation between $x_i^{(\ell)}$ and $y_i^{(j)}$. We transform both p -values to the log scale:

$$S_{\ell j}^{\text{Pearson}} = -\log(p\text{-value}_{\text{Pearson}}), \quad S_{\ell j}^{\text{CA}} = -\log(p\text{-value}_{\text{CA}}),$$

and define the combined statistic

$$S_{\ell j} = \max \{S_{\ell j}^{\text{Pearson}}, S_{\ell j}^{\text{CA}}\}.$$

The overall observed statistic is then

$$S_{\text{obs}} = \max_{\ell, j} S_{\ell j}.$$

Step 3. *Calculating permutation statistics.* For $b = 1, \dots, B$, randomly permute the binary variable associated with the p -values, recompute all statistics $S_{\ell j}^{(b)}$, and record

$$S_{\text{max}}^{(b)} = \max_{\ell, j} S_{\ell j}^{(b)}.$$

Step 4. *Obtaining the permutation p -value.* The omnibus-test p -value is defined as

$$p\text{-value}_{\text{omnibus}} = \frac{1 + \#\{S_{\text{max}}^{(b)} \geq S_{\text{obs}}\}}{B + 1}.$$

A small $p\text{-value}_{\text{omnibus}}$ indicates that small p -values are distributed non-uniformly with respect to the covariate, implying a statistically significant association between the covariate

and the p -values. In our empirical implementation, following [Huang *et al.* \(2020\)](#), we use data-driven p -value cutoffs given by the $\{0.1, 0.5, 1, 5, 10, 20\}\%$ quantiles of the observed p -values, candidate group counts $\{g_1 = 2, g_2 = 4, g_3 = 8\}$, and $B = 999$ permutations.

IC. Technical Trading Rules: Families, Parameter Spaces, and Configurations

This section documents in detail the five families of technical trading rules, parameter configurations, and data construction used in our empirical analysis. It supplements Section 2.2 of the main manuscript, providing the complete specification of the strategies included in our multiple testing framework. The trading rules and their configurations are identical to those used in [Hsu *et al.* \(2016\)](#), and this section draws on that paper. The choice of technical trading rules and the associated parameter spaces were chosen based on a close reading of the academic and practitioner literature of technical trading in FX markets (e.g., [Henderson, 2006](#); [Menkhoff and Taylor, 2007](#); [Hsu *et al.*, 2016](#)).

IC.1. Oscillator trading rules

Oscillator-based indicators aim to detect when recent price movements have been unusually rapid in one direction, suggesting a potential correction or reversal. A widely used example is the *relative strength index* (RSI), defined as:

$$\text{RSI}_t(h) = 100 \left(\frac{U_t(h)}{U_t(h) + D_t(h)} \right) \quad (\text{A1})$$

where $U_t(h)$ denotes the cumulated “up movement” (i.e. the close-to-close increase on a day when the exchange rate has closed higher than the previous day’s closing rate) over the previous h days, and $D_t(h)$ denotes the cumulated absolute “down movement” (the absolute

close-to-close decrease on a day when the exchange rate has closed lower than the previous day's closing rate) over the same period:

$$U_t(h) = \sum_{j=1}^h \iota(s_{t-j+1} - s_{t-j} > 0)(s_{t-j+1} - s_{t-j}) \quad (\text{A2})$$

$$D_t(h) = \sum_{j=1}^h \iota(s_{t-j+1} - s_{t-j} < 0) |s_{t-j+1} - s_{t-j}| \quad (\text{A3})$$

where $\iota(\cdot)$ is an indicator variable that takes the value one when the statement in parentheses is true and zero otherwise. The RSI thus attempts to measure the strength of “up movements” relative to the strength of “down movements” and is normalised to lie between 0 and 100. Common values at which a particular currency is deemed to have been overbought (signalling an imminent downward correction) or oversold (signalling an imminent upward correction) are 70 and 30, respectively.

A standard oscillator trading rule based on the RSI may be expressed as follows:

- **O1:** If $\text{RSI}_t(h)$ moves above $50 + v$ for at least d days and then subsequently moves below $50 + v$, go short the currency. If $\text{RSI}_t(h)$ moves below $50 - v$ for at least d days and then subsequently moves above $50 - v$, go long the currency.

A variation on the standard RSI trading rule imposes a pre-specified holding period for a position:

- **O2:** If $\text{RSI}_t(h)$ moves above $50 + v$ for at least d days and then subsequently moves below $50 + v$, go short the currency for k days and then neutralize the position. If $\text{RSI}_t(h)$ moves below $50 - v$ for at least d days and then subsequently moves above $50 - v$, go long the currency for k days and then neutralize the position.

Note that a trading signal is not generated when the RSI enters the overbought or oversold region (i.e. goes above $50 + v$ or below $50 - v$) but as it exits the region (i.e. crosses $50 + v$ from above or crosses $50 - v$ from below). This is because the currency can remain overbought or oversold for long periods and may become even more overbought or oversold for a while. The oscillator trading rule is designed to allow the spot rate to continue moving in the desired direction until a trend change becomes evident.

We consider a total of 600 oscillator trading rules. The specifications of oscillator trading rules we consider are described as below:

- $h \in \{5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$, $\#h = 10$.
- $v \in \{10, 15, 20, 25\}$, $\#v = 4$.
- $d \in \{1, 2, 5\}$, $\#d = 3$.
- $k \in \{1, 5, 10, 25\}$, $\#k = 4$.
- **O1**: This trading rule has three parameters, h, v and d . The number of versions of O1 considered is therefore $\#O1 = \#h \times \#v \times \#d = 10 \times 4 \times 3 = 120$.
- **O2**: This trading rule has four parameters, h, v, d and k . The number of versions of O2 considered is therefore $\#O2 = \#h \times \#v \times \#d \times \#k = 10 \times 4 \times 3 \times 4 = 480$.

IC.2. Filter rules

A filter rule involves buying a currency against another currency whenever the exchange rate has risen by more than a given percentage above its most recent low and selling it when the rate drops by more than a given percentage below its most recent high:

- **F1:** If the daily closing spot exchange rate of the currency moves up at least x percent above its most recent low and remains so for d days, go long the currency. If the daily closing spot rate moves down below its most recent high at least x percent and remains so for d days, go short the currency.

In this context, we define the most recent high (low) as the most recent closing price that is greater (less) than the j previous daily closing prices, for a given value of j . A variation on this basic filter rule allows for neutral positions rather than requiring that positions always be either long or short after the first position is initiated:

- **F2:** If the daily closing spot exchange rate of the currency moves up at least x percent above its most recent low and remains so for $d(x)$ days, go long the currency until its daily closing spot rate moves down at least y percent below the subsequent high and remains so for $d(y)$ days, at which time sell the currency and neutralize the long position. If the daily closing spot exchange rate of the currency moves down at least x percent below its most recent high and remains so for $d(x)$ days, go short the currency until its daily closing spot rate moves up at least y percent above the subsequent low and remains so for $d(y)$ days, at which time buy the currency and neutralize the short position. y is less than x . $d(y)$ is less than or equal $d(x)$.

Subsequent lows and highs in this context are again defined with respect to the previous j days.

A third, more naive variation on the filter rule imposes a pre-specified holding period that is adhered to ignoring all other signals:

- **F3:** If the daily closing spot exchange rate of the currency moves up at least x percent

above its most recent low and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot exchange rate of the currency moves down at least x percent below its most recent high and remains so for d days, go short the currency for k days and then neutralize the position.

In our empirical work, we examine a total of 2,835 filter rules, the parameterization of which is described as below:

- $x \in \{0.05, 0.1, 0.5, 1.0, 5.0, 10.0, 20.0\}$ in %, $\#x = 7$.
- $y \in \{0.05, 0.1, 0.5, 1.0, 5.0, 10.0, 20.0\}$ in %. Given $y < x$, the number of x - y combinations: $\#(x - y) = 21$.
- $d \in \{0, 1, 2, 3, 4, 5\}$, $\#d = 6$.
- $d(x) \in \{0, 1, 2, 3, 4, 5\}$, $d(y) \in \{0, 1, 2, 3, 4\}$. Given $d(y) \leq d(x)$, the number of $d(x)$ - $d(y)$ combinations: $\#[d(x)-d(y)] = \sum_{j=1}^5 j = 15$.
- $j \in \{1, 2, 5, 10, 20\}$, $\#j = 5$.
- $k \in \{5, 10, 15, 20, 25\}$, $\#k = 5$.
- **F1:** This trading rule has three parameters, x, d and j . The number of versions of F1 considered is therefore $\#F1 = \#x \times \#d \times \#j = 7 \times 6 \times 5 = 210$.
- **F2:** This trading rule has five parameters, $x, y, d(x), d(y)$ and j . The number of versions of F2 considered is therefore $\#F2 = \#(x - y) \times \#[d(x)-d(y)] \times \#j = 21 \times 15 \times 5 = 1,575$.
- **F3:** This trading rule has four parameters, x, d, j and k . The number of versions of F3 considered is therefore $\#F3 = \#x \times \#d \times \#k \times \#j = 7 \times 6 \times 5 \times 5 = 1,050$.

