

CEIS Tor Vergata

RESEARCH PAPER SERIES

Vol. 24, Issue 4, No. 628 – October 2026

Investment Mandates and the Value of Distribution Timing in Currency Portfolios Choice

Mattia Alfero, Stefano Grassi and Francesco Violante

Investment Mandates and the Value of Distribution Timing in Currency Portfolios Choice

Mattia Alfero¹, Stefano Grassi² and Francesco Violante³

October 1, 2026

Abstract

The economic value of richer return forecasts depends on investors' ability to translate distributional information into portfolio decisions. We examine this relationship in currency markets, where observable interest-rate differentials discipline currency expected excess returns under the maintained no-predictability restriction for exchange-rate changes. We combine a multivariate model with time-varying volatility, skewness, and kurtosis (and associated co-moments) with Bayesian online Sequential Monte Carlo estimation, incorporating parameter uncertainty directly into dynamic portfolio choice. In an out-of-sample exercise spanning the period 2006-2025, we consider a US investor allocating across nine major currencies. We compare mean-variance allocations with rules that additionally incorporate skewness and kurtosis under four mandates governing gross exposure, net directional exposure, and short selling. The economic value of distributional information depends critically on the investment mandate through which it can be implemented. The largest and most statistically robust gains arise when investors can adjust total exposure, while smaller and more heterogeneous gains remain under more restrictive mandates. The benchmark-relative results remain economically meaningful after proportional transaction costs, although their magnitude depends on the investment constraints and trading intensity.

Keywords: Currency Carry Trade; Portfolio Choice; Distribution Timing; Higher-Order Moments; Bayesian Estimation.

JEL Classification: G11, F31, G15, G17, C58

¹University of Rome "Tor Vergata", Via Columbia 2, Rome, Italy - E-mail: mattia.alfero@students.uniroma2.eu

²University of Rome "Tor Vergata", Via Columbia 2, Rome, Italy - E-mail: stefano.grassi@uniroma2.it

³IESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Economie Management, F-59000 Lille, France, CREST and ILB - E-mail: f.violante@ieseg.fr.

The authors thank Fabio Troiani, Daniele Massacci, and Alexandre Rubesam for helpful comments and suggestions, as well as participants at the 2026 CAM-Risk Conference (Pavia), the 2025 International Joint Conference CFE-CMStatistics (London), the 2025 Annual Conference of the IAAE (Turin) and the 2025 ICEEE (Palermo). Stefano Grassi gratefully acknowledges financial support from the University of Rome "Tor Vergata" under "Bando Ricerca Scientifica di Ateneo 2024" and "Progetto di ricerca finanziato dal Dipartimento di Eccellenza 2023-2027". The funding organization had no role in the study design, data collection and analysis, decision to publish or preparation of the manuscript.

1 Introduction

Portfolio optimization is a central problem in financial economics and investment management, yet its out-of-sample benefits are often difficult to realize because optimal portfolio weights depend on estimates of expected returns and risk. [DeMiguel et al. \(2009\)](#) show that, across a range of equity datasets, estimation error can substantially erode the theoretical gains from optimized portfolios relative to simple rules such as the equally weighted portfolio. The difficulty is particularly acute for expected returns, whose estimation error can dominate the gains from exploiting cross-sectional differences in expected returns and covariances. Currency markets provide a notable exception. [Ackermann et al. \(2017\)](#) show that interest-rate differentials provide an observable component of expected currency returns and exploit this feature to construct mean-variance currency portfolios that substantially improve upon naive diversification. This makes currency markets a particularly useful setting for asking whether portfolio optimization can extract additional information from characteristics of the conditional return distribution beyond its mean and variance.

The importance of this setting is reinforced by the extensive evidence that currency carry returns reflect common risk exposures and can exhibit substantial crash risk ([Lustig et al., 2011](#); [Menkhoff et al., 2012](#); [Brunnermeier et al., 2008](#)). Exchange-rate returns exhibit time-varying volatility, asymmetry, and tail thickness, and these features can become particularly pronounced during periods of market stress ([Brunnermeier et al., 2008](#)). A growing literature shows that higher-order moments of asset returns are both economically relevant and time-varying ([Hansen, 1994](#); [Harvey and Siddique, 1999](#); [Jondeau and Rockinger, 2003](#)). This motivates studying whether information about the conditional distribution of returns, beyond its first two moments, generates additional economic value for portfolio choice. [Jondeau and Rockinger \(2012\)](#) use forecasts of moments up to the fourth in dynamic asset allocation, and document economically significant gains in international equity allocation. Related work shows that higher moments and their time variation can affect optimal international portfolio choice more broadly ([Guidolin and Timmermann, 2008](#); [Harvey et al., 2010](#); [Gao and Nardari, 2018](#); [Le, 2021](#)). The extent to which such distribution timing generates economic value in currency markets remains, however, largely unexplored.

Currency carry provides a particularly clean environment in which to study this question. Under the maintained short-horizon no-predictability restriction for exchange-rate changes, the conditional mean of a foreign currency excess return is pinned down by the forward-implied interest-rate differential. The investor therefore does not need to estimate the conditional mean from historical excess returns in the same way as in a conventional asset-allocation problem. Instead, the allocation problem can focus on whether time variation in conditional volatility,

skewness, kurtosis, and their associated co-moments contains information that improves portfolio decisions. This distinction is important because it separates the economic value of higher-order distributional information from the first-moment estimation problem that has been central to the criticism of portfolio optimization, see [Best and Grauer \(1991\)](#).

We study this question from the perspective of a CRRA carry-trade investor who allocates across nine foreign currencies while facing different investment mandates. We estimate a flexible conditional-distribution model that allows volatility, asymmetry, and tail thickness to evolve over time and use its predictive co-moments to construct mean-variance (MV) and mean-variance-skewness-kurtosis (MVSK) portfolios. The recursive estimation problem is handled with Sequential Monte Carlo (SMC) methods and an online update that permits the posterior distribution to evolve as the estimation window rolls forward. Parameter uncertainty is incorporated directly into the predictive moments entering the portfolio objective by averaging the conditional moments over the posterior distribution, rather than conditioning on a single parameter estimate. The purpose of this methodology is to obtain a real-time estimate of the conditional distribution rich enough to evaluate the economic value of distribution timing.

A central feature of the analysis is that we evaluate higher-order moments information across a spectrum of portfolio mandates. We consider four feasible sets that restrict the investor's ability to act on the conditional distribution: an unconstrained mandate, a unit-gross-exposure mandate, a dollar-neutral mandate, and a long-only mandate. The comparison holds fixed the information set and conditional mean within each mandate and changes only the information extracted from the predictive distribution. It therefore allows us to distinguish the value of higher-order moments information from the ability to adjust total exposure, net directional exposure, and short positions. This distinction is economically important because information about the conditional distribution can have little value for an investor who cannot adjust the portfolio along the margin through which that information is most useful.

The empirical results show that higher-order moments conditioning generates positive economic value, but that its magnitude and statistical strength depend critically on the investor's feasible set. Relative to MV, MVSK generates positive gains in the performance fee and in the reported risk-adjusted measures across all four mandates and both levels of risk aversion considered. The largest and most statistically robust gains occur under the unconstrained mandate, where the investor can adjust both the composition and the scale of currency exposure. The associated welfare gains are sizeable, amounting to roughly 113 basis points per year (low risk aversion) and 32 basis points per year (high risk aversion). These gains are accompanied by substantial endogenous leverage, with a median realized gross leverage of approximately 12 for the less risk-averse investor. Hence, the unconstrained results should be interpreted as a frictionless diagnostic, rather than representing a naturally implementable investment mandate. They

provide a benchmark against which the constrained mandates reveal how much of the value of distributional timing can be realized under progressively tighter investment restrictions.

When gross exposure is fixed, the gains are smaller but remain positive. Under the unit-gross-exposure mandate, the annualized performance fee is approximately 13 basis points at the lower risk-aversion level and 6 basis points at the higher level, with statistically detectable improvements concentrated at the lower risk-aversion level. The dollar-neutral mandate produces still smaller gains, with performance fees of approximately 3 basis points per year at both risk-aversion levels and little statistical evidence of incremental performance. The long-only mandate provides an important contrast. The performance fee is approximately 19 basis points at the lower risk-aversion level and 24 basis points at the higher level, with statistically significant improvements across the reported risk-adjusted measures at the higher risk-aversion level.

The realized distributions of portfolio returns provide a complementary perspective. MVSK does not uniformly produce higher skewness, lower kurtosis, or lower downside risk. Instead, the higher-order moments objective changes the joint distribution of portfolio payoffs and the resulting trade-off among expected return, dispersion, asymmetry, and tail thickness. In the unconstrained case, for example, MVSK raises realized mean return and skewness but also increases volatility and kurtosis. Under fixed gross exposure, the changes are smaller and more heterogeneous. This pattern is consistent with the portfolio objective, which does not maximize skewness or minimize kurtosis separately, but evaluates their joint contribution to expected CRRA utility. The economic value of higher-order information therefore arises from its use in portfolio choice rather than from a mechanical improvement in every individual distributional statistic.

The results remain economically relevant when the strategies are compared with conventional currency and asset-class benchmarks. MVSK improves upon passive equal weighting in the long-only setting and is competitive with cross-sectional carry benchmarks under the other constrained mandates. The dollar-neutral portfolios in particular achieve Sharpe ratios above several traditional asset-class benchmarks, while the unconstrained portfolios exhibit the strongest realized risk-adjusted performance. At the same time, the correlation analysis shows that optimized currency portfolios have markedly different relationships with traditional asset-class exposures across investment mandates. These comparisons place the magnitude of the currency results in a broader investment context and indicate that the optimized currency strategies generate return dynamics that differ materially from standard asset-class exposures.

Finally, the economic gains are examined after transaction costs. Higher-order moment conditioning does not materially increase turnover relative to the MV benchmark, but implementation costs depend importantly on both trading intensity and the scale of outstanding currency positions. The MVSK advantage survives the baseline transaction-cost adjustment for the unconstrained, dollar-neutral, and long-only mandates at both levels of risk aversion. The gross-only

mandate is more sensitive to implementation costs. The small gross advantage at lower risk aversion is largely absorbed by costs, and the net ranking reverses at higher risk aversion. Robustness checks with higher transaction costs and alternative gross-exposure bounds leave the main comparisons qualitatively unchanged.

This paper contributes to three strands of literature. First, it contributes to the literature on economic evaluation of exchange-rate predictability and currency portfolio choice (Della Corte et al., 2009; Ackermann et al., 2017). We extend the insight that the observable carry component can mitigate first-moment estimation error by asking whether the remaining time variation in the conditional distribution contains additional information for portfolio choice. Second, it contributes to the literature on higher-order moments portfolio allocation and distribution timing (Jondeau and Rockinger, 2006, 2012; Guidolin and Timmermann, 2008; Harvey et al., 2010; Gao and Nardari, 2018; Le, 2021). We bring the distribution-timing framework to currency allocation, where the conditional mean has a particularly transparent economic interpretation, and show that the value of higher-order information depends on the investor’s portfolio constraints. Third, the paper contributes to the application of Bayesian estimation methods to recursive portfolio choice by using an online Sequential Monte Carlo framework to estimate a high-dimensional conditional distribution while retaining parameter uncertainty in the investment decision.

The key implication of our analysis is that higher-order moments acquire economic relevance through their role in portfolio decisions. The economic value of distributional timing is jointly determined by the information contained in the conditional distribution and by the investment mandate through which that information can be implemented. This perspective links the statistical problem of forecasting higher-order moments to the managerial problem of determining when and how that information can affect portfolio exposure. The results are therefore particularly relevant for investors and portfolio managers who actively allocate across currencies and can adjust the scale or composition of their currency exposure.

The remainder of the paper proceeds as follows. Section 2 introduces the carry-trade investor problem, the portfolio construction framework, and the performance criteria. Section 3 describes the conditional-distribution model, Bayesian estimation, and online Sequential Monte Carlo implementation. Section 4 reports the empirical evaluation, including the portfolio comparisons, benchmark analysis, and transaction-cost results. Section 5 concludes. Detailed derivations, estimation procedures, robustness analyses, and additional results are reported in the Appendices.

2 The Carry-Trade Investor’s Portfolio Choice Problem

This section formalizes the portfolio problem faced by a domestic investor allocating wealth across foreign currency positions. A key feature is that currency excess returns combine predetermined

interest-rate differentials, which are observed prior to portfolio selection, and random exchange-rate movements. Currency carry therefore provides a particularly clean setting for studying the value of distributional information in portfolio choice. Under the maintained no-predictability assumption, the conditional mean is determined by an observable interest-rate differential, leaving conditional risk and higher-order distributional characteristics as the main sources of uncertainty relevant for allocation.

Section 2.1 derives currency excess returns and shows how the same payoff can be implemented synthetically with forward contracts. Section 2.2 maps the predictive distribution of currency returns into a CRRA portfolio problem and defines the MV and MVSK allocation rules. Section 2.3 evaluates the performance of the resulting strategies.

2.1 Currency Carry and Exchange Rate Returns Decomposition

Let d denote the domestic currency and let $i = 1, \dots, N$ index foreign currencies. The spot exchange rate $S_{i,t}$ is quoted as units of domestic currency per unit of foreign currency i . A long carry position in foreign currency i borrows one unit of domestic currency at the one-period risk-free interest rate r_t^d , converts the proceeds into $1/S_{i,t}$ units of foreign currency, invests at the corresponding foreign risk-free rate, and converts the payoff $(1 + r_t^i)/S_{i,t}$ back into domestic currency at $t + 1$. This generates the one-period currency excess return

$$y_{t+1}^i = \frac{S_{i,t+1}}{S_{i,t}}(1 + r_t^i) - (1 + r_t^d). \quad (1)$$

In practice, the payoff in equation (1) can be replicated using FX swaps or forward contracts. Let $F_{i,t}$ denote the one-period forward exchange rate. Covered interest parity requires

$$F_{i,t} = S_{i,t} \frac{1 + r_t^d}{1 + r_t^i}. \quad (2)$$

Combining (1) and (2), the currency excess return can be equivalently expressed as

$$y_{t+1}^i = (1 + r_t^d) \left(\frac{S_{i,t+1}}{F_{i,t}} - 1 \right),$$

which is the payoff from a forward position scaled by the domestic gross risk-free rate. This synthetic representation of the long carry position is useful for empirical analysis because it corresponds closely to the implementation of institutional currency positions through FX forwards and facilitates the transaction-cost calculation discussed in Section 4.4.

Define the one-period spot exchange-rate return as

$$x_{t+1}^i = \frac{S_{i,t+1}}{S_{i,t}} - 1. \quad (3)$$

Equation (1) can then be rearranged as

$$y_{t+1}^i = (r_t^i - r_t^d) + (1 + r_t^i)x_{t+1}^i. \quad (4)$$

Equation (4) separates the payoff into two distinct components: a carry component (predetermined interest-rate differential) and a currency component (exchange-rate change scaled by the foreign gross return).

Let $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathcal{F}_t]$ be the expectation conditional to the information set, $\mathcal{F}_t$, representing all information available at time t . Motivated by the limited evidence of short-horizon predictability of exchange-rate changes documented in the literature, see [Meese and Rogoff \(1983\)](#), [Della Corte et al. \(2009\)](#), and [Rossi \(2013\)](#), following [Ackermann et al. \(2017\)](#), we adopt the benchmark restriction $\mathbb{E}_t[x_{t+1}^i] = \mathbf{0}$. Thus, the conditional expected payoff of the carry position in foreign currency i is

$$\mathbb{E}_t[y_{t+1}^i] = (r_t^i - r_t^d),$$

that is the excess return of the foreign investment relative to the domestic funding asset. Thus, the interest-rate differential provides an observable component of expected returns, allowing the portfolio problem to focus on the estimation of conditional risk and higher-order distributional characteristics.

This feature makes currency portfolio choice particularly well suited to studying the value of distributional timing. Unlike conventional asset-allocation problems in which the conditional mean must itself be estimated, and thus estimation uncertainty can severely affect optimal portfolio decisions ([DeMiguel et al., 2009](#)), the carry component provides an observable benchmark for expected currency excess returns. [Ackermann et al. \(2017\)](#) exploit this feature to study MV currency portfolio allocation. We extend this perspective by asking whether time variation in conditional risk and higher-order moments generates economically valuable portfolio decisions beyond those based on conditional means and variances, and whether such gains depend on the investor's leverage, directional, and short-selling constraints.

2.2 Portfolio Optimization with higher-order moments under Alternative Investment Mandates

Let $\mathbf{w}_t = (w_{1,t}, \dots, w_{N,t})'$ denote the vector of currency portfolio weights chosen at time t . The portfolio excess return is

$$y_{t+1}^p = \mathbf{w}_t' \mathbf{y}_{t+1},$$

where

$$\mathbf{y}_{t+1} = \mathbf{d}_t + \mathbf{A}_t \mathbf{x}_{t+1}, \quad \mathbf{d}_t = (r_t^1 - r_t^d, \dots, r_t^N - r_t^d)', \quad \mathbf{A}_t = \text{diag}(1 + r_t^1, \dots, 1 + r_t^N).$$

The investor has CRRA preferences over end-of-period wealth,

$$U(W_{t+1}) = \frac{W_{t+1}^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \quad \gamma \neq 1, \quad (5)$$

where γ denotes relative risk aversion. Normalizing current wealth to $W_t = 1$, end-of-period wealth is

$$W_{t+1} = 1 + y_{t+1}^p.$$

The CRRA specification requires strictly positive end-of-period wealth, $W_{t+1} > 0$. Hence, admissible portfolio choices must ensure that the resulting wealth remains positive over the relevant support of the predictive distribution. This domain restriction becomes particularly important when portfolio leverage is unconstrained.

At each rebalancing date, the investor chooses $\mathbf{w}_t$ using the predictive distribution of $\mathbf{y}_{t+1}$. Because the predictive distribution of currency returns is generally non-Gaussian and expected CRRA utility does not admit a convenient closed-form representation under our setting, we approximate expected utility using a fourth-order Taylor expansion around current wealth

$$\mathbb{E}_t[U(W_{t+1})] \approx \phi_0 + \phi_1 m_{p,t+1}^{(1)} + \phi_2 m_{p,t+1}^{(2)} + \phi_3 m_{p,t+1}^{(3)} + \phi_4 m_{p,t+1}^{(4)}, \quad (6)$$

where $m_{p,t+1}^{(k)} = \mathbb{E}_t[(y_{t+1}^p)^k]$ is the k -th conditional noncentral moment of the portfolio excess return. The fourth-order objective is used as a local polynomial preference approximation, see [Jondeau and Rockinger \(2006\)](#) and [Harvey et al. \(2010\)](#), among others. Conditions for convergence of the Taylor expansion to the expected utility and the order of truncation are extensively

discussed in [Jondeau and Rockinger \(2006\)](#) and references therein. For CRRA utility,

$$\phi_0 = \frac{1}{1-\gamma}, \quad \phi_1 = 1, \quad \phi_2 = -\frac{\gamma}{2}, \quad \phi_3 = \frac{\gamma(\gamma+1)}{6}, \quad \phi_4 = -\frac{\gamma(\gamma+1)(\gamma+2)}{24}.$$

The signs of these coefficients imply that the investor values expected returns and positive skewness, while penalizing dispersion and extreme realizations.

The MVSK allocation rule solves the fourth-order problem

$$\mathbf{w}_t^{MVSK} = \arg \max_{\mathbf{w}_t \in \mathcal{W}} \left\{ \phi_1 m_{p,t+1}^{(1)} + \phi_2 m_{p,t+1}^{(2)} + \phi_3 m_{p,t+1}^{(3)} + \phi_4 m_{p,t+1}^{(4)} \right\}, \quad (7)$$

where $\mathcal{W}$ is the feasible set and the constant term ϕ_0 is omitted because it does not affect portfolio choice. The MV benchmark is obtained by truncating the same expected-utility approximation at second order, which gives

$$\mathbf{w}_t^{MV} = \arg \max_{\mathbf{w}_t \in \mathcal{W}} \left\{ \phi_1 m_{p,t+1}^{(1)} + \phi_2 m_{p,t+1}^{(2)} \right\}. \quad (8)$$

The comparison between equations (7) and (8) therefore constitutes a controlled experiment in the value of distributional information. Conditional means and information sets are held fixed, and the same investment mandate is imposed on both rules, while the MVSK investor additionally exploits time variation in conditional skewness, kurtosis, and their associated co-moments.

To identify the channels through which distributional information affects portfolio choice, we vary three economically distinct margins: the scale of total exposure, net directional exposure, and short-selling. This produces four mandates, each corresponding to an economically distinct feasible set. Since the same feasible set is imposed on both the MV and MVSK allocation rules, any within-set performance difference is attributable to the additional distributional information contained in higher-order moments.

The first set is unconstrained,

$$\mathcal{W}^U = \mathbb{R}^N,$$

and serves as the frictionless benchmark. More precisely, $\mathbb{R}^N$ characterizes the absence of explicit portfolio restrictions on leverage, net exposure, and short selling, while admissibility is governed by the positive-wealth condition required by CRRA utility. Following the classical portfolio-choice framework of [Markowitz \(1959\)](#), this benchmark represents the allocation that would be chosen in the absence of implementation restrictions. Under a second-order approximation, for example, the unconstrained solution corresponds to the familiar tangency portfolio.

We interpret the unconstrained solution as a frictionless diagnostic of the exposure implied

by the estimated conditional distribution. The constrained mandates then provide economically interpretable alternatives that impose explicit limits on total exposure, net directional exposure, and short selling, and can therefore be interpreted as restrictions imposed on the frictionless optimum rather than as alternative investment opportunities.

This benchmark is also economically relevant in currency markets because institutional investors obtain leveraged exposure through derivative instruments such as FX forwards and swaps, which are central to the implementation of currency positions in institutional foreign exchange markets, see [Lyons \(2001\)](#). Unlike retail investors, professional currency managers face leverage limits determined by funding conditions, collateral requirements, and internal risk controls rather than by mechanical regulatory caps. Hence, the unconstrained solution should be interpreted as a frictionless benchmark for desired exposure, rather than as an implementable mandate. Its realized leverage provides a diagnostic of how strongly the estimated conditional distribution would induce the investor to scale the position in the absence of institutional constraints.

The second set imposes a unit gross-exposure constraint,

$$\mathcal{W}^G = \left\{ \mathbf{w} \in \mathbb{R}^N : \sum_{i=1}^N |w_i| = 1 \right\},$$

where total notional exposure is fixed at wealth, ruling out mechanical leverage while leaving the net directional position unconstrained. This mandate corresponds to a diversified currency portfolio with fixed unit gross exposure, as in [Ackermann et al. \(2017\)](#).

The third set is dollar-neutral,

$$\mathcal{W}^{DN} = \left\{ \mathbf{w} \in \mathbb{R}^N : \sum_{i=1}^N |w_i| = 1, \sum_{i=1}^N w_i = 0 \right\},$$

which offsets long and short positions in net terms. By removing directional exposure to the domestic currency, this mandate isolates the cross-sectional component of the carry trade and provides a direct comparison with conventional currency carry portfolios.

The fourth set is long-only,

$$\mathcal{W}^L = \left\{ \mathbf{w} \in \mathbb{R}^N : \sum_{i=1}^N |w_i| = 1, \sum_{i=1}^N w_i = 1 \right\}.$$

The two restrictions jointly imply a fully invested, long-only allocation. This mandate captures an investor unable to use leverage or short selling and is consistent with the constraints underlying the conventional 1/N benchmark.

Together, these four mandates trace the channels through which higher-order moments information can generate economic value. Exposure scaling in response to conditional risk and

higher-order moments is available only under $\mathcal{W}^U$, whereas the remaining mandates fix total notional exposure and differ in the permitted cross-sectional reallocations. Specifically, $\mathcal{W}^G$ allows unconstrained net currency exposure, $\mathcal{W}^{DN}$ isolates relative-value currency positions, and $\mathcal{W}^L$ additionally rules out short positions. Comparing performance across feasible sets therefore allows us to distinguish whether gains from conditioning on higher-order moments arise from leverage adjustment, cross-sectional positioning, or both.

2.3 Economic Evaluation

We evaluate the economic value of conditioning portfolio decisions on the predictive distribution using four groups of performance measures, summarized in Table 1. The first group characterizes the realized distribution of portfolio excess returns. The second evaluates risk-adjusted performance and welfare. The third focuses on downside and tail risk. The fourth captures implementation characteristics.

First, we characterize the realized distribution of portfolio excess returns using the mean excess return (μ_p), volatility (σ_p), semi-standard deviation (SSD_p), skewness (SK_p), and kurtosis (KU_p). These statistics show whether conditioning on higher-order moments changes the shape of the realized return distribution rather than merely its first two moments. In particular, the comparison between MV and MVSK strategies reveals whether incorporating conditional skewness, kurtosis, and the associated co-moments translates into systematic differences in realized distributional characteristics.

Second, we evaluate economic performance using the Sharpe ratio (SR; [Sharpe, 1966](#)), the adjusted Sharpe ratio (aSR; [Pezier and White, 2006](#); see also [Bacon, 2023](#)), and the performance fee (Φ ; [Fleming et al., 2001](#)). The SR summarizes mean return relative to total volatility, while the aSR accounts additionally for skewness and excess kurtosis. The performance fee is our primary welfare criterion because it is derived from the same CRRA preferences that determine the portfolio allocation. Specifically, Φ translates the utility difference between two strategies into a monetary measure: the constant per-period fee that an investor would pay to switch from the MV benchmark to the competing MVSK strategy while remaining indifferent in expected utility.

Third, we evaluate downside and tail risk using the modified Sharpe ratio (mSR; [Favre and Galeano, 2002](#); [Gregoriou and Gueyie, 2003](#)), the Sortino ratio ($\mathcal{S}$ -ratio; [van der Meer et al., 2001](#)), and the Omega ratio (Ω -ratio; [Keating and Shadwick, 2002](#)). These measures are particularly relevant for currency portfolios because carry-trade returns can exhibit pronounced asymmetry and left-tail risk ([Brunnermeier et al., 2008](#)). They provide complementary information about whether the welfare gains associated with higher-order moments conditioning reflect improved left-tail protection, reduced downside dispersion, or broader changes in the return distribution.