IC.3. Moving average trading rules

This class of trading rules seeks to capture sustained price trends and to identify potential trend reversals or the emergence of new trends. In a simple moving average (MA) rule, the local trend is approximated by a moving average of past prices, and a trading signal is generated when the spot rate crosses this moving average—indicating a possible shift in trend direction and prompting the initiation or closure of positions. More commonly, traders employ pairs of moving averages with differing horizons. In such cases, a bullish signal arises when a short-term moving average crosses above a longer-term moving average, while a bearish signal is indicated when the short-term moving average crosses below the longer-term average.

If we define the simple moving average of the exchange rate over j days as

$$\text{MA}_t(j) = \frac{1}{j} \sum_{i=0}^{j-1} s_{t-i} \quad (\text{A4})$$

then a single moving average trading rule may be expressed as follows:

- **MA1:** If the daily closing spot exchange rate of the currency moves up at least x percent above $\text{MA}_t(q)$ and remains so for d days, go long the currency until its daily closing spot rate moves down at least x percent below $\text{MA}_t(q)$ and remains so for d days, at which time go short the currency. If the daily closing spot exchange rate of the currency moves down at least x percent below $\text{MA}_t(q)$ and remains so for d days, go short the currency until the daily closing spot rate moves up at least x percent above $\text{MA}_t(q)$ and remains so for d days, at which time go long the currency.

As before, a simple variation that pre-specifies the holding period, ignoring all other signals

during the holding period, would be as follows:

- **MA2:** If the daily closing spot exchange rate of the currency moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot exchange rate of the currency moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency for k days and then neutralize the position.

A double moving average trading rule may be formulated in the following fashion:

- **MA3:** If $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency until $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, at which time go short the currency. If $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency until $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, at which time go long the currency. p is less than q .

And the pre-specified holding period version is of the form:

- **MA4:** If $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency for k days and then neutralize the position. If $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency for k days and then neutralize the position. p is less than q .

Finally, we consider a triple moving average rule:

- **MA5:** If the daily closing spot exchange rate of the currency moves up at least x percent above any two of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go

long the currency with one third of the risk budget (currency position=+1/3). If the daily closing spot rate of the currency moves up at least x percent above all three of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go long the currency with the whole of the risk budget (currency position=+1). If the daily closing spot exchange rate of the currency moves down at least x percent below any two of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go short the currency with one third of the risk budget (currency position=-1/3). If the daily closing spot rate moves down at least x percent below all three of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go short the currency with the whole of the risk budget (currency position=-1). n is less than p , which is less than q .

Clearly, the trading rule never indicates a neutral position—for example if the spot rate is above only one moving average, it must be below two of them and a one-third short position is indicated. We consider a total of 12,870 moving average trading rules, the details of the specifications are described as below:

- $q \in \{2, 5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$, $\#q = 11$.
- $p \in \{2, 5, 10, 15, 20, 25, 50, 100, 150, 200\}$. Given $p < q$, the number of p - q combinations:

$$\#(p - q) = \sum_{j=1}^{10} j = 55.$$
- $n \in \{2, 5, 10, 15, 20, 25, 50, 100, 150\}$. Given $n < p < q$, the number of n - p - q combinations: $\#(n - p - q) = 165$.
- $x \in \{0, 0.05, 0.1, 0.5, 1.0, 5.0\}$ in %, $\#x = 6$.
- $d \in \{0, 2, 3, 4, 5\}$, $\#d = 5$.

- $k \in \{5, 10, 25\}$, $\#k = 3$.
- **MA1:** This trading rule has three parameters, q, x and d . The number of versions of MA1 considered is therefore $\#MA1 = \#q \times \#x \times \#d = 11 \times 6 \times 5 = 330$.
- **MA2:** This trading rule has four parameters, q, x, d and k . The number of versions of MA2 considered is therefore $\#MA2 = \#q \times \#x \times \#d \times \#k = 11 \times 6 \times 5 \times 3 = 990$.
- **MA3:** This trading rule has four parameters, p, q, x and d . The number of versions of MA3 considered is therefore $\#MA3 = \#(p - q) \times \#x \times \#d = 55 \times 6 \times 5 = 1,650$.
- **MA4:** This trading rule has five parameters, p, q, x, d and k . The number of versions of MA4 considered is therefore $\#MA4 = \#(p - q) \times \#x \times \#d \times \#k = 55 \times 6 \times 5 \times 3 = 4,950$.
- **MA5:** This trading rule has five parameters, n, p, q, x and d . The number of versions of MA5 considered is therefore $\#MA5 = \#(n - p - q) \times \#x \times \#d = 165 \times 6 \times 5 = 4,950$.

IC.4. Support-resistance trading rules

Support-resistance trading rules aim to identify price levels at which the exchange rate exhibits persistent difficulty advancing (resistance) or declining (support). The underlying idea is that these levels represent psychological or structural barriers in the market. While conceptually related to filter rules, support-resistance strategies differ in their signal generation: rather than responding to breaks of recent highs or lows, they trigger trades when the exchange rate surpasses established support or resistance levels by a specified percentage threshold.

The support and resistance levels have to be pre-specified. We define a resistance level as the highest closing rate of the j previous closing rates and we define a support level as

the lowest closing rate of j previous closing rates:

- **SR1:** If the daily closing spot exchange rate of the currency moves up at least x percent above the highest closing of the j previous closing rates and remains so for d days, go long the currency. If the daily closing spot rate moves down at least x percent below the lowest closing of the j previous closing rates and remains so for d days, go short the currency.

The pre-specified holding period version of the support-resistance rule is also analogous to the corresponding filter rule:

- **SR2:** If the daily closing spot exchange rate of the currency moves up at least x percent above the highest closing of the j previous closing rates and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot rate moves down at least x percent below the lowest closing of the j previous closing rates and remains so for d days, go short the currency for k days and then neutralize the position.

The details of our parameterization of a total of 1,890 support-resistance trading rules are provided as below:

- $x \in \{0.05, 0.1, 0.5, 1.0, 2.5, 5.0, 10.0\}$ in %, $\#x = 7$.
- $d \in \{0, 1, 2, 3, 4, 5\}$, $\#d = 6$.
- $j \in \{2, 5, 10, 15, 20, 25, 50, 100, 250\}$, $\#j = 9$.
- $k \in \{1, 5, 10, 25\}$, $\#k = 4$.

- **SR1:** This trading rule has three parameters, x, d and j . The number of versions of SR1 considered is therefore $\#SR1 = \#x \times \#d \times \#j = 7 \times 6 \times 9 = 378$.
- **SR2:** This trading rule has four parameters, x, d, j and k . The number of versions of SR2 considered is therefore $\#SR2 = \#x \times \#d \times \#j \times \#k = 7 \times 6 \times 9 \times 4 = 1,512$.

IC.5. Channel breakout trading rules

A trading channel is said to occur when the highest and lowest daily closing exchange rates over a given historical window remain within a narrow range, suggesting that dynamic support and resistance levels are moving together over time. These evolving upper and lower bounds define the channel. Channel breakout strategies generate trading signals when the exchange rate moves beyond these bounds, based on the premise that such a breakout is likely to initiate a sustained trend in the direction of the breach.

Formally, a $c\%$ trading channel is identified when the maximum daily closing exchange rate over the past j days lies within $c\%$ of the corresponding minimum. The resulting upper bound on day t is computed as $c\%$ above the lowest close over the prior j days, while the lower bound is $c\%$ below the highest close in the same window. A signal is triggered when the spot rate crosses either threshold.

- **CB1:** If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves up at least x percent above the upper bound of the channel and remains so for d days, go long the currency. If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves down at least x percent below the lower bound of the channel and remains so for d days, go short the currency.

And the pre-specified holding period version is:

- **CB2:** If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves up at least x percent above the upper bound of the channel and remains so for d days, go long the currency for k days and then neutralize the position. If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves down at least x percent below the lower bound of the channel and remains so for d days, go short the currency for k days and then neutralize the position.