Table 1: Performance measures

Measure	Economic interpretation	Definition
Distributional characteristics		
Mean	Average payoff	$\mu_p = T^{-1} \sum_{t=1}^T y_t^p$
Volatility	Total risk	$\sigma_p = \left[T^{-1} \sum_{t=1}^T (y_t^p - \mu_p)^2 \right]^{1/2}$
Semi-Standard Deviation	Downside dispersion	$SSD_p = \left[T^{-1} \sum_{t=1}^T \min(y_t^p - \mu_p, 0)^2 \right]^{1/2}$
Skewness	Distributional asymmetry	$SK_p = T^{-1} \sum_{t=1}^T ((y_t^p - \mu_p)/\sigma_p)^3$
Kurtosis	Tail thickness	$KU_p = T^{-1} \sum_{t=1}^T ((y_t^p - \mu_p)/\sigma_p)^4$
Economic/financial performance diagnostics		
Sharpe Ratio	Return relative to total risk, Sharpe (1966)	$SR_p = \frac{\mu_p}{\sigma_p}$
adjusted Sharpe Ratio	Non-Gaussian risk-adjusted return, (Pezier and White, 2006 ; Bacon, 2023)	$aSR_p = SR_p \left[1 + \frac{SK_p}{6} SR_p - \frac{KU_p - 3}{24} SR_p^2 \right]$
Performance fee	Utility-based welfare gain, Fleming et al. (2001)	$T^{-1} \sum_{t=1}^T \frac{(1 + y_t^p - \Phi)^{1-\gamma}}{1-\gamma} = \bar{U}_b$
Downside/tail risk measures		
modified Sharpe Ratio	Tail-risk-adjusted return, Favre and Galeano (2002) ; Gregoriou and Gueyie (2003)	$mSR_{q,p} = \frac{\mu_p}{mVaR_{q,p}}$
Omega ratio	Distribution-wide performance, Keating and Shadwick (2002)	$\Omega_p\text{-ratio} = \frac{\sum_{t=1}^T \max(y_t^p - \tau, 0)}{\sum_{t=1}^T \max(\tau - y_t^p, 0)}$
Sortino ratio	Downside-risk-adjusted performance, van der Meer et al. (2001)	$\mathcal{S}_p\text{-ratio} = \frac{\mu_p - \tau}{\left[T^{-1} \sum_{t=1}^T \min(y_t^p - \tau, 0)^2 \right]^{1/2}}$
Implementation		
Gross portfolio turnover	Trading intensity	$TO = (T-1)^{-1} \sum_{t=1}^{T-1} \sum_{i=1}^N w_{i,t+1} - w_{i,t} $
Normalized Herfindahl–Hirschman index	Portfolio concentration, Hirschman (1980) ; Herfindahl (1997)	$HHI = T^{-1} \sum_{t=1}^T \sum_{i=1}^N \bar{w}_{i,t} ^2$

Notes: The table reports the measures used in the empirical analysis. Compact formulas are shown for reference; full definitions, annualization conventions, and auxiliary quantities are reported in Appendix A. Kurtosis is the fourth standardized central moment, so the Gaussian benchmark is three. Normalized portfolio weights are denoted by $\bar{w}_{i,t}$. The modified Value-at-Risk, $mVaR_{q,p}$, is computed using the Cornish–Fisher expansion. The level q in the $mVaR_{q,p}$ and $mSR_{q,p}$ is set to 5%. The target τ in the Ω -ratio and $\mathcal{S}$ -ratio is set to 0.

In particular, the Ω -ratio provides a distribution-wide assessment that complements the moment-based portfolio objective.

Fourth, we assess implementation characteristics using gross portfolio turnover (TO) and portfolio concentration. Gross turnover captures the trading intensity induced by dynamic portfolio rebalancing and is used in the transaction-cost analysis of Section 4.4. Portfolio concentration is

measured using the Normalized Herfindahl–Hirschman Index (HHI) (Hirschman, 1980; Herfindahl, 1997) evaluated conditional on each mandate’s exposure scale. Because the unconstrained investor $\mathcal{W}^U$ can take leveraged positions, and squared weights are mechanically driven by leverage, the HHI is computed from normalized portfolio weights, that is $\bar{w}_{i,t} = w_{i,t}/G_t$, $i = 1, \dots, N$, and $\sum_{i=1}^N |\bar{w}_{i,t}| = 1$, where $G_t = \sum_{j=1}^N |w_{j,t}|$ denotes the gross exposure of the portfolio at time t . Since the other mandates are constrained to unit gross exposure, i.e. $G_t = 1$, they are unaffected by the normalization. The HHI, therefore, measures the concentration of gross currency exposure independently of its overall scale. Gross exposure G_t captures the extent of leverage, whereas the HHI captures how that exposure is distributed across currencies. This distinction is particularly important for the unconstrained investor, for whom the scale of exposure can vary endogenously with the conditional distribution.

All measures are computed from realized out-of-sample portfolio excess returns. Unless otherwise stated, the same realized returns and evaluation periods are used across competing allocation rules. Formal definitions, annualization conventions, and auxiliary quantities are reported in Appendix A.

3 Estimating the Conditional Distribution of Currency Returns

The portfolio problem in Section 2 requires the predictive distribution of currency excess returns. Since the conditional mean is determined by the observed interest-rate differentials, the modeling effort concentrates on conditional volatility, skewness, and tail thickness, and then maps the resulting predictive co-moments into the MV and MVSK allocation rules.

Section 3.1 specifies the return model. Section 3.2 describes the Bayesian SMC estimator and the online update used in the recursive portfolio exercise. Section 3.3 explains how posterior parameter uncertainty is incorporated into the predictive moments used by the investor.

3.1 Conditional Distribution Model

We construct an autoregressive conditional density model that captures the asymmetric and fat-tailed characteristics of returns, by combining a multivariate skewed Student- t (MSk- t) distribution and dynamic covariance, shape, and skew parameters. This modeling approach was pioneered by Hansen (1994) and Harvey and Siddique (1999), and extended to time-varying higher-order co-moments for multivariate variables by Sahu et al. (2003), Jondeau and Rockinger (2009), Jondeau and Rockinger (2012), and Ghalanos et al. (2015).

Let $\mathbf{x}_{t+1} = (x_{t+1}^1, \dots, x_{t+1}^N)'$ denote the vector of exchange-rate net returns defined in equation

(3). We construct the model as follows

$$\mathbf{x}_{t+1} = \mathbf{H}_{t+1}^{1/2} \mathbf{z}_{t+1}, \quad \mathbf{z}_{t+1} \sim \text{MSk-}t(\boldsymbol{\eta}_{t+1}, \boldsymbol{\lambda}_{t+1}),$$

where $\mathbf{H}_{t+1}^{1/2}$ is a suitable matrix square root factorization of the conditional covariance matrix $\mathbf{H}_{t+1} = \mathbb{E}_t[\mathbf{x}_{t+1}\mathbf{x}_{t+1}']$, $\mathbf{z}_{t+1}$ is a vector of independent, zero-mean and unit-variance innovations, $\boldsymbol{\eta}_{t+1}$ and $\boldsymbol{\lambda}_{t+1}$ are the vectors of tail-thickness and asymmetry parameters, respectively. Following Hansen (1994), Sahu et al. (2003), and Harvey and Siddique (1999), we assume that the components of $\mathbf{z}_{t+1}$ are cross-sectionally independent so that the dependence structure between the elements of $\mathbf{x}_{t+1}$ is fully captured by $\mathbf{H}_{t+1}$ alone.

The conditional covariance matrix of $\mathbf{x}_{t+1}$ is decomposed as

$$\mathbf{H}_{t+1} = (\mathbf{I} \odot \mathbf{H}_{t+1})^{1/2} \boldsymbol{\Gamma}_{t+1} (\mathbf{I} \odot \mathbf{H}_{t+1})^{1/2},$$

where $\odot$ represents the Hadamard product, such that the matrix $(\mathbf{I} \odot \mathbf{H}_{t+1})$ is diagonal and contains the conditional marginal variances, $h_{ii,t+1} = \mathbb{E}_t[x_{i,t+1}^2]$, on the main diagonal, and $\boldsymbol{\Gamma}_{t+1}$ is the conditional correlation matrix. We now specify the dynamics of the latent processes. Each marginal variance $h_{ii,t+1}$ follows a GARCH(1, 1) recursion (Bollerslev, 1986)

$$h_{ii,t+1} = \omega_i + \alpha_i x_{i,t}^2 + \beta_i h_{ii,t}, \quad i = 1, \dots, N.$$

The conditional correlation matrix $\boldsymbol{\Gamma}_{t+1}$ follows the DCC recursion of Engle (2002)

$$\begin{aligned} \boldsymbol{\Gamma}_{t+1} &= (\mathbf{I} \odot \mathbf{Q}_{t+1})^{-1/2} \mathbf{Q}_{t+1} (\mathbf{I} \odot \mathbf{Q}_{t+1})^{-1/2}, \\ \mathbf{Q}_{t+1} &= (1 - \delta_1 - \delta_2) \bar{\mathbf{Q}} + \delta_1 \mathbf{u}_t \mathbf{u}_t' + \delta_2 \mathbf{Q}_t, \end{aligned}$$

where $\mathbf{u}_t = (\mathbf{I} \odot \mathbf{H}_t)^{-1/2} \mathbf{x}_t$, and $\bar{\mathbf{Q}} = \mathbb{E}[\mathbf{u}_t \mathbf{u}_t']$. Parameter restrictions ensuring symmetry and positive definiteness of $\mathbf{H}_{t+1}$, as well as weak stationarity of $\mathbf{x}_t$, are given in Bollerslev (1986) and Engle (2002). In the empirical application, $\bar{\mathbf{Q}}$ is estimated using a sample estimator.

We model latent shape parameters for $i = 1, \dots, N$ as

$$\eta_{i,t+1} = \underline{\eta} + \frac{\bar{\eta} - \underline{\eta}}{1 + \exp(-\tilde{\eta}_{i,t+1})}, \quad \lambda_{i,t+1} = -1 + \frac{2}{1 + \exp(-\tilde{\lambda}_{i,t+1})}, \quad (9)$$

where the auxiliary processes $\tilde{\eta}_{i,t+1}$ and $\tilde{\lambda}_{i,t+1}$ follow asymmetric GARCH-type recursions driven

by the standardized innovation $z_{i,t}$, that is

$$\begin{aligned}\tilde{\eta}_{i,t+1} &= \gamma_{i,0} + \gamma_{i,1}^+ |z_{i,t}| \mathbf{1}_{\{z_{i,t} > 0\}} + \gamma_{i,1}^- |z_{i,t}| \mathbf{1}_{\{z_{i,t} < 0\}} + \gamma_{i,2} \tilde{\eta}_{i,t}, \\ \tilde{\lambda}_{i,t+1} &= \kappa_{i,0} + \kappa_{i,1}^+ z_{i,t} \mathbf{1}_{\{z_{i,t} > 0\}} + \kappa_{i,1}^- z_{i,t} \mathbf{1}_{\{z_{i,t} < 0\}} + \kappa_{i,2} \tilde{\lambda}_{i,t},\end{aligned}$$

with stationarity conditions $|\gamma_{i,2}| < 1$ and $|\kappa_{i,2}| < 1$. The logistic maps in equation (9) restrict $\lambda_{i,t+1} \in (-1, 1)$, which ensures a valid density, and $\eta_{i,t+1} \in (\underline{\eta}, \bar{\eta})$, with bounds chosen to reflect the region of interest. More precisely, the lower bound ensures existence of moments least up to order $\underline{\eta}$, while bounding $\eta_{i,t+1}$ from above substantially improves identifiability. We restrict $\eta_{i,t+1} \in (4, 30)$, see Hansen (1994) for additional details.

Finally, the MSk- t joint probability density function of $\mathbf{z}_{t+1}$ is

$$g(\mathbf{z}_{t+1} | \boldsymbol{\eta}_{t+1}, \boldsymbol{\lambda}_{t+1}) = \prod_{i=1}^N b_{i,t+1} c_{i,t+1} \left[1 + \frac{1}{\eta_{i,t+1} - 2} \left(\frac{b_{i,t+1} z_{i,t+1} + a_{i,t+1}}{\delta_{i,t+1}} \right)^2 \right]^{-\frac{(\eta_{i,t+1}+1)}{2}},$$

where the auxiliary parameters

$$a_{i,t+1} = 4\lambda_{i,t+1} c_{i,t+1} \frac{\eta_{i,t+1} - 2}{\eta_{i,t+1} - 1}, \quad b_{i,t+1} = (1 + 3\lambda_{i,t+1}^2 - a_{i,t+1}^2)^{1/2},$$

ensure $\mathbf{z}_{t+1}$ is centered and unit variance,

$$c_{i,t+1} = \frac{\Gamma[(\eta_{i,t+1} + 1)/2]}{\sqrt{\pi(\eta_{i,t+1} - 2)} \Gamma(\eta_{i,t+1}/2)}, \quad i = 1, \dots, N,$$

and the piecewise scale factor is

$$\delta_{i,t+1} = \begin{cases} 1 - \lambda_{i,t+1}, & z_{i,t+1} < -a_{i,t+1}/b_{i,t+1}, \\ 1 + \lambda_{i,t+1}, & \text{otherwise.} \end{cases}$$

The joint probability density function of $\mathbf{z}_{t+1}$ is related to the density of $\mathbf{x}_{t+1}$ by

$$g(\mathbf{x}_{t+1} | \mathbf{H}_{t+1}, \boldsymbol{\eta}_{t+1}, \boldsymbol{\lambda}_{t+1}) = |\mathbf{H}_{t+1}|^{-1/2} g(\mathbf{z}_{t+1} | \boldsymbol{\eta}_{t+1}, \boldsymbol{\lambda}_{t+1}),$$

see Jondeau and Rockinger (2006) for additional details. This parameterization guarantees that $\mathbf{z}_{t+1}$ has zero mean and unit variance for all admissible $\boldsymbol{\eta}_{t+1}$ and $\boldsymbol{\lambda}_{t+1}$.

We now express the moments entering equation (6). Let $\mathbf{S}_{t+1}$ and $\mathbf{K}_{t+1}$ denote the conditional

third and fourth central co-moments of exchange rate returns. We use the matrix representation $\mathbf{S}_{t+1} \in \mathbb{R}^{N \times N^2}$ and $\mathbf{K}_{t+1} \in \mathbb{R}^{N \times N^3}$, and Appendix C derives their expressions. Since $\mathbf{y}_{t+1} = \mathbf{d}_t + \mathbf{A}_t \mathbf{x}_{t+1}$, and $\mathbf{d}_t$ and $\mathbf{A}_t$ are known at time t , the central co-moments of currency excess returns $\mathbf{y}_{t+1}$ are obtained by deterministic scaling of the corresponding co-moments of $\mathbf{x}_{t+1}$

$$\mathbf{H}_{t+1}^y = \mathbf{A}_t \mathbf{H}_{t+1} \mathbf{A}_t, \quad \mathbf{S}_{t+1}^y = \mathbf{A}_t \mathbf{S}_{t+1} (\mathbf{A}_t \otimes \mathbf{A}_t), \quad \mathbf{K}_{t+1}^y = \mathbf{A}_t \mathbf{K}_{t+1} (\mathbf{A}_t \otimes \mathbf{A}_t \otimes \mathbf{A}_t). \quad (10)$$

The conditional mean and central moments of the portfolio excess return are, therefore,

$$\begin{aligned} \mu_{p,t+1} &= \mathbf{w}_t' \mathbf{d}_t, & s_{p,t+1} &= \mathbf{w}_t' \mathbf{S}_{t+1}^y (\mathbf{w}_t \otimes \mathbf{w}_t), \\ h_{p,t+1} &= \mathbf{w}_t' \mathbf{H}_{t+1}^y \mathbf{w}_t, & \kappa_{p,t+1} &= \mathbf{w}_t' \mathbf{K}_{t+1}^y (\mathbf{w}_t \otimes \mathbf{w}_t \otimes \mathbf{w}_t). \end{aligned} \quad (11)$$

The noncentral moments required in equation (6) are then

$$\begin{aligned} m_{p,t+1}^{(1)} &= \mu_{p,t+1}, & m_{p,t+1}^{(3)} &= s_{p,t+1} + 3 h_{p,t+1} \mu_{p,t+1} + \mu_{p,t+1}^3, \\ m_{p,t+1}^{(2)} &= h_{p,t+1} + \mu_{p,t+1}^2, & m_{p,t+1}^{(4)} &= \kappa_{p,t+1} + 4 s_{p,t+1} \mu_{p,t+1} + 6 h_{p,t+1} \mu_{p,t+1}^2 + \mu_{p,t+1}^4. \end{aligned} \quad (12)$$

3.2 Bayesian Estimation

Let Θ collect the GARCH parameters $\{\omega_i, \alpha_i, \beta_i\}_{i=1}^N$, the DCC parameters (δ_1, δ_2) , and the coefficients of the shape dynamics $\{\gamma_{i,\cdot}, \kappa_{i,\cdot}\}_{i=1}^N$. Define the likelihood

$$\mathcal{L}(\mathbf{x}_{1:T} | \Theta) = \prod_{t=1}^T g(\mathbf{x}_t | \mathbf{H}_t(\Theta), \boldsymbol{\eta}_t(\Theta), \boldsymbol{\lambda}_t(\Theta)).$$

The posterior distribution is

$$p(\Theta | \mathbf{x}_{1:T}) = \frac{\mathcal{L}(\mathbf{x}_{1:T} | \Theta) p(\Theta)}{\int \mathcal{L}(\mathbf{x}_{1:T} | \Theta) p(\Theta) d\Theta}.$$

The posterior is high-dimensional and nonlinear as the likelihood depends on GARCH recursions, DCC recursions, and logistic transformations in equation (9). Standard Markov Chain Monte Carlo (MCMC) can be difficult to tune in this environment, especially when the posterior is irregular or multimodal. We therefore use SMC, which represents the posterior by a weighted swarm of particles and is well suited for parallel computation and marginal-likelihood evaluation (Del Moral et al., 2006), and which has proven effective for GARCH-type models (Li et al., 2021).

For an initial estimation window, the SMC builds a sequence of tempered distributions

$$p_n(\Theta|\mathbf{x}_{1:T}) \propto \mathcal{L}(\mathbf{x}_{1:T}|\Theta)^{\phi_n} p(\Theta), \quad 0 = \phi_1 < \phi_2 < \dots < \phi_{N_\phi} = 1, \quad n = 1, \dots, N_\phi,$$

where N_ϕ is the maximum number of tempering steps. The algorithm starts from the prior ($\phi_1 = 0$), progressively introduces the likelihood through the tempering schedule ($\phi_n \rightarrow 1$), and is based on correction, selection, and mutation steps. The correction step updates particle weights, the selection step resamples particles when the effective sample size falls below a threshold, and the mutation step rejuvenates particles. The algorithm yields the posterior estimates and the marginal likelihood, which is useful for model comparison.

The recursive portfolio exercise requires posterior updates at every rebalancing date t . Re-estimating the model from the prior at each estimation window would be computationally costly; therefore, we use an online SMC update as in [Cai et al. \(2021\)](#). Using a rolling estimation framework, let $\mathcal{T}_t = 1, \dots, T$ be the old window and $\mathcal{T}_{t'} = 2, \dots, T+1$ the new window, so when the window rolls forward, new observations enter, and old observations leave. The online SMC replaces the old likelihood with the new one, that is

$$p_n(\Theta|\mathbf{x}_{\mathcal{T}_{t'}}) \propto \mathcal{L}(\mathbf{x}_{\mathcal{T}_{t'}}|\Theta)^{\phi_n} \mathcal{L}(\mathbf{x}_{\mathcal{T}_t}|\Theta)^{1-\phi_n} p(\Theta), \quad 0 = \phi_1 < \dots < \phi_{N_\phi} = 1, \quad n = 1, \dots, N_\phi. \quad (13)$$

At $\phi_1 = 0$, the target is the posterior based on the old window. After moving through the tempering schedule, at $\phi_{N_\phi} = 1$, the target is the posterior based on the new window. This defines a tempering bridge between the posterior based on the old window and that based on the new window, allowing the particle approximation from the former to be efficiently transformed into the latter. This formulation makes the online estimator consistent with the rolling-window design used in the empirical analysis.

Following [Jondeau and Rockinger \(2012\)](#) the priors are designed to concentrate mass on economically plausible regions of the parameter space, while remaining diffuse enough that inference is driven by the data. For the GARCH parameters, we use Beta ($\mathcal{B}$) priors that favor persistent but stable volatility dynamics

$$\omega_i \sim \mathcal{B}(2, 10), \quad \alpha_i \sim \mathcal{B}(2, 20), \quad \beta_i \sim \mathcal{B}(25, 1.5), \quad i = 1, \dots, N.$$

For the DCC parameters, we impose stationarity by assigning priors to δ_1 and to aggregate persistence

$$\delta_1 \sim \mathcal{B}(1.5, 100), \quad \delta_1 + \delta_2 \sim \mathcal{B}(100, 1.5),$$

For shape dynamics, we use centered Gaussian ($\mathcal{N}$) priors:

$$\gamma_{i,0}, \kappa_{i,0} \sim \mathcal{N}(0, 2), \quad \gamma_{i,1}^+, \gamma_{i,1}^-, \kappa_{i,1}^+, \kappa_{i,1}^- \sim \mathcal{N}(0, 2), \quad \gamma_{i,2}, \kappa_{i,2} \sim \mathcal{N}(0, 0.3).$$

These priors allow the data to generate time-varying skewness and kurtosis without imposing such variation a priori. Draws that violate stationarity restrictions are excluded. Details are reported in Appendix B.

3.3 Predictive Moments and Parameter Uncertainty

At each rebalancing date, the estimator delivers a particle approximation to the posterior,

$$\left\{ \Theta_t^{(j)}, \bar{\mathcal{V}}_t^{(j)} \right\}_{j=1}^M, \quad \sum_{j=1}^M \bar{\mathcal{V}}_t^{(j)} = 1,$$

where $\Theta_t^{(j)}$ are the parameter draws and $\bar{\mathcal{V}}_t^{(j)}$ are the associated normalized weights for $j = 1, \dots, M$. For each particle j , we compute the predictive second, third and fourth co-moments of exchange-rate changes from equation (10). For a candidate portfolio $\mathbf{w}_t$, the corresponding central portfolio moments are given in equation (11). The noncentral moments used in expected utility follow from the standard identities linking central and noncentral moments, and are given in equation (12).

Parameter uncertainty enters the investor's problem through the posterior distribution of the predictive moments. The important simplification in currency markets is that the conditional mean of currency excess returns does not depend on Θ , that is

$$\mu_{p,t+1}(\Theta) = \mathbf{w}_t' \mathbf{d}_t.$$

Thus, posterior uncertainty in Θ affects the allocation only through the second, third and fourth conditional co-moments. This is the sense in which the currency setting mitigates the first-moment estimation problem that often undermines optimized portfolio performance in equity applications (Chopra et al., 1993; DeMiguel et al., 2009; Ackermann et al., 2017).

We incorporate parameter uncertainty by maximizing posterior expected utility obtained by averaging the conditional noncentral moments over the posterior distribution of the model parameters as in Harvey et al. (2010). Using the fourth-order utility approximation in equation

(6), the Bayesian objective can be written as

$$\max_{\mathbf{w}_t \in \mathcal{W}} \sum_{k=1}^4 \phi_k \bar{m}_{p,t+1}^{(k)}, \quad \bar{m}_{p,t+1}^{(k)} = \int m_{p,t+1}^{(k)}(\boldsymbol{\Theta}) p(\boldsymbol{\Theta} | \mathcal{F}_t) d\boldsymbol{\Theta}.$$

The SMC approximation is

$$\bar{m}_{p,t+1}^{(k)} \approx \sum_{j=1}^M \bar{v}_t^{(j)} m_{p,t+1}^{(k)}(\boldsymbol{\Theta}_t^{(j)}), \quad k = 1, \dots, 4. \quad (14)$$

Relative to a plug-in rule based on the posterior mean of $\boldsymbol{\Theta}$, this procedure preserves the effect of parameter uncertainty and the nonlinear mapping from model parameters to predictive comoments. If the final particle swarm has been resampled and all weights are set equal, equation (14) reduces to a simple Monte Carlo average.

4 Empirical Evaluation

This section evaluates the economic value of higher-order distributional information for currency portfolio choice. We first describe the currency investment opportunity set and the out-of-sample portfolio construction. We then compare the MV and MVSK allocation rules across alternative investment mandates, focusing on how higher-order information affects realized portfolio distributions, exposure, and economic performance. Next, we place the optimized currency strategies in the context of conventional currency and asset-class benchmarks to assess their economic magnitude and return distinctiveness. Finally, we examine whether the principal performance differences remain economically meaningful after accounting for transaction costs.

4.1 The Currency Investment Opportunity

We study dynamic currency allocation from the perspective of a US investor. The risky investment universe consists of nine liquid foreign currencies: the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), Swedish krona (SEK), and Singapore dollar (SGD). The US dollar is the domestic numeraire and funding asset.

As discussed in Section 2.1, under the maintained no-predictability restriction, the conditional mean of currency excess returns is given by the forward-implied interest-rate differential $d_{i,t}$, observed at time t , while the econometric model in Section 3.1 is used to characterize the exchange-rates conditional covariance, co-skewness, and co-kurtosis.

Daily exchange-rate returns provide the information state used to characterize short-horizon variation in the conditional distribution. The model is estimated using rolling windows of 2000 observations, corresponding to approximately eight years of data. Portfolios are formed on Friday each week (or the previous common business day across all currencies), using the most recent predictive distribution. The realized portfolio return is then computed using Equation (1) over the holding period. This design allows the allocation rule to incorporate the daily evolution of the conditional distribution while limiting portfolio turnover and associated implementation costs, as well as the computational effort. Portfolio decisions and performance measures are evaluated out-of-sample over the allocation period spanning from January 2006 to September 2025 (4934 observations, 1030 portfolio formation dates).

The observations from January 1998 to December 2005 are used to initialize the first rolling estimation window. Posterior mean and 90% credible intervals of the model parameters at the initial batch estimation and at the final online update are reported in Appendix B.

Daily spot and one-week forward exchange rates are obtained from Refinitiv. The domestic risk-free rate r_t^d is proxied by the three-month US Treasury bill rate, also from Refinitiv, and converted to a simple one-week rate. The one-week foreign interest rate r_t^i is inferred from covered interest parity using Equation (2).

To put the economic magnitude of the currency strategies in perspective, We also report the performance of non-currency benchmarks. They consist of the S&P 500 total return index, the US 10-year Treasury total return index, the S&P GSCI total return index, an Equity-Bond (EB) portfolio allocated 50% to the S&P 500 and 50% to US 10-year Treasuries, and an Equity-Bond-Commodity (EBC) portfolio allocated 50% to the S&P 500, 40% to US 10-year Treasuries, and 10% to the S&P GSCI. All benchmark series are obtained from Refinitiv.

Table 2 reports summary statistics for the nine risky currency positions over the out-of-sample evaluation period. The table reports summary statistics for the one-week forward-implied interest-rate differential relative to the US dollar, and for the weekly spot exchange-rate return. Mean and volatility figures are annualized, while skewness and kurtosis are computed over the weekly holding frequency.

The summary statistics highlight three features of the data. First, forward-implied interest-rate differentials vary substantially across currencies, generating an economically relevant predetermined component of expected currency excess returns. Second, spot exchange-rate changes display pronounced departures from Gaussianity, with several currencies exhibiting substantial skewness, and kurtosis well above the Gaussian benchmark of three. Third, the spot component varies considerably across currencies and can offset the interest-rate component over long evaluation periods, so that currencies with higher interest-rate differentials do not mechanically deliver higher realized excess returns. These unconditional facts, by themselves, do not establish

Table 2: Summary statistics for currency returns

Currency	d^i	x^i			
	Mean	Mean	Volatility	Skewness	Kurtosis
AUD	1.193	0.072	12.448	-1.045	14.093
CAD	-0.067	-0.466	8.617	-0.525	7.365
CHF	-1.930	2.876	10.459	1.850	30.353
EUR	-0.983	0.213	8.934	-0.109	4.947
GBP	-0.232	-0.908	9.456	-0.383	6.662
JPY	-2.042	-0.789	9.892	0.467	4.849
NOK	0.201	-1.382	12.503	-0.476	10.847
SEK	-0.860	-0.350	11.651	0.044	4.578
SGD	-0.584	1.366	5.155	-0.181	5.332

Notes: Summary statistics for the one-week forward-implied interest-rate differential relative to the US dollar obtained from covered interest parity (d^i), and the weekly spot exchange-rate return (x^i) over the out-of-sample allocation period from January 2006 to September 2025. Mean and standard deviation are annualized and expressed in percentage points.

time variation in higher-order conditional moments. They do, however, show that a Gaussian MV benchmark is restrictive for describing the empirical distribution of currency payoffs and motivate a predictive framework that allows for time-varying asymmetry and tail thickness.

4.2 Distributional Timing and Portfolio Performance

We first examine whether incorporating time-varying higher-order moments into portfolio choice translates into economically meaningful differences in realized portfolio outcomes. Table 3 reports the realized distributional characteristics of the MV and MVSF portfolios across the four investment mandates and the two levels of relative risk aversion, $\gamma \in \{5, 15\}$. The comparison is particularly informative because, within each mandate, the MV and MVSF strategies use the same information set and conditional mean vector and are subject to the same investment constraints. The only difference is that MVSF additionally conditions on higher-order moments.