We analyse a total of 3,000 channel breakout rules, the parameterization of which is described as below:

- $x \in \{0.05, 0.1, 0.5, 1.0, 5.0\}$ in %, $\#x = 5$.
- $d \in \{0, 1, 2\}$, $\#d = 3$.
- $j \in \{5, 10, 15, 20, 25, 50, 100, 200\}$, $\#j = 8$.
- $c \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$ in %, $\#c = 5$.
- $k \in \{1, 5, 10, 25\}$, $\#k = 4$.
- **CB1:** This trading rule has four parameters, x, d, j and c . The number of versions of CB1 considered is therefore $\#CB1 = \#x \times \#d \times \#j \times \#c = 5 \times 3 \times 8 \times 5 = 600$.
- **CB2:** This trading rule has five parameters, x, d, j, c and k . The number of versions of CB2 considered is therefore $\#CB2 = \#x \times \#d \times \#j \times \#c \times \#k = 5 \times 3 \times 8 \times 5 \times 4 = 2,400$.

IC.6. Total technical trading rules

The total number of different calibrations of trading rules considered is therefore: $(\#O1 + \#O2) + (\#F1 + \#F2 + \#F3) + (\#MA1 + \#MA2 + \#MA3 + \#MA4 + \#MA5) + (\#SR1 +$

$$\#SR2) + (\#CB1 + \#CB2) = 600 + 2,835 + 12,870 + 1,890 + 3,000 = 21,195.$$

ID. Performance of *mfFDR* under Different Numbers of Tests

In the main manuscript, our simulation exercises use $n = 10,000$ tests, which is close to the number of hypotheses considered in the empirical application. To assess the performance of *mfFDR* more broadly, we conduct an additional robustness analysis using smaller numbers of tests. Specifically, we repeat the simulations for $n = 500$, $1,000$, and $5,000$.

Figure IV compares the performance of *mfFDR* with that of the benchmark *stdFDR* procedure in terms of false discovery rate control and power. Several findings emerge.

First, under weak and very weak signal settings, *mfFDR* controls the false discovery rate well across the range of target levels, regardless of the number of tests. Second, under strong signals, *mfFDR* shows some over-rejection at high target levels when the number of tests is small. In particular, when $n = 500$, *mfFDR* controls the false discovery rate well up to a target of about 0.2. When $n = 1,000$ and $n = 5,000$, this range expands to approximately 0.3 and 0.4, respectively. These results suggest that when the number of tests is relatively small, *mfFDR* is most reliable at lower false-discovery targets, especially in strong-signal environments.

That said, this limitation is unlikely to be practically important in most finance and economics applications, where researchers rarely work with such small test universes and where very high false-discovery targets are of limited interest. Finally, in terms of power, *mfFDR* consistently outperforms *stdFDR* in all cases, regardless of the number of tests.

IE. Correlation and Estimation Error in Covariates

In this section, we examine the performance of $mfFDR$ when the informative covariates are correlated.²³ Starting from the two-covariate design in the previous section, we use a Cholesky transformation to construct new covariates with pairwise correlation approximately equal to r . The simulated p -values are then generated as in the baseline design. For brevity, we report results only for the sine specification of the null-proportion function.²⁴ To assess the effect of covariate correlation on performance, we consider values of r ranging from 0.1 to 0.8.

Figure V reports the results. Panel A shows the false discovery rate of $mfFDR$ under correlated covariates, while Panels B and C report the power gains of $mfFDR$ relative to $fFDR$ and the standard FDR procedure of Storey (2002) ($stdFDR$), respectively. In terms of false discovery rate control, $mfFDR$ performs well when $r < 0.7$, regardless of the signal strength. This range covers the vast majority of cases in our empirical application: as shown in Section 5 of the main manuscript, more than 95% of the empirical pairwise correlations have absolute value below 0.7. When $r \geq 0.7$, $mfFDR$ continues to control the false discovery rate well under weak and very weak signals, and remains asymptotically well behaved under strong signals.

Because the Cholesky transformation changes the effective signal strength across correlation settings, the absolute power levels of $mfFDR$ are not directly comparable across values

²³The results for negatively correlated covariates are similar.

²⁴Because the transformed covariates are no longer uniformly distributed on $[0, 1]$, we modify the null-proportion function slightly to ensure that it remains well defined:

$$\pi_0(u, v) = \min \{1, \max \{0, \sin(\pi(u + v)/2)\}\}.$$

of r . To address this issue, we focus instead on relative power. Panels B and C therefore plot the differences in power between $mfFDR$ and its benchmarks for each value of r . Values above zero indicate that $mfFDR$ is more powerful than the competing procedure. This is true in all cases considered. At the 20% false discovery target used in the main empirical analysis, the power gains can be substantial, reaching roughly 30% relative to $fFDR$ and as much as 70% relative to $stdFDR$, depending on the correlation level and signal strength.

IF. Performance of $mfFDR$ under Noisy Covariates

As noted in Section IE, the covariates used in empirical applications are typically estimated rather than directly observed, and therefore contain noise. In this section, we examine whether such estimation error materially affects the performance of $mfFDR$.

We consider a simple setting in which the input covariates are measured with noise. Specifically, we use the covariates $u = (u_1, \dots, u_n)$ and $v = (v_1, \dots, v_n)$ from the baseline simulations in Section 2.2 to generate the data, but apply $mfFDR$ using the noisy covariates

$$u'_i = u_i + \eta_i \tag{F.1}$$

and

$$v'_i = v_i + \epsilon_i, \tag{F.2}$$

where η_i and ϵ_i are independently drawn from $N(0, \sigma^2)$ for $i = 1, \dots, n$.

To study the effect of noise, we consider three values of σ :

$$\sigma \in \left\{ \frac{0.5}{\sqrt{12}}, \frac{1.0}{\sqrt{12}}, \frac{1.5}{\sqrt{12}} \right\}.$$

These values are chosen relative to the standard deviation of the original covariates, since u and v are uniformly distributed on $[0, 1]$ and therefore each has standard deviation $1/\sqrt{12}$. Larger values of σ correspond to less informative observed covariates and hence to a weaker relation between the covariates used by the procedure and the underlying null status of the hypotheses.

Figure VI reports the performance of $mfFDR$ under noisy covariates in terms of false discovery rate control and power, for target levels up to 0.45.

The results show that $mfFDR$ continues to control the false discovery rate well across the full range of target values considered. To assess the impact of noise on power, we compare $mfFDR$ based on the noisy covariates (u', v') with $fFDR$ based on the noisy univariate covariate u' and with $stdFDR$. As expected, measurement error in the covariates reduces the power of $mfFDR$ relative to the case of perfectly observed covariates. Even so, the method remains substantially more powerful than both $fFDR$ and $stdFDR$. These results suggest that $mfFDR$ is robust to realistic levels of covariate noise, an important property for empirical applications in which informative covariates must themselves be estimated from the data.

IG. Performance of $mfFDR$ under non IID settings

In a particular category of trading rules, the rules with close parameters tend to have highly correlated returns. This leads to a dependency among the testing statistics or p -values and, potentially, the covariates of the corresponding hypotheses. This section shows that our method is robust under this type of dependence. Specifically, we are generating data such that the hypotheses are partitioned into groups with the same size k . In each

group, the values of the concerning component (p -value or covariates) of testing hypotheses are mutually dependent at the same level, which is characterized by a covariance matrix Σ . The values corresponding to the hypotheses from different groups are independent. In the follows, we will elaborately describe this for each of the components: the p -value and the covariates.

IG.1. Dependence in p -values

For the p -value case, the data-generating process is as follows. The simulated covariates, $\pi_0(u, v)$, and the true status of null hypotheses H are as described in Section 2.2. For the sake of space, we present only the results for the sine form of $\pi_0(u, v)$. To account for the dependence in hypotheses, we first partition the true null hypotheses into groups of size k . Secondly, we generate the z -scores for the null hypotheses of each group from the multivariate normal distribution $\mathcal{N}(0, \Sigma)$. We process similarly for the false null ones, but the z -scores of each group are drawn from $\mathcal{N}(2, \Sigma)$. To simplify, the matrix $\Sigma = (\Sigma_{ij})_{k \times k}$ is set as $\Sigma_{ii} = 1$ and $\Sigma_{ij} = c$ for $i \neq j = 1, \dots, k$ for some c . By considering various values of the parameters k and c , we reveal the impact of the dependence at different levels on the performance of the $mfFDR$. Here, we choose $c \in \{0, 0.25, 0.5, 0.75, 0.9\}$ where the case $c = 0$ indicates the absence of the dependence among p -values and will be used as a benchmark for a comparison purpose, and $k \in \{10, 100, 500\}$. The parameter k represents the dependence scale. A larger k indicates the presence of more hypotheses that are mutually dependent.²⁵ Finally, the p -value of each (two-sided) test is calculated from its z -score by using the cumulative distribution function of the standard normal distribution. Note that,

²⁵This type of dependent setting is also studied in Storey (2003) with $k = 10$. Here, we extend the cases of k to study its impact on the method's performance.

in this setting, the covariates convey only the information on the probability of being true null of the hypotheses (i.e. via $\pi_0(u, v)$) and not on the p -values (i.e. we do not generate p -values via $f_{alt}(p|(u, v))$, but via the dependent z -scores, which does not depend on the covariates) as in the simulations conducted in the previous sections.