Several patterns emerge. First, MVSF generally produces higher mean excess returns than MV, although the magnitude of the difference varies across mandates. The largest differences occur under the unconstrained mandate $\mathcal{W}^U$. The mean excess return increases from 21.51% to 23.93% for $\gamma = 5$ and from 7.17% to 7.94% for $\gamma = 15$. The corresponding changes are much smaller when gross exposure is fixed, although the long-only mandate produces somewhat larger mean differences than the other fixed-exposure mandates. Thus, the gains from higher-order moments conditioning are not independent of the investor’s ability to adjust the scale and composition of the portfolio.

Second, the higher mean returns under MVSF are not accompanied by a uniform reduction in volatility. Under $\mathcal{W}^U$, volatility increases from 24.29% to 25.60% for $\gamma = 5$ and from 8.10%

Table 3: Return distribution and portfolio diagnostics

γ	Strategy	μ_p		σ_p		SSD $_p$		SK $_p$		KU $_p$		TO		HHI	
		MV	MVSK	MV	MVSK	MV	MVSK	MV	MVSK	MV	MVSK	MV	MVSK	MV	MVSK
$\mathcal{W}^U$ (Unconstrained)															
5		21.51	23.93	24.29	25.60	14.65	14.82	0.80	1.10	10.48	11.20	15.56	16.19	0.20	0.20
15		7.17	7.94	8.10	8.54	4.88	4.95	0.80	1.09	10.48	11.62	5.19	5.38	0.20	0.20
$\mathcal{W}^G$ (Unit gross exposure)															
5		2.55	2.64	6.12	5.98	4.22	4.08	0.07	0.18	8.27	9.15	1.26	1.27	0.57	0.55
15		1.58	1.60	3.52	3.45	2.44	2.38	-0.15	-0.15	6.69	6.88	1.22	1.22	0.36	0.35
$\mathcal{W}^{DN}$ (Dollar-neutral)															
5		2.50	2.52	4.19	4.16	2.87	2.84	-0.23	-0.13	6.27	6.87	1.24	1.24	0.40	0.40
15		1.88	1.90	3.06	3.02	2.11	2.07	-0.41	-0.37	5.09	5.37	1.19	1.19	0.30	0.29
$\mathcal{W}^L$ (Long-only)															
5		0.88	1.08	7.24	7.25	5.12	5.14	-0.13	-0.18	4.96	5.18	0.72	0.72	0.23	0.23
15		0.30	0.53	6.33	6.33	4.50	4.46	-0.14	-0.09	5.24	5.33	0.52	0.51	0.20	0.20

Notes: MV denote the mean-variance and MVSK denote the mean-variance-skewness-kurtosis allocation rules, respectively. The Table reports portfolio Mean (μ_p), Volatility (σ_p), and Semi-Standard Deviation (SSD_p) annualized and expressed in percentage points. Skewness (SK_p) and Kurtosis (KU_p) are reported at portfolio return frequency to avoid vanishing effects. TO denotes gross portfolio turnover per rebalancing period. HHI denotes the Normalized Herfindahl–Hirschman concentration index.

to 8.54% for $\gamma = 15$, together with small increases in semi-volatility. By contrast, volatility and semi-volatility decline slightly under $\mathcal{W}^G$ and $\mathcal{W}^{DN}$, while remaining essentially unchanged under $\mathcal{W}^L$. The results therefore do not suggest that the MVSK gains arise simply from lower variance. Rather, MVSK changes the overall trade-off among expected return and the different dimensions of portfolio risk.

Third, the realized higher-order moments do not move uniformly in the direction that would mechanically imply improved performance. Under $\mathcal{W}^U$, MVSK increases realized skewness by approximately 0.30 at both risk-aversion levels, but also increases kurtosis. Under $\mathcal{W}^G$ and $\mathcal{W}^{DN}$, skewness generally becomes more favorable, while kurtosis also increases. Under the long-only mandate, the change in skewness also depends on risk aversion. It becomes slightly more negative for $\gamma = 5$ but less so for $\gamma = 15$. These results are consistent with the portfolio objective. The MVSK investor does not maximize skewness or minimize kurtosis alone, but jointly evaluates all moments through the fourth-order approximation to expected CRRA utility. Consequently, the relevant question is whether the joint distributional information improves portfolio choice, rather than whether every individual distributional characteristic improves separately.

The differences across mandates are particularly informative about the portfolio channels through which higher-order moments information can operate. Under $\mathcal{W}^U$, the investor can adjust both portfolio composition and total exposure, whereas the remaining mandates fix gross exposure and therefore remove the leverage-scaling margin. The substantially larger mean differences

under $\mathcal{W}^U$ are therefore consistent with exposure scaling being an important channel through which higher-order moments information is monetized. At the same time, the positive, although smaller, changes in mean return under the fixed gross-exposure mandates show that higher-order information retains value when leverage adjustment is removed, implying that cross-sectional reallocations also play a role.

The dollar-neutral and long-only mandates provide further variation in the margins through which this information can be exploited. Under $\mathcal{W}^{DN}$, the investor can reallocate across currencies but cannot take a net directional position, while $\mathcal{W}^L$ additionally rules out short positions. The resulting differences in realized means and higher moments indicate that the effect of higher-order moments conditioning depends on the set of portfolio adjustments available to the investor rather than on the conditional distribution alone.


Turnover and concentration provide little evidence that MVSK systematically generates more concentrated or more actively traded portfolios. HHI is unchanged under $\mathcal{W}^U$ and $\mathcal{W}^L$ and declines modestly under $\mathcal{W}^G$ and $\mathcal{W}^{DN}$. Turnover is very similar between MV and MVSK under the constrained mandates and increases substantially under $\mathcal{W}^U$. However, considering that turnover mechanically increases with leverage, when normalizing $\mathcal{W}^U$'s TO by gross exposure, the figures align with those of the other mandates. Indeed, TO ratios unconstrained-to-constrained mandates match $\mathcal{W}^U$'s gross exposure distribution reported in Figure 1. These results indicate similar trading intensity per unit of exposure. As such they do not by themselves rule out higher absolute trading costs under the more highly leveraged unconstrained portfolios. The latter remains an important implementation consideration, further discussed in Section 4.4.

Taken together, the table shows that higher-order moments conditioning changes the realized return distribution in economically meaningful but non-uniform ways. The largest differences in mean return arise when the investor can adjust total exposure, while smaller differences remain when gross exposure is fixed. The next question is therefore whether these changes in the joint return distribution translates into economically valuable improvements in risk-adjusted performance and investor welfare.

Figure 1 reports realized gross leverage under $\mathcal{W}^U$, which serves as a frictionless diagnostic of the exposure implied by the estimated conditional distribution. For a given degree of risk aversion, the distributions of gross exposure of MV and MVSK are very close, although the latter exhibits a longer right tail. For $\gamma = 5$, median leverage is approximately 12, with an interquartile range of $[8, 17]$, while leverage is substantially lower for $\gamma = 15$. Leverage of a similar broad order has been documented for institutional currency and multi-strategy funds using derivatives (Packer et al., 2024). These results should nevertheless be interpreted as an upper-bound assessment of the scope for distributional timing when exposure can be freely scaled.

Table 4 summarizes the economic and statistical evidence. Across the four mandates, two

Figure 1: Distribution of realized gross exposure for $\mathcal{W}^U$



Notes: Results are shown for mean-variance (MV) and mean-variance-skewness-kurtosis (MVSK) rules for risk aversion $\gamma = 5$ and $\gamma = 15$ over the out-of-sample evaluation period. Boxes span the interquartile range, the horizontal line denotes the median, and whiskers extend to the most extreme observations within 1.5 times the interquartile range; points beyond this range are plotted as outliers.

levels of risk aversion, and six performance measures, all 48 MVSK–MV point estimates are positive. This consistency is informative about the direction of the estimated effects, although it does not imply that the gains are statistically distinguishable from zero in every portfolio environment. In fact, although the fourth-order CRRA approximation assigns a positive weight to the third portfolio moment and a negative weight to the fourth moment (Scott and Horvath, 1980), an out-of-sample improvement is not guaranteed in the presence of parameter and model uncertainty (DeMiguel et al., 2009).

The strongest evidence is obtained under $\mathcal{W}^U$. The annualized performance fee is approximately 113 basis points for $\gamma = 5$ and 32 basis points for $\gamma = 15$. The corresponding SR gains are 0.049 and 0.044, and all six reported performance differences are statistically significant at least at the 10% level. Taken together with Table 3, these results indicate that the additional information contained in higher-order conditional moments can have substantial economic value when the investor can adjust both the composition and scale of currency exposure. At the same time, the magnitude of these gains should be interpreted jointly with the high leverage documented in Figure 1. Financing costs, collateral requirements, and leverage-dependent margin costs are not incorporated here and would attenuate the net economic value of the unconstrained strategy.

The gains are substantially smaller under the constrained mandates. Under $\mathcal{W}^G$, the performance fee is approximately 13 basis points for $\gamma = 5$ and 6 basis points for $\gamma = 15$. All point estimates remain positive, but the performance fee and conventional SR differences are not statistically significant. For $\gamma = 5$, the adjusted Sharpe, modified Sharpe, Omega, and Sortino

Table 4: Performance gains from distribution timing

	Φ	ΔSR_p	ΔaSR_p	$\Delta\text{mSR}_{5\%,p}$	$\Delta\Omega_p\text{-ratio}$	$\Delta\mathcal{S}_p\text{-ratio}$
$\mathcal{W}^U$ (Unconstrained)						
$\gamma = 5$	113.057**	0.049**	0.041**	0.138**	0.050**	0.145**
$\gamma = 15$	32.106*	0.044**	0.025*	0.140**	0.049**	0.135*
$\mathcal{W}^G$ (Unit gross exposure)						
$\gamma = 5$	12.539	0.023	0.024*	0.028*	0.012*	0.040*
$\gamma = 15$	6.495	0.017	0.017	0.012	0.009	0.028
$\mathcal{W}^{DN}$ (Dollar-neutral)						
$\gamma = 5$	2.563	0.009	0.009	0.016	0.005	0.019
$\gamma = 15$	3.121	0.012	0.012	0.012	0.006	0.023
$\mathcal{W}^L$ (Long-only)						
$\gamma = 5$	19.422	0.027	0.027	0.016	0.011	0.038
$\gamma = 15$	24.249*	0.036*	0.036*	0.023*	0.014*	0.052*

Notes: Φ is the annualized performance fee, expressed in basis points. Sharpe ratio (SR_p), adjusted Sharpe ratio (aSR_p), modified Sharpe ratio ($\text{mSR}_{q,p}$), Omega ratio ($\Omega_p\text{-ratio}$) and Sortino ratio ($\mathcal{S}_p\text{-ratio}$) are annualized and expressed in ratio units. ΔSR_p , ΔaSR_p , $\Delta\text{mSR}_{5\%,p}$, $\Delta\Omega_p\text{-ratio}$, and $\Delta\mathcal{S}_p\text{-ratio}$ denote the difference between Mean-Variance-Skewness-Kurtosis (MVSK) and Mean-Variance (MV) in the corresponding performance measure. Positive values indicate better performance under MVSK. The $\text{mSR}_{q,p}$ uses a $q = 5\%$ Value-at-Risk threshold. Statistical significance refers to whether MVSK improves upon MV. Test statistics and bootstrap procedures are reported in Appendix E. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

differences are significant at the 10% level, whereas none of the performance differences is statistically significant for $\gamma = 15$. Thus, fixing gross exposure substantially reduces the economic value of higher-order moments conditioning, although the point estimates remain positive.

Under $\mathcal{W}^{DN}$, the estimated gains are smaller still. The performance fees are approximately 3 basis points per year for both levels of risk aversion, and none of the reported performance differences is statistically significant at conventional levels. The positive point estimates therefore provide, at most, weak evidence of an economic gain once both leverage and net directional exposure are removed.

The long-only results are different. For $\gamma = 5$, the performance fee is approximately 19 basis points, although none of the individual performance differences is statistically significant. For $\gamma = 15$, the estimated performance fee rises to approximately 24 basis points, and all six performance differences are significant at the 10% level. Thus, despite the absence of leverage and short selling, higher-order moments information can still generate economically meaningful gains for the more risk-averse long-only investor considered here.

The contrast between $\mathcal{W}^G$ and $\mathcal{W}^L$ also shows that the value of higher-order moments information does not vary monotonically with risk aversion across mandates. Under $\mathcal{W}^G$, statistically detectable gains arise only for the lower risk-aversion investor, whereas under $\mathcal{W}^L$ they arise only

for the higher risk-aversion investor. This asymmetry is consistent with the feasible set shaping how higher-order moments information can be translated into portfolio positions. Under $\mathcal{W}^G$, the investor can adjust net directional and cross-sectional exposure within a fixed gross-exposure budget, whereas under the long-only mandate, the response is confined to reallocations within the foreign-currency basket.

Taken together, the evidence indicates that the economic value of distribution timing is positive but strongly mandate-dependent. The largest and most statistically robust gains occur under the unconstrained mandate, where the investor can adjust total exposure as well as portfolio composition. The gains become much smaller when gross exposure is fixed and are weakest under the dollar-neutral mandate, while the long-only results show that economically meaningful gains can nevertheless survive when leverage and short selling are prohibited. In conclusion, our results show that higher-order conditional distributional information can generate improvements in investor welfare and risk-adjusted performance, with the magnitude and statistical strength of the gains depending critically on the feasible set and the investor’s risk aversion.

4.3 Currency and Asset-Class Benchmarks

We next benchmark the performance of MVSK against alternative investment strategies. We first compare MVSK with currency benchmarks designed to reflect the constraints of each investment mandate, and then with traditional equity, bond, and commodity investments. These comparisons provide context for the economic magnitude and return characteristics of the optimized currency strategies.

4.3.1 Currency Benchmarks

Table 5 compares MVSK with a benchmark appropriate to each feasible set. For the long-only mandate $\mathcal{W}^L$, the natural benchmark is the equally weighted (EW) currency portfolio. For the gross-only and dollar-neutral mandates, $\mathcal{W}^G$ and $\mathcal{W}^{DN}$, we use the long-four/short-four carry strategy (L4S4), which ranks currencies by their forward discount and takes equal-weighted long positions in the four highest- yielding currencies and short positions in the four lowest-yielding currencies. This adapts the standard cross-sectional carry construction to our nine-currency universe (Lustig et al., 2011; Menkhoff et al., 2012). To ease comparability with the unconstrained mandate, we report L4S4* which is the leverage-scaled version of L4S4, such that its gross exposure is 10. This level of exposure is motivated by Figure 1. It is worth noting that leverage scales mean, volatility, and turnover linearly, thus leaving all other measures unaffected. Because the same scale is used to benchmark MVSK for both levels of risk aversion, L4S4* serves as a common descriptive reference rather than as a risk-matched benchmark for each MVSK portfolio.

Table 5: Return distribution and portfolio diagnostics: MVSK and Currency Benchmarks

γ	Strategy	μ_p	σ_p	SK_p	KU_p	SR_p	aSR _p	Ω_p -ratio	$\mathcal{S}_p$ -ratio	TO	HHI
$\mathcal{W}^U$ (Unconstrained)											
	L4S4*	13.091	30.873	-0.716	10.645	0.424	0.378	1.192	0.586	6.959	0.125
5	MVSK	23.925	25.595	1.096	11.204	0.935	0.815	1.478	1.614	16.189	0.201
15	MVSK	7.941	8.539	1.095	11.622	0.930	0.799	1.476	1.603	5.376	0.203
$\mathcal{W}^G$ (Unit gross exposure)											
	L4S4	1.309	3.087	-0.716	10.645	0.424	0.378	1.192	0.586	0.696	0.125
5	MVSK	2.635	5.981	0.184	9.149	0.441	0.425	1.189	0.646	1.266	0.552
15	MVSK	1.603	3.450	-0.148	6.883	0.465	0.443	1.194	0.675	1.219	0.350
$\mathcal{W}^{DN}$ (Dollar-neutral)											
	L4S4	1.309	3.087	-0.716	10.645	0.424	0.378	1.192	0.586	0.696	0.125
5	MVSK	2.516	4.159	-0.134	6.866	0.605	0.561	1.257	0.887	1.243	0.397
15	MVSK	1.897	3.017	-0.371	5.371	0.629	0.580	1.263	0.916	1.189	0.295
$\mathcal{W}^L$ (Long-only)											
	EW	-0.315	7.371	-0.059	5.624	-0.043	-0.043	0.984	-0.060	0.000	0.111
5	MVSK	1.078	7.247	-0.184	5.177	0.149	0.148	1.058	0.210	0.717	0.234
15	MVSK	0.534	6.328	-0.089	5.332	0.084	0.084	1.032	0.120	0.512	0.200

Notes: EW denotes the equally weighted currency portfolio. L4S4 denotes the long-four/short-four carry portfolio, which takes equal-weighted long positions in the four highest-yielding currencies and short positions in the four lowest-yielding currencies. L4S4* denotes L4S4 scaled to a gross exposure of 10. Benchmark statistics do not depend on risk aversion and are therefore reported once per mandate. Portfolio Mean (μ_p) and Volatility (σ_p) are annualized and expressed in percentage points. Skewness (SK_p) and Kurtosis (KU_p) are reported at portfolio return frequency to avoid vanishing effects. Sharpe ratio (SR_p), adjusted Sharpe ratio (aSR_p), Omega ratio (Ω_p -ratio) and Sortino ratio ($\mathcal{S}_p$ -ratio) are annualized and reported in ratio units. TO denotes gross portfolio turnover per rebalancing period. HHI denotes the normalized Herfindahl-Hirschman concentration index.

For the long-only investor, EW generates a negative annualized mean excess return of -0.32% and a SR of -0.043 . Both MVSK portfolios have positive mean returns and SRs. For $\gamma = 5$, the mean return is 1.08% with a SR of 0.149 , while for $\gamma = 15$ the corresponding values are 0.53% and 0.084 . Volatility is similar to or slightly below that of EW, although MVSK allocations are more concentrated, with HHI of 0.234 ($\gamma = 5$) and 0.200 ($\gamma = 15$) which correspond to an effective number of positions ($1/\text{HHI}$) of about 4 and 5, respectively, and require positive turnover. The comparison therefore favors MVSK, while also indicating that the improvement comes with a more active allocation than passive equal weighting.

The comparison with L4S4 is more nuanced under the remaining fixed-gross-exposure mandates. Under $\mathcal{W}^G$, MVSK produces higher mean returns than L4S4, but also higher volatility. The SRs of 0.441 and 0.465 are only moderately above the L4S4 value of 0.424 . aSR and $\mathcal{S}$ -ratios are also higher, while the Ω -ratio is slightly lower for $\gamma = 5$. MVSK is also substantially more concentrated than L4S4, with HHI between 0.35 and 0.55 , which corresponds to an effective number of positions of about 3 and 2 respectively, and requires higher turnover. The gross-only comparison therefore points to a broadly similar risk-adjusted profile, with modest improvements in some dimensions rather than uniform gains across all measures.

The dollar-neutral comparison is more favorable. The MVSK SRs are 0.605 and 0.629, compared with 0.424 for L4S4, and the aSR, Ω -ratio, and $\mathcal{S}$ -ratio are higher at both levels of risk aversion. MVSK also combines higher mean returns with less negative skewness and lower kurtosis than L4S4. These results indicate that conditioning on higher-order moments can materially reshape the cross-sectional currency portfolio even when gross and net exposure are fixed. This comes with greater concentration and turnover. MVSK has HHI between 0.30 and 0.40 and roughly 1.2 turnover per rebalancing period, versus 0.125 and 0.696 for L4S4.

The largest differences arise under the unconstrained mandate $\mathcal{W}^U$. MVSK produces Sharpe ratios of approximately 0.93, compared with 0.424 for L4S4*, and also records substantially higher aSR, Ω -ratio, and $\mathcal{S}$ -ratio. Because the Sharpe-type measures are invariant to a common scaling of portfolio weights in the absence of nonlinear costs, the difference cannot be attributed mechanically to the larger nominal scale of L4S4*. At the same time, L4S4* is not separately matched to each MVSK portfolio in terms of leverage or volatility, and the two strategies have materially different return distributions. The comparison is therefore best interpreted as evidence of the magnitude of the realized risk-adjusted performance difference rather than as a controlled test of relative performance. MVSK is also more concentrated than L4S4*, with HHI of about 0.2 versus 0.125, and requires substantially higher turnover per unit of exposure than L4S4* for $\gamma = 5$, but lower than L4S4* for $\gamma = 15$.

Taken together, the currency-benchmark comparisons show that MVSK improves substantially over passive equal weighting in the long-only setting and is competitive with conventional cross-sectional carry benchmarks under the other mandates, with the largest risk-adjusted differences occurring under the unconstrained and dollar-neutral mandates. At the same time, the magnitude and composition of the differences vary across feasible sets. The gross-only comparison is more modest, whereas the long-only gains come with greater concentration and active rebalancing.

4.3.2 Asset-Class Benchmarks

Table 6 places the currency strategies in the broader context of traditional asset classes, namely S&P 500, the US 10-year Treasury, the S&P GSCI, and two balanced portfolios. The EB portfolio allocates 50% to equities and 50% to bonds, while EBC allocates 50% to equities, 40% to bonds, and 10% to commodities.

Among the traditional benchmarks, EB has the strongest overall risk-adjusted profile on several measures. Its SR is 0.679, its Ω -ratio is 1.324, and its $\mathcal{S}$ -ratio is 0.970. The S&P 500 has a higher mean return of 9.88% but substantially higher volatility, negative skewness, and high kurtosis. The GSCI performs poorly over the evaluation period, with a negative mean return and SR.

Table 6: Return distribution and portfolio diagnostics: Asset-Class benchmarks

Benchmark	μ	σ	SK	KU	SR	aSR	Ω -ratio	$\mathcal{S}$ -ratio
S&P 500	9.885	17.269	-1.081	15.68	0.572	0.414	1.269	0.782
US 10Y	1.754	7.406	0.348	4.750	0.237	0.239	1.091	0.353
GSCI	-1.726	22.108	-0.549	5.830	-0.078	-0.079	0.971	-0.104
EB	5.819	8.566	-0.592	15.28	0.679	0.473	1.324	0.970
EBC	5.471	9.424	-0.880	15.67	0.581	0.428	1.273	0.807

Notes: All benchmark performance measures are expressed in terms of excess returns. The Table reports the benchmark Mean (μ), and Volatility (σ) annualized and expressed in percentage points. Skewness (SK), Kurtosis (KU) are reported at portfolio return frequency to avoid vanishing effects. Sharpe ratio (SR), adjusted Sharpe ratio (aSR), Omega ratio (Ω -ratio) and Sortino ratio ($\mathcal{S}$ -ratio) are annualized and reported in ratio units. S&P 500 denotes the S&P 500 total return index; US 10Y denotes the US 10-year Treasury total return index; and GSCI denotes the S&P GSCI total return index. EB denotes a portfolio composed of 50% S&P 500 and 50% US 10Y. EBC denotes a portfolio composed of 50% S&P 500, 40% US 10Y, and 10% GSCI.

The unconstrained MVSK portfolios have substantially higher realized SR, aSR, Ω -ratio, and $\mathcal{S}$ -ratio than all five asset classes benchmark. Their SRs of approximately 0.93, for example, are well above the 0.679 achieved by EB. These differences provide a useful benchmark for the economic magnitude of the currency results, while interpreting their relative performance should account for differences in exposure, leverage, financing requirements, and return distributions.

Because of the limits imposed on gross and net exposure, the dollar-neutral MVSK portfolios provide a more stringent comparison. Their SRs of 0.605 and 0.629 exceed those of the S&P 500, US 10-year Treasury, GSCI, and EBC, although they remain below the 0.679 of EB. Their aSRs exceed that of EB, while their Ω -ratio and $\mathcal{S}$ -ratio remain below it. The dollar-neutral strategy therefore compares favorably with traditional diversified portfolios on several dimensions, while its relative performance varies across the risk-adjusted measures.

The gross-only MVSK portfolios, with SRs of 0.441 and 0.465, outperform the Treasury and commodity benchmarks but remain below the S&P 500 and the balanced portfolios. The long-only MVSK portfolios have yet lower SRs, although both generate positive mean excess returns. Overall, the asset-class comparison confirms that the economic profile of MVSK depends materially on the investment mandate. The strongest risk-adjusted results occur when the portfolio can exploit either flexible exposure or dollar-neutral relative-value positioning, whereas the gross-only and long-only implementations are more modest.

4.3.3 Correlation with Currency and Asset-Class Benchmarks

Table 7 reports unconditional correlations between realized MVSK returns and the currency and asset-class benchmarks. These correlations provide complementary evidence on the extent to which the optimized currency strategies co-move with standard investment opportunities.

The long-only MVSK portfolios are highly correlated with the equally weighted currency

Table 7: Correlation with currency and asset-class benchmarks

Benchmark	$\mathcal{W}^U$		$\mathcal{W}^G$		$\mathcal{W}^{DN}$		$\mathcal{W}^L$	
	(Unconstrained)		(Unit gross exposure)		(Dollar-neutral)		(Long-only)	
	$\gamma = 5$	$\gamma = 15$	$\gamma = 5$	$\gamma = 15$	$\gamma = 5$	$\gamma = 15$	$\gamma = 5$	$\gamma = 15$
EW	-0.077	-0.078	-0.060	-0.099	0.169	0.125	0.946	0.953
L4S4	0.398	0.398	0.496	0.514	0.639	0.620	0.374	0.277
SP500	0.103	0.103	0.204	0.192	0.309	0.263	0.468	0.416
US 10y	-0.060	-0.060	-0.252	-0.202	-0.213	-0.162	0.054	0.116
GSCI	0.039	0.039	0.135	0.100	0.230	0.185	0.420	0.387
EB	0.078	0.078	0.097	0.106	0.219	0.195	0.495	0.470
EBC	0.085	0.085	0.140	0.135	0.270	0.233	0.544	0.508

Notes: EW denotes the equally weighted currency portfolio. L4S4 denotes the long-four short-four carry strategy. SP500 denotes the S&P 500 total return index, US 10y denotes the US 10-year Treasury total return index, and GSCI denotes the S&P GSCI total return index. EB denotes a balanced equity-bond portfolio composed of 50% S&P 500 and 50% US 10Y. EBC denotes a balanced equity-bond-commodity portfolio composed of 50% S&P 500, 40% US 10Y, and 10% GSCI.

portfolio, with correlations of approximately 0.95 at both levels of risk aversion. They also display the highest correlations with the equity, commodity, and balanced asset-class benchmarks, ranging from approximately 0.39 to 0.54. This pattern is consistent with the long-only mandate retaining substantial directional exposure to the foreign-currency basket.

Correlations with traditional asset classes are substantially lower under the other three mandates. For $\mathcal{W}^U$, correlations with the S&P 500, US 10-year Treasury, GSCI, EB, and EBC range from approximately -0.06 to 0.10 . The gross-only and dollar-neutral portfolios display somewhat higher positive correlations with risky asset classes, but these remain well below those of the long-only strategy. Their correlations with Treasury returns are negative.

The relationship with currency benchmarks also varies across mandates. The dollar-neutral portfolio has the highest correlations with L4S4, at approximately 0.63, followed by the gross-only portfolios. The unconstrained portfolios have an L4S4 correlation of about 0.40, while as noted the long-only portfolios are more closely related to EW than to the cross-sectional carry benchmark. The optimized portfolios, therefore, differ not only in exposure scale but also in the type of currency risk they capture.

The relatively low correlations of the unconstrained, gross-only, and dollar-neutral strategies with traditional asset classes point to potentially valuable diversification opportunities when combined with broader portfolios. They also highlight that the optimized currency strategies have return dynamics that differ materially from standard asset-class exposures.