Figure VII presents the performance of $mfFDR$. Panel A of the figure shows that the considering dependence type does not affect the FDR control of the method when the FDR target is less than 40%. When k is small, the first two sub-figures of Panel B show an insignificant difference between the power of $mfFDR$ with different levels of dependence among z -scores in each group. Therefore, the lines corresponding to different values of c are virtually identical and covered by a blue line. Their powers are only distinguishable in the third sub-figure, where k is large. In this case, the power is decreasing concerning c , i.e. the level of dependence among the z -scores.

IG.2. Dependence in covariates

Next, in our dependent covariates, the data-generating process is as follows. We partition hypotheses into n/k equal groups, i.e., each group contains k hypotheses. Each of two covariates u and v is generated via a copula so that its values of those in the same group are correlated with a coefficient c (see Joe, 1997).²⁶ The two covariates are generated independently of each other. Once the covariates are generated, the $\pi_0(u, v)$, the true status of null hypotheses H , and the p -values are as described in Section 2.2. Note that, since the status of true null is implicitly non i.i.d. as it is generated using the non i.i.d. covariates. For the sake of space, we present only the results for the sine form of $\pi_0(u, v)$. Similar to the

²⁶One can easily generate such random variables using the freely available R package `Copula` of Hofert et al. (2025).

p -value case, we choose $c \in \{0, 0.25, 0.5, 0.75, 0.9\}$ and $k \in \{10, 100, 500\}$.

Figure VIII shows the performance of the $mfFDR$. Panel A of the figure shows that the presence of non i.i.d. covariates virtually does not affect the FDR control when the FDR target is up to 40% except the most extremely dependent case ($k = 500$ and $c = 0.9$) - where the $mfFDR$ controls well for targets up to 30% and does asymptotically so for higher targets. When k is small, the first two sub-figures of Panel B show an insignificant difference between the power of $mfFDR$ with different levels of dependence in each covariate. The powers are only distinguishable in the third sub-figure, where k is large. Thereby, the power is decreasing with respect to c , i.e. the level of dependence within groups of each covariate.


IIH. Performance of $mfFDR^+$ -Based Portfolios under Alternative FDR Targets

In this section, we examine the performance of $mfFDR^+$ -based portfolios under alternative FDR targets, both at the individual-currency level and for the pooled portfolio across all currencies. Specifically, we repeat the main empirical exercise with the FDR target τ varying from 10% to 40%. The results for individual-currency portfolios are reported in Table I, while the corresponding results for the pooled cross-currency portfolios are reported in Table II. For ease of comparison, we also reproduce the benchmark case of $\tau = 0.20$.

Overall, portfolio performance is stable across the range of FDR targets considered. For the pooled cross-currency portfolios, higher target levels are associated with slightly higher Sharpe ratios but somewhat lower annual returns. This pattern suggests that, when the FDR target is set more conservatively, the resulting portfolios tend to exhibit higher daily-return volatility.

Figure IV: Performance comparison of *mfFDR* and the standard *FDR* procedure of Storey (2002) (*stdFDR*) for different numbers of tests (n). Panel A reports results when $\pi_0(u, v)$ has the sine form, and Panel B reports results when $\pi_0(u, v)$ is monotonic in each covariate.

Panel A: $\pi_0(u, v)$ is sine-shaped.



Panel B: $\pi_0(u, v)$ is monotonic in each covariate.






Figure V: Performance of $mfFDR$ under correlated covariates. Panel A reports false discovery rate control. Panels B and C report the power gains of $mfFDR$ relative to $fFDR$ and the standard FDR procedure of Storey (2002) ($stdFDR$), respectively.

Panel A: False discovery rate control of $mfFDR$ under correlated covariates



Panel B: Power gain of $mfFDR$ over $fFDR$ under correlated covariates



Panel C: Power gain of $mfFDR$ over $stdFDR$ under correlated covariates





Figure VI: Performance of $mfFDR$ under noisy covariates. The figure compares $mfFDR$, $fFDR$, and the standard FDR procedure of Storey (2002) ($stdFDR$) when the covariates used by the covariate-adaptive procedures are measured with noise. For $mfFDR$, the inputs are $u' = u + \eta$ and $v' = v + \epsilon$, where $\eta, \epsilon \sim N(0, \sigma^2)$ and $\sigma \in \{0.5/\sqrt{12}, 1.0/\sqrt{12}, 1.5/\sqrt{12}\}$. For $fFDR$, the noisy covariate is u' . Panel A reports results for the sine specification of $\pi_0(u, v)$, and Panel B for the monotonic specification.

Panel A: $\pi_0(u, v)$ is a sine function.



Panel B: $\pi_0(u, v)$ is monotonic in each covariate.





Figure VII: Performance of the $mfFDR$ under correlated p -values. Panel A (B) presents the FDR control (power) of the $mfFDR$ under different levels of dependencies among p -value of the tests.

Panel A: FDR control of the $mfFDR$ under dependent p -values



Panel B: Power of the $mfFDR$ under dependent p -values





Figure VIII: Performance of the $mfFDR$ under non i.i.d. covariates. Panel A (B) presents the FDR control (power) of the $mfFDR$ under different levels of dependencies among values of each covariate of the tests.

Panel A: FDR control of the $mfFDR$ under non i.i.d. covariates



Panel B: Power of the $mfFDR$ under non i.i.d. covariates




Table I: Performance of $mfFDR$ based portfolios on individual currency with varying FDR target. The table shows annualized Sharpe ratios of the $mfFDR$ based portfolio with FDR target $\tau = \{0.1, 0.2, 0.3, 0.4\}$ based on portfolios' returns before (left side) and after transaction cost (right side). The final row shows the average Sharpe ratio across 30 portfolios corresponding to the 30 currencies. The numbers in parentheses are the corresponding p -values. “*”, “**” and “***”, respectively, indicate statistical significance at levels of 10%, 5% and 1%.