Overall, the benchmark analysis provides three observations. First, MVSK compares favorably with passive equal weighting in the long-only setting and is generally competitive with the diversified carry benchmark under the other mandates. Second, the unconstrained and dollar-neutral implementations deliver the strongest realized risk-adjusted profiles relative to traditional

asset-class benchmarks, while the gross-only and long-only portfolios are more modest. Third, the optimized currency portfolios exhibit materially different correlations with standard currency and asset-class benchmarks across investment mandates, indicating that portfolio constraints affect not only performance but also the type of risk exposure generated by the allocation rule.

4.4 Transaction Costs and Gross-to-Net Performance

We next assess whether the economic differences documented above survive implementation costs. As discussed in Section 2.1, currency positions are implemented through forward contracts. A forward is generally synthesized by a spot transaction followed by an FX swap, reversing that spot transaction. In the absence of rebalancing, rolling over the currency forward requires rolling over only the FX swap, thus avoiding the initial spot transaction. This is the most cost-effective implementation because the investor only pays the much lower transaction costs for the FX swap. Spot transaction costs occur only if the weight in a currency changes, and only to the extent of the marginal change in the weight. Thus, at each rebalancing date, the total cost can be decomposed into two components: changes in target positions generate rebalancing costs, while maintaining the outstanding forward exposure generates rollover costs.

Let c_i^{spot} and c_i^{swap} denote the proportional bid-ask half-spreads for spot and swap transactions in currency i , respectively, and let $\Delta w_{i,t} = w_{i,t} - w_{i,t-1}$ denote the change in the target portfolio weight between two consecutive rebalancing dates. We use a reduced-form proportional-cost specification in which rebalancing costs are proportional to changes in target positions, whereas rollover costs are proportional to the post-rebalancing outstanding exposure. The total transaction cost incurred at date t is

$$TC_t = \sum_{i=1}^N \left(c_i^{\text{spot}} |\Delta w_{i,t}| + c_i^{\text{swap}} |w_{i,t}| \right).$$

For the gross-only, long-only, and dollar-neutral mandates, gross exposure is fixed at one, so rollover costs primarily reflect the cost of maintaining unit currency exposure after rebalancing. Under $\mathcal{W}^U$, rollover costs instead vary with the endogenous scale of the portfolio and therefore increase with realized gross exposure. We use constant transaction-cost assumptions across currencies and over time. Specifically, the proportional spot transaction cost c_i^{spot} is set to 1 basis point and the proportional swap transaction cost c_i^{swap} to 0.1 basis point. The spot-cost assumption is consistent with the order of magnitude of transaction costs reported for liquid currency markets in [Ackermann et al. \(2017\)](#) and [Della Corte et al. \(2009\)](#). More precisely, the 1-basis-point assumption is of the same order as transaction costs reported for liquid G10 currency markets over the post-2010 electronic-trading sample and it is conservative for the major

pairs, whose quoted half-spreads are routinely below one basis point. Darvas (2009) reports larger estimates but reflects an earlier trading environment and a wider currency set. As a robustness check, Appendix F.1 doubles both cost assumptions and shows that the main rankings between MVSK and the benchmark are unchanged. Appendix F.2 examines the sensitivity of the currency-benchmark comparisons to the amount of gross exposure the investor is permitted to take. We consider a spectrum of gross-exposure caps, filling the gap between the unit gross exposure mandate and the unconstrained one.

The costs incurred at rebalancing date t are deducted from the portfolio return earned over the subsequent holding period as $y_{t+1}^{p,\text{net}} = y_{t+1}^p - TC_t$. Figure 2 reports the resulting net cumulative wealth of MVSK at $\gamma = 5$ and $\gamma = 15$ against the relevant currency benchmark for each investment mandate, while Table 8 reports the corresponding net return-distribution and risk-adjusted performance measures.


The main result is that the economic value of higher-order moments conditioning largely survives the implementation costs considered here. The MVSK strategy has higher reported risk-adjusted performance than its benchmark at both risk-aversion levels under $\mathcal{W}^U$, $\mathcal{W}^{DN}$, and $\mathcal{W}^L$, while under $\mathcal{W}^G$ the ranking survives for $\gamma = 5$ but reverses at $\gamma = 15$. The extent of the gross-to-net adjustment varies substantially across mandates, reflecting differences in turnover and, most importantly, the scale of outstanding currency exposure.

The unconstrained investor exhibits the largest gross-to-net cost adjustment. Its net Sharpe ratio remains approximately 0.59 at both levels of risk aversion, compared with 0.294 for the net L4S4* benchmark. The high level of gross exposure and turnover nevertheless generates a substantial absolute transaction-cost drag relative to the constrained mandates. Figure 2 shows the corresponding separation in cumulative wealth, particularly for $\gamma = 5$, while the $\gamma = 15$ strategy compounds more moderately. The persistence of the MVSK advantage after these costs indicates that the gross performance differences under $\mathcal{W}^U$ are not driven solely by costless rebalancing, although the high leverage remains an important consideration for the net economic interpretation.

The dollar-neutral investor displays a similar but more moderate pattern. Net Sharpe ratios are 0.44 and 0.41 for $\gamma = 5$ and $\gamma = 15$, respectively, compared with 0.294 for L4S4. The corresponding cumulative wealth series also show a persistent separation between MVSK and the benchmark over the evaluation period. Because gross and net exposure are fixed, these results provide particularly direct evidence that the economic value of higher-order moments conditioning can survive implementation costs even when the leverage and directional-exposure margins are removed.

The long-only results also survive the cost adjustment. Net Sharpe ratios are 0.092 and 0.036 for $\gamma = 5$ and $\gamma = 15$, respectively, while the equally weighted benchmark remains negative at

Figure 2: Cumulative wealth across investor types



Notes: Net-of-cost cumulative wealth for the MVSK strategy at risk aversion $\gamma = 5$ and $\gamma = 15$ against the long-four/short-four carry benchmark L4S4 (L4S4* for L4S4 scaled to a gross exposure of 10 used for W^U), and the equally weighted currency portfolio EW (benchmark for W^L) and across the four investment mandates. Shaded bands denote, from left to right, the Global Financial Crisis, the European sovereign debt crisis, the 2014–2016 oil price crash, the COVID-19 shock, and the 2022 Ukraine war and USD shock. Vertical dashed lines denote the 2011 Swiss franc floor introduction, the January 2015 SNB shock, the June 2016 Brexit referendum, and the August 2024 yen carry-trade unwind.

−0.050. MVSK therefore retains positive net risk-adjusted performance despite the additional turnover and concentration relative to passive equal weighting. The cumulative wealth paths show the common weakness of the long-only currency exposure during the middle and later part of the sample, while MVSK remains above EW throughout this episode and into the subsequent

Table 8: Return distribution and performance diagnostics net of transaction costs: MVSK and Currency Benchmarks

γ	Strategy	μ_p	σ_p	SK_p	KU_p	SR_p	aSR_p	Ω_p -ratio	$\mathcal{S}_p$ -ratio
$\mathcal{W}^U$ (Unconstrained)									
	L4S4*	9.082	30.860	-0.709	10.625	0.294	0.276	1.129	0.403
5	MVSK	15.077	25.496	1.041	11.239	0.591	0.581	1.279	0.971
15	MVSK	5.002	8.506	1.038	11.646	0.588	0.575	1.278	0.965
$\mathcal{W}^G$ (Unit gross exposure)									
	L4S4	0.908	3.086	-0.709	10.625	0.294	0.276	1.129	0.403
5	MVSK	1.947	5.978	0.185	9.135	0.326	0.320	1.136	0.472
15	MVSK	0.939	3.447	-0.154	6.875	0.272	0.267	1.110	0.388
$\mathcal{W}^{DN}$ (Dollar-neutral)									
	L4S4	0.908	3.086	-0.709	10.625	0.294	0.276	1.129	0.403
5	MVSK	1.840	4.158	-0.128	6.841	0.442	0.424	1.182	0.639
15	MVSK	1.249	3.015	-0.368	5.337	0.414	0.397	1.166	0.591
$\mathcal{W}^L$ (Long-only)									
	EW	-0.366	7.371	-0.059	5.624	-0.050	-0.050	0.981	-0.069
5	MVSK	0.667	7.248	-0.183	5.175	0.092	0.092	1.036	0.129
15	MVSK	0.226	6.329	-0.089	5.332	0.036	0.036	1.013	0.050

Notes: Results are net of transaction costs. Portfolio Mean (μ_p) and Volatility (σ_p) are annualized and expressed in percentage points. Skewness (SK_p) and Kurtosis (KU_p) are reported at portfolio return frequency to avoid vanishing effects. Sharpe ratio (SR_p), adjusted Sharpe ratio (aSR_p), Omega ratio (Ω_p -ratio) and Sortino ratio ($\mathcal{S}_p$ -ratio) are annualized and reported in ratio units. Benchmark definitions are given in the notes to Table 5. Benchmark statistics do not depend on risk aversion and are therefore reported once per mandate.

recovery.

The gross-only mandate illustrates the limits of the gross performance gains once implementation costs are taken into account. At $\gamma = 5$, the net SR of MVSK is 0.326, slightly above the 0.294 of L4S4. At $\gamma = 15$, however, the MVSK ratio falls to 0.272, below the benchmark. The gross advantage is therefore largely absorbed by implementation costs for the more risk-averse investor. This result is consistent with the weaker gross-of-cost performance documented for $\mathcal{W}^G$ in Section 4.2. Modest gross performance differences can be sensitive to turnover and rollover costs even when the underlying allocation rule contains economically useful information.

Overall, the net results show that the value of distribution timing survives implementation costs for the unconstrained, dollar-neutral, and long-only investors at both levels of risk aversion, while the evidence is weaker under the gross-only mandate. The size of the gross-to-net adjustment is largest when exposure and trading intensity are high, but the main economic conclusion remains strongest for the unconstrained and dollar-neutral strategies, survives on a smaller scale for the long-only strategy, and is more sensitive to implementation costs under the gross-only mandate. The final assessment therefore depends jointly on the information value documented above and on the investment mandate through which that information is implemented.

5 Conclusion

This paper examines whether time variation in higher-order conditional distributional characteristics contains economically valuable information for currency portfolio choice.

Currency carry provides a particularly clean setting because under the maintained no-predictability assumption, the conditional mean of currency excess returns is tied to an observable forward-implied interest-rate differential. The portfolio problem can consequently isolate whether the information contained in conditional higher order moments and their associated co-moments improves investment decisions beyond mean and variance.

We combine a flexible conditional-distribution model that allows volatility, asymmetry, and tail thickness of exchange-rate returns to evolve over time, with a fourth-order approximation to CRRA expected utility and compare mean-variance and mean-variance-skewness-kurtosis allocation rules. The comparison is conducted under alternative investment mandates that constrain different margins of portfolio choice: total exposure, net directional exposure, and short selling. This design allows us to distinguish the economic value of higher-order distributional information from the investor's ability to act on that information.

The central finding is that the economic value of conditioning on higher-order moments is mandate-dependent. Relative to mean-variance, the allocation rule which also includes the third and the fourth moments produces higher point estimates for all reported risk-adjusted performance measures across the four mandates and the two levels of risk aversion considered. The largest and most statistically robust gains arise under the unconstrained mandate, where the investor can adjust both the composition and the scale of currency exposure. Median realized gross leverage is substantial in this frictionless benchmark, highlighting the importance of distinguishing the information value of the strategy from the constraints required for implementation. When gross exposure is fixed, the gains are smaller and their statistical strength varies across mandates. The dollar-neutral mandate provides the weakest statistical evidence, while the long-only mandate shows that economically meaningful gains can still arise when both leverage and short selling are prohibited, particularly for the more risk-averse investor.

The mandate comparisons also provide evidence on how higher-order information is monetized. The largest differences between the four- and the two-moments allocation rules arise when the investor can adjust total exposure, suggesting that exposure scaling is an important channel through which conditional distributional information is converted into economic value. Positive but smaller gains under fixed gross-exposure mandates indicate that cross-sectional reallocations retain some value even when leverage adjustment is removed. More generally, the results show that higher-order moments conditioning does not improve every realized feature of the return distribution in isolation. Rather, the investor exploits the joint conditional distribution through

the fourth-order expected-utility objective, with the resulting trade-off among expected return, dispersion, asymmetry, and tail thickness shaped by the feasible set and risk aversion.

Benchmark comparisons place these gains in a broader investment context. The higher-moment portfolios are economically competitive with conventional currency strategies tailored to the corresponding mandates and achieve competitive risk-adjusted performance relative to traditional equity, bond, commodity, and balanced portfolios on several risk-adjusted dimensions, especially under the unconstrained and dollar-neutral mandates. The correlation analysis further shows that the optimized currency strategies have markedly different relationships with conventional asset-class exposures across mandates. These results suggest that the economic role of higher-order moments in currency allocation extends beyond improving the standalone return-risk profile and can also reflect a distinct pattern of portfolio exposure.

The transaction-cost analysis provides an important implementation perspective. Higher-order moments conditioning does not materially increase turnover relative to the mean-variance benchmark, but the economic impact of implementation costs depends on the scale of the positions and the investment mandate. The advantage of the mean-variance-skewness-kurtosis allocation survives the baseline transaction-cost adjustment for the unconstrained, dollar-neutral, and long-only investors at both levels of risk aversion, whereas the gross-only mandate is more sensitive to implementation costs. Additional robustness checks show that the main comparisons remain stable under higher transaction-cost assumptions and alternative gross-exposure bounds.

The broader lesson is that the economic value of conditional distribution information in currency markets is both positive and mandate-dependent. Higher-order moments are valuable not simply because they describe non-Gaussian currency returns, but because investors can use their time variation to alter portfolio decisions. The extent to which this information translates into welfare gains depends on the feasible set through which the investor can implement it.

An important direction for future research is to examine how the economic value of distributional timing varies across alternative portfolio decision horizons and implementation frequencies. A complementary extension is to investigate how optimized currency portfolios relate to broader currency and multi-asset risk factors. Evaluating factor-adjusted performance, conducting portfolio-spanning tests, and examining interactions with dollar, carry, value, momentum, volatility, and funding-liquidity factors could help clarify the sources of the performance and diversification patterns documented here.