Country	Before transaction cost				After transaction cost			
	$\tau = 0.1$	$\tau = 0.2$	$\tau = 0.3$	$\tau = 0.4$	$\tau = 0.1$	$\tau = 0.2$	$\tau = 0.3$	$\tau = 0.4$
Australia	0.17 (0.21)	0.18 (0.17)	0.18 (0.17)	0.18 (0.17)	0.13 (0.34)	0.14 (0.29)	0.14 (0.30)	0.14 (0.31)
Canada	0.19 (0.17)	0.17 (0.23)	0.17 (0.25)	0.17 (0.24)	0.13 (0.37)	0.11 (0.46)	0.11 (0.47)	0.11 (0.46)
Germany/E.U.	0.48 (0.00)***	0.49 (0.00)***	0.48 (0.00)***	0.48 (0.00)***	0.45 (0.00)***	0.46 (0.00)***	0.45 (0.00)***	0.45 (0.00)***
Japan	0.46 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.44 (0.00)***	0.43 (0.00)***	0.41 (0.00)***	0.41 (0.00)***	0.41 (0.00)***
New Zealand	0.32 (0.01)***	0.34 (0.00)***	0.35 (0.00)***	0.35 (0.00)***	0.27 (0.02)**	0.28 (0.02)**	0.29 (0.01)***	0.29 (0.01)***
Norway	0.20 (0.15)	0.18 (0.17)	0.18 (0.19)	0.18 (0.18)	0.15 (0.29)	0.13 (0.36)	0.13 (0.38)	0.13 (0.37)
Sweden	0.37 (0.00)***	0.40 (0.00)***	0.40 (0.00)***	0.40 (0.00)***	0.33 (0.02)**	0.36 (0.01)***	0.36 (0.00)***	0.36 (0.00)***
Switzerland	0.26 (0.06)*	0.26 (0.06)*	0.25 (0.07)*	0.25 (0.07)*	0.22 (0.11)	0.22 (0.12)	0.22 (0.12)	0.21 (0.13)
U.K.	0.29 (0.05)**	0.29 (0.05)**	0.28 (0.05)**	0.28 (0.06)*	0.26 (0.07)*	0.25 (0.08)*	0.24 (0.10)*	0.24 (0.10)*
Argentina	0.35 (0.06)*	0.35 (0.06)*	0.34 (0.06)*	0.35 (0.06)*	0.32 (0.10)*	0.31 (0.11)	0.29 (0.11)	0.30 (0.12)
Columbia	0.56 (0.00)***	0.60 (0.00)***	0.62 (0.00)***	0.61 (0.00)***	0.48 (0.01)***	0.53 (0.01)***	0.55 (0.00)***	0.54 (0.01)***
India	0.35 (0.03)**	0.34 (0.03)**	0.34 (0.03)**	0.34 (0.04)**	0.31 (0.06)*	0.30 (0.07)*	0.30 (0.07)*	0.30 (0.07)*
Indonesia	0.33 (0.02)**	0.34 (0.02)**	0.33 (0.02)**	0.33 (0.02)**	0.23 (0.10)*	0.24 (0.09)*	0.24 (0.09)*	0.24 (0.10)*
Israel	0.52 (0.00)***	0.54 (0.00)***	0.54 (0.00)***	0.54 (0.00)***	0.37 (0.02)**	0.39 (0.01)***	0.39 (0.01)***	0.39 (0.01)***
Philippines	0.64 (0.00)***	0.62 (0.00)***	0.62 (0.00)***	0.62 (0.00)***	0.48 (0.01)***	0.46 (0.01)***	0.46 (0.01)***	0.46 (0.01)***
Romania	0.19 (0.31)	0.23 (0.19)	0.24 (0.16)	0.25 (0.15)	0.09 (0.65)	0.12 (0.51)	0.13 (0.47)	0.13 (0.46)
Russia	0.47 (0.01)***	0.51 (0.01)***	0.52 (0.01)***	0.52 (0.01)***	0.44 (0.02)**	0.48 (0.01)***	0.49 (0.01)***	0.49 (0.01)***
Slovak	0.16 (0.40)	0.18 (0.35)	0.19 (0.33)	0.19 (0.31)	0.11 (0.51)	0.13 (0.48)	0.13 (0.46)	0.14 (0.43)
Brazil	0.36 (0.07)*	0.36 (0.07)*	0.35 (0.08)*	0.35 (0.08)*	0.33 (0.10)*	0.33 (0.10)*	0.32 (0.10)*	0.32 (0.10)*
Chile	0.49 (0.01)***	0.46 (0.02)**	0.45 (0.02)**	0.45 (0.02)**	0.41 (0.04)**	0.38 (0.06)*	0.37 (0.06)*	0.37 (0.06)*
Czech	0.07 (0.64)	0.12 (0.46)	0.13 (0.43)	0.14 (0.41)	0.01 (0.85)	0.06 (0.65)	0.07 (0.61)	0.08 (0.60)
Hungary	-0.12 (0.56)	-0.11 (0.60)	-0.11 (0.60)	-0.11 (0.62)	-0.17 (0.39)	-0.17 (0.41)	-0.16 (0.41)	-0.16 (0.43)
Korea	0.30 (0.15)	0.29 (0.16)	0.29 (0.16)	0.29 (0.16)	0.24 (0.23)	0.23 (0.25)	0.23 (0.26)	0.23 (0.26)
Mexico	0.10 (0.48)	0.10 (0.49)	0.09 (0.50)	0.10 (0.49)	0.05 (0.71)	0.04 (0.73)	0.04 (0.74)	0.04 (0.73)
Poland	-0.02 (0.90)	0.02 (0.92)	0.02 (0.88)	0.03 (0.85)	-0.07 (0.73)	-0.03 (0.89)	-0.03 (0.91)	-0.02 (0.93)
Singapore	0.28 (0.05)**	0.32 (0.02)**	0.32 (0.02)**	0.32 (0.02)**	0.11 (0.52)	0.15 (0.32)	0.15 (0.32)	0.15 (0.33)
South Africa	0.19 (0.19)	0.23 (0.14)	0.25 (0.10)*	0.25 (0.09)*	0.07 (0.61)	0.10 (0.49)	0.12 (0.43)	0.12 (0.43)
Taiwan	0.73 (0.00)***	0.73 (0.00)***	0.73 (0.00)***	0.73 (0.00)***	0.66 (0.00)***	0.66 (0.00)***	0.65 (0.00)***	0.65 (0.00)***
Thailand	0.46 (0.02)**	0.47 (0.02)**	0.46 (0.02)**	0.46 (0.02)**	0.37 (0.05)**	0.38 (0.04)**	0.38 (0.04)**	0.38 (0.04)**
Turkey	0.51 (0.00)***	0.51 (0.00)***	0.51 (0.00)***	0.50 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.44 (0.00)***
Average	0.32	0.33	0.33	0.33	0.25	0.26	0.26	0.26

Table II: Performance of $mfFDR^+$ based portfolio with various FDR targets. The table shows the annualized Sharpe ratios and mean returns (before and after transaction cost) and break-even point by implementing the $mfFDR^+$ on all strategies in all currencies to control the FDR at 10% (Panel A), 20% (Panel B), 30% (Panel C) and 40% (Panel D). The selected out-performing strategies are then combined by currency. The fund is allocated to trade on each of the currencies having out-performing strategies weighted by the size of the selected out-performing trading rules. The first row presents the numbers of the whole sample period while the rest are those of sub-periods.

Period	Excess SR	Net SR	Excess Return (%)	Net Return (%)	Break-even (bps)
Panel A: FDR target of 10%					
Whole Period	0.98	0.88	4.36	3.91	63
1973-1980	1.40	1.31	4.89	4.58	76
1981-1990	2.11	2.01	8.28	7.89	141
1991-2000	0.81	0.72	4.93	4.38	84
2001-2010	0.62	0.47	2.53	1.89	34
2011-2020	0.33	0.25	1.28	0.99	10
Panel B: FDR target of 20%					
Whole Period	1.06	0.95	3.80	3.40	60
1973-1980	1.45	1.35	4.47	4.18	69
1981-1990	2.08	1.97	7.30	6.93	128
1991-2000	0.92	0.81	4.15	3.64	72
2001-2010	0.69	0.53	2.37	1.82	34
2011-2020	0.29	0.21	0.88	0.63	14
Panel C: FDR target of 30%					
Whole Period	1.11	0.99	3.54	3.15	57
1973-1980	1.49	1.39	4.29	4.01	76
1981-1990	2.02	1.91	6.69	6.33	118
1991-2000	1.00	0.87	3.77	3.28	66
2001-2010	0.72	0.56	2.28	1.75	34
2011-2020	0.32	0.23	0.83	0.60	8
Panel D: FDR target of 40%					
Whole Period	1.12	1.00	3.34	2.97	54
1973-1980	1.47	1.37	4.11	3.84	75
1981-1990	1.98	1.87	6.30	5.94	112
1991-2000	1.06	0.92	3.57	3.10	63
2001-2010	0.73	0.57	2.20	1.70	34
2011-2020	0.29	0.20	0.72	0.50	7

II. Portfolios Conditional on Rule Categories

Table 7 in the main manuscript shows that the prevalence of profitable rules differs markedly across trading-rule categories. This suggests that rule category is itself informative about expected performance. Because category is discrete rather than continuous, however, it cannot be incorporated directly as an input into the *mfFDR* procedure. To examine its informativeness, we therefore repeat the basket-of-currencies exercise separately within each rule category.

Specifically, for each category, we pool all rule–currency combinations belonging to that category and apply *mfFDR* to this restricted universe. This yields five category-specific sets of rules: filter rules, moving average rules, RSI rules, support-resistance rules, and channel breakout rules. The largest set is moving average rules, with the number of rule–currency combinations ranging from $9 \times 2,870 = 25,830$ to $30 \times 2,870 = 86,100$, while the smallest is RSI rules, ranging from $9 \times 600 = 5,400$ to $30 \times 600 = 18,000$. We then construct five *mfFDR*-based portfolios, one for each category, each controlling the *FDR* at 20%.

Table III reports the out-of-sample performance of these category-specific portfolios. Three findings stand out. First, portfolios based on filter rules are the most profitable before transaction costs, while portfolios based on moving average rules remain among the strongest performers after transaction costs. Second, the lower break-even costs for filter and support-resistance portfolios, relative to those for channel breakout and moving average portfolios, indicate that the former trade at higher frequencies and therefore incur higher transaction costs. This suggests that, when trading costs are high, investors should be cautious about allocating heavily to filter and support-resistance rules. Third, the RSI category

performs worst throughout the sample, generating negative profits in all subperiods as well as over the full sample. This indicates that even those RSI rules identified as outperforming in sample do not exhibit persistent out-of-sample profitability. Accordingly, investors constructing an $mfFDR$ portfolio over the full rule universe may wish to exclude RSI rules altogether.

Table III: Performance of the $mfFDR^+$ portfolios implemented on each category. The table shows the annualized Sharpe ratios (SR) and mean returns (before and after transaction cost) of portfolios generated by implementing the $mfFDR^+$ on each category of trading rule (on all currencies) to control the FDR at 20%. The last column is the related break-even point. In each panel, we present the mentioned metrics for the portfolio implemented on a category in the whole sample period (first row) and sub-samples (next five rows).

Period	Excess SR	Net SR	Excess Return (%)	Net Return (%)	Break-even (bps)
Panel A: Channel Breakout					
Whole Period	0.88	0.80	2.64	2.40	72
1973-1980	0.82	0.76	2.05	1.88	62
1981-1990	2.03	1.95	5.81	5.57	184
1991-2000	0.98	0.89	3.45	3.14	99
2001-2010	0.47	0.36	1.43	1.11	37
2011-2020	0.12	0.06	0.33	0.18	7
Panel B: Filter Rule					
Whole Period	1.26	0.91	4.41	3.19	23
1973-1980	1.77	1.53	5.62	4.86	36
1981-1990	2.17	1.81	7.14	5.97	41
1991-2000	0.98	0.66	4.90	3.27	27
2001-2010	1.02	0.44	2.95	1.28	14
2011-2020	0.67	0.37	1.71	0.95	7
Panel C: Moving Average					
Whole Period	1.01	0.92	3.80	3.47	71
1973-1980	1.39	1.32	4.49	4.24	89
1981-1990	1.99	1.91	7.42	7.12	156
1991-2000	0.92	0.83	4.23	3.82	89
2001-2010	0.65	0.53	2.43	1.97	41
2011-2020	0.18	0.11	0.57	0.37	9
Panel D: RSI					
Whole Period	-0.46	-0.96	-0.50	-1.04	-
1973-1980	-0.73	-1.16	-0.63	-0.99	-
1981-1990	-0.20	-0.70	-0.25	-0.84	-
1991-2000	-1.04	-1.67	-1.05	-1.70	-
2001-2010	-0.15	-0.73	-0.19	-0.90	-
2011-2020	-0.40	-0.74	-0.40	-0.74	-
Panel E: Support-resistance					
Whole Period	1.00	0.80	3.13	2.48	32
1973-1980	1.37	1.20	3.64	3.18	41
1981-1990	2.06	1.83	6.02	5.34	64
1991-2000	0.87	0.68	3.68	2.87	41
2001-2010	0.55	0.24	1.54	0.67	15
2011-2020	0.34	0.19	0.87	0.49	8

IJ. Macro Variables as Predictors

In this section, we revisit the regressions reported in Section 6.3 of the main manuscript using lagged macro variables in place of contemporaneous values. The results are reported in Table IV. Several lagged variables—the 3-month T-bill rate, the 5-year note yield, inflation, idiosyncratic volatility, and illiquidity—significantly predict the profitability of the selected outperforming rules. At the same time, the overall goodness of fit remains low, indicating that lagged macro variables explain only a modest share of the time-series variation in portfolio returns.

IK. Publications and the Returns on *mfFDR*-Based Portfolios

Figure IX plots the number of publications whose titles, abstracts, or keywords contain at least one of the following phrases: “algorithmic trading,” “artificial intelligence in trading,” “machine learning in trading,” “neural networks in trading,” and “high-frequency trading.” We also plot the out-of-sample Sharpe ratio of the *mfFDR*-based portfolio reported in Table 6.²⁷

The figure reveals a striking inverse relation between the profitability of technical trading and the growth of research activity related to advanced trading technologies. Prior to the 1990s, publication counts in these areas were low, while the out-of-sample Sharpe ratio of the selected technical rules was high. Beginning in the 1990s, the number of such publications rose markedly, coinciding with a sharp decline in the performance of the selected rules. This pattern continued in the subsequent two decades, during which publication activity expanded


²⁷The publication data were exported from Scopus in January 2024.

Table IV: Regression of $mfFDR^+$ -based portfolio returns on lagged macro variables. The table reports country fixed-effects regressions of monthly returns on the lagged macro variables. The dependent variable is the return on the $mfFDR^+$ -based portfolio, before and after transaction costs. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Excess return	Net return
3-month T-Bill	0.079*** (0.018)	0.079*** (0.017)
5-year note	-0.107*** (0.037)	-0.105*** (0.038)
10-year bond	0.032 (0.043)	0.026 (0.042)
Inflation	0.123** (0.050)	0.123** (0.050)
Dividend-to-price ratio	-0.030 (0.026)	-0.028 (0.026)
Price-to-earnings ratio	-0.00001 (0.00001)	-0.00001 (0.00001)
Stock market capitalization	-0.007* (0.004)	-0.007* (0.004)
Unemployment gap	-0.007 (0.011)	-0.008 (0.011)
Idiosyncratic volatility	-0.144*** (0.045)	-0.155*** (0.046)
Illiquidity	-0.020*** (0.007)	-0.022*** (0.007)
Aggregate volatility	0.003 (0.003)	0.003 (0.003)
Observations	6,588	6,588
Adjusted R^2	0.006	0.007

further and technical-trading profitability continued to weaken. Since the spread of high-frequency and machine-learning-based trading is also associated with deeper market liquidity and faster information processing, the figure is consistent with our earlier evidence on the role of illiquidity and with the broader interpretation that advances in trading technology have eroded the profitability of slower-moving technical strategies.

Figure IX: Profitability of technical rules and computational power. The figure plots the total number of Scopus publications related to advanced trading technologies (left axis) together with the out-of-sample Sharpe ratio of the $mfFDR$ -based portfolio (right axis).



II. Performance of *mfFDR*-Based Portfolios When Mean Excess Return Is the Testing Metric

As a robustness check, we repeat the main empirical analysis using mean excess return, rather than the Sharpe ratio, as the performance metric in hypothesis testing. That is, the object tested is now $\phi = \text{mean excess return}$. Table V reports the out-of-sample performance of the resulting *mfFDR*-based portfolios at the individual-currency level, both before and after transaction costs, for FDR targets of 10%, 20%, 30%, and 40%. Table VI reports the corresponding results when the method is implemented jointly across all currencies. Tables VII and VIII then examine how the returns on these portfolios are related to computational power, financial openness, and the broader set of macroeconomic variables considered in the main text.

Overall, the results are very similar to those obtained when the Sharpe ratio is used as the testing metric. The performance of the *mfFDR*-selected portfolios remains economically meaningful, is stable across alternative FDR targets, and exhibits the same broad relationships with market environment and macroeconomic conditions. These findings indicate that the main results are not driven by the particular choice of the Sharpe ratio as the testing criterion.

Table V: Performance of $mfFDR$ -Based Portfolios on Individual Currencies with Varying FDR Targets When Mean Return Is the Testing Metric. The table reports annualized Sharpe ratios of the $mfFDR$ -based portfolio for FDR targets $\tau \in \{0.1, 0.2, 0.3, 0.4\}$, computed from portfolio returns before and after transaction costs. The final row reports the average Sharpe ratio across the 30 currencies. Numbers in parentheses are p -values. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Countries	Before transaction costs				After transaction costs			
	$\tau = 0.1$	$\tau = 0.2$	$\tau = 0.3$	$\tau = 0.4$	$\tau = 0.1$	$\tau = 0.2$	$\tau = 0.3$	$\tau = 0.4$
Australia	0.20 (0.11)	0.20 (0.10)*	0.20 (0.11)	0.19 (0.11)	0.16 (0.19)	0.16 (0.20)	0.15 (0.21)	0.15 (0.22)
Canada	0.19 (0.15)	0.18 (0.17)	0.18 (0.17)	0.18 (0.17)	0.13 (0.31)	0.12 (0.39)	0.12 (0.41)	0.12 (0.38)
Germany/E.U.	0.45 (0.00)***	0.46 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.42 (0.00)***	0.42 (0.00)***	0.42 (0.00)***	0.41 (0.01)***
Japan	0.46 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.45 (0.00)***	0.43 (0.00)***	0.41 (0.01)***	0.42 (0.01)***	0.41 (0.01)***
New Zealand	0.33 (0.01)***	0.35 (0.00)***	0.35 (0.00)***	0.36 (0.00)***	0.27 (0.03)**	0.29 (0.02)**	0.30 (0.02)**	0.30 (0.02)**
Norway	0.21 (0.09)*	0.19 (0.14)	0.18 (0.14)	0.18 (0.14)	0.16 (0.20)	0.14 (0.27)	0.13 (0.29)	0.13 (0.29)
Sweden	0.37 (0.01)***	0.40 (0.01)***	0.40 (0.01)***	0.40 (0.01)***	0.32 (0.01)***	0.35 (0.01)***	0.35 (0.01)***	0.35 (0.01)***
Switzerland	0.30 (0.04)**	0.30 (0.02)**	0.30 (0.03)**	0.29 (0.03)**	0.26 (0.06)*	0.26 (0.06)*	0.26 (0.07)*	0.26 (0.07)*
U.K.	0.32 (0.02)**	0.30 (0.04)**	0.30 (0.05)**	0.30 (0.05)**	0.28 (0.05)**	0.27 (0.06)*	0.26 (0.07)*	0.26 (0.07)*
Argentina	0.22 (0.29)	0.24 (0.25)	0.24 (0.23)	0.24 (0.23)	0.18 (0.37)	0.19 (0.33)	0.20 (0.31)	0.20 (0.31)
Colombia	0.52 (0.01)***	0.57 (0.00)***	0.57 (0.00)***	0.57 (0.00)***	0.45 (0.03)**	0.50 (0.01)***	0.51 (0.01)***	0.50 (0.01)***
India	0.37 (0.02)**	0.35 (0.03)**	0.35 (0.03)**	0.35 (0.03)**	0.33 (0.05)**	0.31 (0.06)*	0.31 (0.06)*	0.30 (0.06)*
Indonesia	0.33 (0.02)**	0.34 (0.02)**	0.33 (0.02)**	0.33 (0.02)**	0.24 (0.10)*	0.24 (0.10)*	0.24 (0.10)*	0.24 (0.10)*
Israel	0.30 (0.04)**	0.31 (0.03)**	0.31 (0.03)**	0.31 (0.03)**	0.12 (0.38)	0.13 (0.32)	0.13 (0.33)	0.13 (0.33)
Philippines	0.66 (0.00)***	0.64 (0.00)***	0.63 (0.00)***	0.63 (0.00)***	0.49 (0.01)***	0.47 (0.01)***	0.46 (0.02)**	0.47 (0.02)**
Romania	0.16 (0.39)	0.20 (0.28)	0.20 (0.25)	0.20 (0.26)	0.05 (0.78)	0.08 (0.65)	0.09 (0.62)	0.09 (0.62)
Russia	0.40 (0.09)*	0.41 (0.08)*	0.42 (0.08)*	0.42 (0.08)*	0.38 (0.10)*	0.39 (0.08)*	0.40 (0.08)*	0.40 (0.08)*
Slovak	0.14 (0.42)	0.15 (0.40)	0.16 (0.37)	0.17 (0.36)	0.10 (0.58)	0.10 (0.56)	0.11 (0.54)	0.12 (0.50)
Brazil	0.35 (0.08)*	0.35 (0.09)*	0.34 (0.09)*	0.34 (0.09)*	0.32 (0.11)	0.31 (0.11)	0.31 (0.11)	0.31 (0.11)
Chile	0.44 (0.03)**	0.40 (0.04)**	0.40 (0.05)**	0.40 (0.05)**	0.36 (0.08)*	0.32 (0.11)	0.31 (0.12)	0.31 (0.12)
Czech	0.07 (0.64)	0.11 (0.51)	0.12 (0.47)	0.12 (0.46)	0.01 (0.88)	0.05 (0.71)	0.06 (0.67)	0.06 (0.64)
Hungary	-0.13 (0.49)	-0.12 (0.54)	-0.12 (0.55)	-0.12 (0.57)	-0.18 (0.35)	-0.18 (0.36)	-0.17 (0.37)	-0.17 (0.38)
Korea	0.32 (0.12)	0.30 (0.14)	0.29 (0.15)	0.29 (0.15)	0.26 (0.19)	0.24 (0.21)	0.24 (0.21)	0.24 (0.21)
Mexico	0.08 (0.55)	0.08 (0.55)	0.08 (0.57)	0.08 (0.56)	0.03 (0.78)	0.03 (0.80)	0.03 (0.81)	0.03 (0.80)
Poland	-0.05 (0.78)	-0.01 (0.98)	-0.01 (0.98)	0.00 (0.98)	-0.10 (0.59)	-0.05 (0.78)	-0.06 (0.78)	-0.05 (0.81)
Singapore	0.27 (0.05)**	0.32 (0.03)**	0.32 (0.03)**	0.32 (0.03)**	0.11 (0.54)	0.15 (0.35)	0.14 (0.37)	0.14 (0.37)
South Africa	0.17 (0.24)	0.21 (0.16)	0.23 (0.12)	0.24 (0.11)	0.05 (0.71)	0.09 (0.53)	0.11 (0.46)	0.11 (0.45)
Taiwan	0.78 (0.00)***	0.77 (0.00)***	0.77 (0.00)***	0.77 (0.00)***	0.70 (0.00)***	0.70 (0.00)***	0.69 (0.00)***	0.69 (0.00)***
Thailand	0.46 (0.02)**	0.47 (0.01)***	0.47 (0.01)***	0.47 (0.01)***	0.37 (0.05)**	0.39 (0.04)**	0.39 (0.04)**	0.38 (0.04)**
Turkey	0.48 (0.00)***	0.48 (0.00)***	0.48 (0.00)***	0.48 (0.00)***	0.42 (0.00)***	0.42 (0.00)***	0.42 (0.00)***	0.42 (0.00)***
Average	0.31	0.31	0.31	0.31	0.24	0.24	0.24	0.24

Table VI: Performance of $mfFDR^+$ -Based Portfolios with Alternative FDR Targets When Mean Return Is the Testing Metric. The table reports annualized Sharpe ratios, mean returns before and after transaction costs, and break-even transaction costs for portfolios formed by applying $mfFDR^+$ to the full cross-currency universe using mean return as the testing metric.

Period	Excess Sharpe Ratio	Net Sharpe Ratio	Excess Return	Net Return	Break-even
Panel A: FDR target of 10%					
Whole Period	1.00	0.90	4.48	4.04	65
1973–1980	1.43	1.34	5.01	4.69	78
1981–1990	2.09	1.99	8.36	7.96	141
1991–2000	0.85	0.76	5.27	4.71	89
2001–2010	0.63	0.47	2.51	1.89	34
2011–2020	0.35	0.28	1.37	1.08	11
Panel B: FDR target of 20%					
Whole Period	1.08	0.97	3.87	3.47	60
1973–1980	1.45	1.36	4.47	4.18	75
1981–1990	2.07	1.97	7.33	6.95	128
1991–2000	0.95	0.84	4.31	3.80	75
2001–2010	0.69	0.53	2.35	1.80	34
2011–2020	0.34	0.26	1.03	0.78	10
Panel C: FDR target of 30%					
Whole Period	1.12	1.00	3.55	3.17	57
1973–1980	1.47	1.38	4.24	3.97	75
1981–1990	2.02	1.91	6.71	6.34	119
1991–2000	1.02	0.89	3.85	3.37	68
2001–2010	0.72	0.55	2.24	1.73	34
2011–2020	0.34	0.25	0.88	0.65	9
Panel D: FDR target of 40%					
Whole Period	1.13	1.00	3.36	2.99	55
1973–1980	1.46	1.36	4.06	3.80	75
1981–1990	1.99	1.87	6.33	5.97	113
1991–2000	1.08	0.94	3.62	3.16	65
2001–2010	0.73	0.56	2.18	1.68	33
2011–2020	0.31	0.22	0.76	0.54	8

IM. Performance of Technical Trading from an Alpha Perspective

So far, we have selected portfolios of trading rules by testing whether either the Sharpe ratio or the mean return is equal to zero. In this section, we instead evaluate performance from an asset-pricing perspective by using alpha as the testing metric. Specifically, we

Table VII: Regressions of $mfFDR$ -Based Portfolio Returns When Mean Return Is the Testing Metric. The table reports country fixed-effects regressions of annual portfolio returns on *Publication*, *Openness*, and macroeconomic controls when mean return is used as the testing metric. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	<i>Dependent variable</i>					
	Excess return			Net return		
Publication	−0.0072** (0.0030)		−0.0069** (0.0032)	−0.0067** (0.0029)		−0.0064** (0.0032)
Openness		−0.0074** (0.0031)	−0.0064* (0.0033)		−0.0070** (0.0031)	−0.0062* (0.0032)
GDP growth	−0.0025 (0.0017)	−0.0024 (0.0017)	−0.0026 (0.0017)	−0.0023 (0.0016)	−0.0023 (0.0016)	−0.0024 (0.0016)
Stock market development	−0.0005*** (0.0001)	−0.0005*** (0.0001)	−0.0005*** (0.0001)	−0.0004*** (0.0001)	−0.0005*** (0.0001)	−0.0005*** (0.0001)
International trade	0.0001 (0.0002)	0.0002 (0.0003)	0.0001 (0.0002)	0.0001 (0.0002)	0.0001 (0.0002)	0.0001 (0.0002)
Unemployment rate	−0.0019 (0.0017)	−0.0014 (0.0015)	−0.0021 (0.0016)	−0.0024 (0.0016)	−0.0019 (0.0014)	−0.0025 (0.0016)
Observations	790	772	772	790	772	772
Adjusted R^2	0.0323	0.0302	0.0350	0.0255	0.0237	0.0275

estimate the following three-factor model, which is widely used in the FX literature:

$$r_{i,t} = \alpha_i + b_i RX_t + h_i HML_t + m_i MOM_t + \eta_{i,t}, \quad (\text{M.1})$$

where RX_t , HML_t , and MOM_t denote the daily returns on the dollar, carry, and momentum factors, respectively. The dollar factor, RX_t , is the cross-sectional average return across all currencies on day t . The carry factor, HML_t , captures the return to going long currencies with relatively high interest rates and short currencies with relatively low interest rates, based on lagged monthly interest-rate differentials. The momentum factor, MOM_t , captures the continuation in currency returns by going long recent winners and short recent losers. We construct this factor using a 25-day lag, which corresponds approximately to one month in our sample and is therefore consistent with the standard monthly implementation of FX

Table VIII: Regression of $mfFDR$ -Based Portfolio Returns on Macro Variables When Mean Return Is the Testing Metric. The table reports country fixed-effects regressions of monthly returns on macro variables when mean return is used as the testing metric. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses.

	Excess return	Net return
3-month T-Bill	0.040 (0.025)	0.041 (0.025)
5-year note	-0.028 (0.056)	-0.027 (0.056)
10-year bond	-0.0003 (0.037)	-0.006 (0.037)
Inflation	0.056 (0.049)	0.056 (0.049)
Dividend-to-price ratio	-0.046 (0.028)	-0.043 (0.028)
Price-to-earnings ratio	-0.00002** (0.00001)	-0.00002** (0.00001)
Stock market capitalization	-0.021*** (0.005)	-0.021*** (0.005)
Unemployment gap	0.007 (0.013)	0.007 (0.013)
Idiosyncratic volatility	-0.108* (0.066)	-0.121* (0.067)
Illiquidity	0.029*** (0.010)	0.027*** (0.010)
Aggregate volatility	0.010*** (0.003)	0.010*** (0.003)
Observations	6,609	6,609
R^2	0.032	0.030

momentum strategies.

We then construct the $mfFDR$ portfolio using the same rolling-window procedure as in the main manuscript, now applying $mfFDR^+$ to the pooled cross-currency universe while controlling the FDR at 20% and using alpha as the testing metric. The four covariates are ρ , α_{mk} , β_{mk} , and the R^2 from regression (M.1).²⁸

Table IX reports the resulting portfolio performance in terms of annualized Sharpe ratio,

²⁸Following Amihud and Goyenko (2013), we use the R^2 from the same factor model used to estimate alpha.

mean return, and three-factor alpha, both before and after transaction costs. The final column reports the break-even transaction cost in basis points.²⁹ Overall, the portfolio performs even better than in the Sharpe-ratio-based analysis. Over the full sample, it delivers an annualized Sharpe ratio of 1.12 before transaction costs and 1.01 after transaction costs, together with annualized alpha of 3.92% and 3.47%, respectively. Relative to Table 6 in the main text, the subperiod results are also generally stronger, with the exception of 2001–2010.

Table IX: Performance of $mfFDR^+$ Portfolios Based on Testing Alpha. The table reports annualized Sharpe ratios, mean returns, and three-factor alphas, before and after transaction costs, for portfolios formed by applying $mfFDR^+$ to the full cross-currency universe and using alpha as the testing metric. The final column reports the corresponding break-even transaction cost in basis points.

Period	Sharpe ratio		Return (%)		Alpha (%)		Break-even (bps)
	Excess	Net	Excess	Net	Excess	Net	
Whole Period	1.12	1.01	4.27	3.83	3.92	3.47	64
1973–1980	1.54	1.44	4.44	4.14	4.22	3.92	73
1981–1990	2.34	2.23	8.54	8.13	7.79	7.39	146
1991–2000	0.90	0.80	4.68	4.15	2.54	2.01	81
2001–2010	0.66	0.45	2.17	1.50	2.63	1.96	27
2011–2020	0.47	0.39	1.60	1.32	0.78	0.50	20

As in Sections 6.1–6.3, we then partition the selected strategies by currency and study whether the resulting portfolio returns are related to computational power, financial openness, and macroeconomic conditions. Tables X and XI report the corresponding country fixed-effects panel regressions at the annual and monthly frequencies, respectively.

The results broadly mirror those obtained when the Sharpe ratio is used as the testing metric. In particular, *Publication*, *Openness*, and stock market development remain statistically significant, and *Publication* emerges as the most robust explanatory variable. In addition, the unemployment rate is significant in several specifications. Monthly macro re-

²⁹Annualized alpha is obtained by multiplying daily alpha by 252, the conventional number of trading days in a year.

gressions yield results that are again broadly similar to those in the main text, although illiquidity is no longer consistently significant, even though its estimated coefficient remains positive.

Table X: Regression of *mfFDR*-Based Portfolio Returns on *Publication*, *Openness*, and Macro Factors When Alpha Is the Testing Metric. The table reports country fixed-effects regressions of annual portfolio returns on *Publication*, *Openness*, and macroeconomic controls. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	<i>Dependent variable</i>					
	Excess return			Net return		
Publication	−0.0054*** (0.0015)		−0.0049*** (0.0016)	−0.0047*** (0.0015)		−0.0042*** (0.0016)
Openness		−0.0078*** (0.0030)	−0.0071** (0.0030)		−0.0075** (0.0029)	−0.0070** (0.0030)
GDP growth	−0.0007 (0.0012)	−0.0007 (0.0012)	−0.0008 (0.0012)	−0.0005 (0.0013)	−0.0005 (0.0012)	−0.0006 (0.0012)
Stock market development	−0.0004*** (0.0001)	−0.0004*** (0.0001)	−0.0004*** (0.0001)	−0.0003*** (0.0001)	−0.0003** (0.0001)	−0.0003** (0.0001)
International trade	−0.0001 (0.0001)	−0.0001 (0.0001)	−0.0001 (0.0001)	−0.0001 (0.0001)	−0.0001 (0.0002)	−0.0001 (0.0002)
Unemployment rate	−0.0027* (0.0015)	−0.0023 (0.0014)	−0.0027* (0.0015)	−0.0032** (0.0014)	−0.0028** (0.0013)	−0.0032** (0.0014)
Observations	790	772	772	790	772	772
Adjusted R^2	0.0109	0.0098	0.0118	0.0045	0.0038	0.0049

Table XI: Regression of $mfFDR$ -Based Portfolio Returns on Macro Variables When Alpha Is the Testing Metric. The table reports country fixed-effects regressions of monthly portfolio returns on macro variables when alpha is used as the testing metric. Heteroskedasticity- and autocorrelation-consistent standard errors are reported in parentheses.

	Excess return	Net return
3-month T-Bill	0.027 (0.021)	0.026 (0.021)
5-year note	-0.006 (0.069)	0.001 (0.070)
10-year bond	-0.010 (0.053)	-0.019 (0.051)
Inflation	0.023 (0.052)	0.023 (0.053)
Dividend-to-price ratio	-0.023 (0.035)	-0.018 (0.035)
Price-to-earnings ratio	-0.00003*** (0.00001)	-0.00003*** (0.00001)
Stock market capitalization	-0.015*** (0.004)	-0.015*** (0.005)
Unemployment gap	-0.006 (0.012)	-0.006 (0.012)
Idiosyncratic volatility	-0.145* (0.077)	-0.167** (0.077)
Illiquidity	0.008* (0.004)	0.005 (0.004)
Aggregate volatility	0.013*** (0.003)	0.013*** (0.003)
Observations	6,609	6,609
Adjusted R^2	0.020	0.020