References

Ackermann, F., Pohl, W., and Schmedders, K. (2017). Optimal and Naive Diversification in Currency Markets. *Management Science*, 63:3347–3360.

- Andrews, D. W. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, 59:817–858.
- Ardia, D. and Boudt, K. (2015). Testing Equality of Modified Sharpe Ratios. *Finance Research Letters*, 13:97–104.
- Bacon, C. R. (2023). *Practical Portfolio Performance Measurement and Attribution*. John Wiley & Sons Ltd.
- Best, M. J. and Grauer, R. R. (1991). On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results. *The Review of Financial Studies*, 4:315–342.
- Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31:307–327.
- Brunnermeier, M. K., Nagel, S., and Pedersen, L. H. (2008). Carry Trades and Currency Crashes. *NBER Macroeconomics Annual*, 23:313–348.
- Cai, M., Del Negro, M., Herbst, E., Matlin, E., Sarfati, R., and Schorfheide, F. (2021). Online Estimation of DSGE Models. *The Econometrics Journal*, 24:C33–C58.
- Chopra, V. K., Ziemba, W. T., et al. (1993). The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice. *Journal of Portfolio Management*, 19:6–11.
- Darvas, Z. (2009). Leveraged Carry Trade Portfolios. *Journal of Banking & Finance*, 33:944–957.
- Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo Samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68:411–436.
- Della Corte, P., Sarno, L., and Tsiakas, I. (2009). An Economic Evaluation of Empirical Exchange Rate Models. *The Review of Financial Studies*, 22:3491–3530.
- DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial studies*, 22:1915–1953.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, 13:253–263.
- Engle, R. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, 20:339–350.

- Favre, L. and Galeano, J.-A. (2002). Mean-Modified Value-at-Risk Optimization with Hedge Funds. *Journal of Alternative Investments*, 5:21–25.
- Fleming, J., Kirby, C., and Ostdiek, B. (2001). The Economic Value of Volatility Timing. *The Journal of Finance*, 56:329–352.
- Gao, X. and Nardari, F. (2018). Do Commodities Add Economic Value in Asset Allocation? New Evidence from Time-Varying Moments. *Journal of Financial and Quantitative Analysis*, 53:365–393.
- Ghalanos, A., Rossi, E., and Urga, G. (2015). Independent Factor Autoregressive Conditional Density Model. *Econometric Reviews*, 34:594–616.
- Gregoriou, G. N. and Gueyie, J.-P. (2003). Risk-Adjusted Performance of Funds of Hedge Funds Using a Modified Sharpe Ratio. *The Journal of Wealth Management*, 6:77–83.
- Guidolin, M. and Timmermann, A. (2008). International Asset Allocation Under Regime Switching, Skew, and Kurtosis Preferences. *The Review of Financial Studies*, 21:889–935.
- Hansen, B. E. (1994). Autoregressive Conditional Density Estimation. *International Economic Review*, 35:705–730.
- Harvey, C. R., Liechty, J. C., Liechty, M. W., and Müller, P. (2010). Portfolio Selection with Higher Moments. *Quantitative Finance*, 10:469–485.
- Harvey, C. R. and Siddique, A. (1999). Autoregressive Conditional Skewness. *Journal of financial and quantitative analysis*, 34:465–487.
- Herfindahl, O. C. (1997). *Concentration in the Steel Industry*. Columbia university.
- Hirschman, A. O. (1980). *National Power and the Structure of Foreign Trade*, volume 105. Univ of California Press.
- Jondeau, E. and Rockinger, M. (2003). Conditional Volatility, Skewness, and Kurtosis: Existence, Persistence, and Comovements. *Journal of Economic Dynamics and Control*, 27:1699–1737.
- Jondeau, E. and Rockinger, M. (2006). Optimal Portfolio Allocation under Higher Moments. *European Financial Management*, 12:29–55.
- Jondeau, E. and Rockinger, M. (2009). The Impact of Shocks on Higher Moments. *Journal of Financial Econometrics*, 7:77–105.

- Jondeau, E. and Rockinger, M. (2012). On the Importance of Time Variability in Higher Moments for Asset Allocation. *Journal of Financial Econometrics*, 10:84–123.
- Keating, C. and Shadwick, W. F. (2002). A Universal Performance Measure. *Journal of Performance Measurement*, 6:59–84.
- Le, T. H. (2021). International Portfolio Allocation: The Role of Conditional Higher Moments. *International Review of Economics & Finance*, 74:33–57.
- Ledoit, O. and Wolf, M. (2008). Robust Performance Hypothesis Testing with the Sharpe Ratio. *Journal of Empirical Finance*, 15:850–859.
- Li, D., Clements, A., and Drovandi, C. (2021). Efficient Bayesian estimation for GARCH-type Models via Sequential Monte Carlo. *Econometrics and Statistics*, 19:22–46.
- Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common Risk Factors in Currency Markets. *The Review of Financial Studies*, 24:3731–3777.
- Lyons, R. K. (2001). *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge, MA.
- Markowitz, H. (1959). Portfolio Selection: Efficient Diversification of Investments. *Journal of Finance*, 14.
- Meese, R. A. and Rogoff, K. (1983). Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? *Journal of International Economics*, 14:3–24.
- Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2012). Carry Trades and Global Foreign Exchange Volatility. *The Journal of Finance*, 67:681–718.
- Packer, F., Schrimpf, A., Sushko, V., and Zarra, N. (2024). Hedge Fund Exposure to the Carry Trade. *BIS Quarterly Review*. Box extracted from the Overview chapter “Carry off, carry on”.
- Pezier, J. and White, A. (2006). The Relative Merits of Investable Hedge Fund Indices and of Funds of Hedge Funds in Optimal Passive Portfolios. ICMA Centre Discussion Papers in Finance icma-dp2006-10, Henley Business School, University of Reading.
- Politis, D. N. and Romano, J. P. (1992). A Circular Block-Resampling Procedure for Stationary Data. In LePage, R. and Billard, L., editors, *Exploring the Limits of Bootstrap*, pages 263–270. Wiley, New York.
- Rossi, B. (2013). Exchange Rate Predictability. *Journal of Economic Literature*, 51:1063–1119.

- Sahu, S., Dey, D., and Branco, M. (2003). A New Class of Multivariate Skew Distributions with Applications to Bayesian Regression Models. *Canadian Journal of Statistics*, 31:129–150.
- Scott, R. C. and Horvath, P. A. (1980). On the Direction of Preference for Moments of Higher Order Than the Variance. *Journal of Finance*, 35:915–919.
- Sharpe, W. F. (1966). Mutual Fund Performance. *The Journal of Business*, 39:119–138.
- van der Meer, R., Sortino, F., and Plantinga, A. (2001). The Impact of Downside Risk on Risk-Adjusted Performance of Mutual Funds in the Euronext Markets. *Available at SSRN 277352*.

A Performance Measures

This appendix defines the performance measures used in the empirical analysis. For a given strategy p , let y_t^p be the realized portfolio excess return (either gross or net of transaction costs).

Denote by $\mathcal{H}$ the number of rebalancing periods per year. The sample mean and standard deviation of $\{y_t^p\}_{t=1}^T$ are

$$\mu_p = \frac{1}{T} \sum_{t=1}^T y_t^p, \quad \sigma_p = \left[\frac{1}{T} \sum_{t=1}^T (y_t^p - \mu_p)^2 \right]^{1/2}.$$

When reported in the tables, the mean excess return and volatility are annualized as $\text{ER}_p = \mathcal{H}\mu_p$, and $\text{Vol}_p = \sqrt{\mathcal{H}}\sigma_p$.

Realized distributional moments. The Semi-Standard Deviation (SSD) measures downside dispersion around the sample mean

$$\text{SSD}_p = \left[\frac{1}{T} \sum_{t=1}^T \min(y_t^p - \mu_p, 0)^2 \right]^{1/2}.$$

In the empirical tables, SSD_p is annualized by multiplying by $\sqrt{\mathcal{H}}$. The realized Skewness (SK) and Kurtosis (KU) are

$$\text{SK}_p = \frac{1}{T} \sum_{t=1}^T \left(\frac{y_t^p - \mu_p}{\sigma_p} \right)^3, \quad \text{KU}_p = \frac{1}{T} \sum_{t=1}^T \left(\frac{y_t^p - \mu_p}{\sigma_p} \right)^4.$$

SK_p and KU_p are reported in the tables at portfolio return frequency to avoid vanishing effects.

Sharpe-type and utility-based measures. The Sharpe Ratio (SR) is

$$\text{SR}_p = \frac{\mu_p}{\sigma_p}.$$

Because the SR is based only on the first two moments, we also report the adjusted Sharpe Ratio (aSR) of [Pezier and White \(2006\)](#), as described in [Bacon \(2023\)](#)

$$\text{aSR}_p = \text{SR}_p \left[1 + \frac{\text{SK}_p}{6} \text{SR}_p - \frac{\text{KU}_p - 3}{24} \text{SR}_p^2 \right].$$

This adjustment rewards positive skewness and penalizes excess kurtosis.

The per-period performance fee (Fleming et al., 2001) is the constant return Φ that equates the average CRRA utility of the portfolio a , to that of the benchmark b

$$\frac{1}{T} \sum_{t=1}^T \frac{(1 + y_t^a - \Phi)^{1-\gamma}}{1-\gamma} = \bar{U}_b(\gamma), \quad \bar{U}_b(\gamma) = \frac{1}{T} \sum_{t=1}^T \frac{(1 + y_t^b)^{1-\gamma}}{1-\gamma}.$$

A positive Φ is the maximum fee an investor with risk aversion γ would pay to access portfolio a in place of portfolio b . The equation above is solved numerically, then the performance fee is annualized by multiplying by the factor $\mathcal{H}$

Downside and tail-risk measures. The modified Sharpe Ratio (mSR) replaces volatility with a Cornish-Fisher (CF) tail-risk adjustment. Let z_q denote the q -quantile of the standard Normal distribution. The CF adjusted quantile is

$$z_{q,p}^{CF} = z_q + \frac{1}{6}(z_q^2 - 1)SK_p + \frac{1}{24}(z_q^3 - 3z_q)(KU_p - 3) - \frac{1}{36}(2z_q^3 - 5z_q)SK_p^2.$$

The corresponding modified VaR (mVaR) is

$$\text{mVaR}_{q,p} = -(\mu_p + \sigma_p z_{q,p}^{CF}),$$

and the mSR is

$$\text{mSR}_{q,p} = \frac{\mu_p}{\text{mVaR}_{q,p}}.$$

This measure is sensitive to the effect of skewness and excess kurtosis on left-tail risk.

The Omega ratio (Ω -ratio) compares gains and losses relative to a target return τ

$$\Omega_p(\tau)\text{-ratio} = \frac{\sum_{t=1}^T \max(y_t^p - \tau, 0)}{\sum_{t=1}^T \max(\tau - y_t^p, 0)}.$$

The lower partial moment of order n around the same target is

$$\text{LPM}_{n,p}(\tau) = \frac{1}{T} \sum_{t=1}^T \max(\tau - y_t^p, 0)^n.$$

The Kappa ratio ($\mathcal{K}$ -ratio) of order n is

$$\mathcal{K}_{n,p}(\tau) - \text{ratio} = \frac{\mu_p - \tau}{\text{LPM}_{n,p}(\tau)^{1/n}}.$$

The Sortino ratio ($\mathcal{S}$ -ratio) is the special case $\mathcal{K}_{2,p}(\tau)$, typically evaluated at $\tau = 0$.

Implementation measures. Let $w_{i,t}$ be the target portfolio weight in currency i chosen at rebalancing date t , the average Turnover (TO) is

$$\text{TO}_p = \frac{1}{T-1} \sum_{t=1}^{T-1} \sum_{i=1}^N |w_{i,t+1} - w_{i,t}|.$$

TO measures the strategy's trading intensity and is the quantity that links the portfolio rule to spot rebalancing costs.

Portfolio concentration is measured by the normalized Herfindahl-Hirschman Index (HHI) evaluated conditional on each mandate's exposure scale

$$\text{HHI}_p = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N |\bar{w}_{i,t}|^2, \quad \bar{w}_{i,t} = \frac{w_{i,t}}{\sum_{j=1}^N |w_{j,t}|},$$

such that $\sum_{i=1}^N \bar{w}_{i,t}$ and $0 < \text{HHI}_p < 1$. The normalized HHI measures the concentration of gross currency exposure independently of its overall scale, making the dispersion comparable across portfolios with different levels of leverage. Gross exposure captures the extent of leverage, whereas the HHI captures how that exposure is distributed across portfolio constituents. In our analysis this distinction is important for the unconstrained mandate, for whom the scale of exposure can vary endogenously with the conditional distribution, whereas it becomes irrelevant for the other mandates (gross exposure constrained to unity). A higher value indicates a more concentrated allocation, while a lower value indicates a more diversified allocation across currencies.

Reporting conventions. Mean excess returns and performance fees are annualized using the factor $\mathcal{H}$. σ_p , SSD_p , SR_p , aSR_p , mSR_p , and $\mathcal{S}$ -ratio are annualized using the factor $\sqrt{\mathcal{H}}$. SK_p , KU_p , Ω -ratio, TO, and HHI are reported at the portfolio return frequency unless otherwise stated. In the empirical tables, the default target for Ω - and $\mathcal{K}$ -ratios is $\tau = 0$.

B Bayesian Estimation

This appendix details the Bayesian estimation procedure. It describes the SMC sampler used for the initial estimation window, the online SMC update used in the recursive exercise, the adaptive tempering schedule, and the main implementation details.

B.1 Likelihood and Posterior

Let $\mathbf{x}_{1:T} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ be the vector of observed exchange-rate changes. The parameter vector is denoted by Θ and collects the GARCH coefficients, the DCC coefficients, and the coefficients in the dynamics of η_t and λ_t . The likelihood is

$$\mathcal{L}(\mathbf{x}_{1:T}|\Theta) = \prod_{t=1}^T g(\mathbf{x}_t|\mathbf{H}_t(\Theta), \eta_t(\Theta), \lambda_t(\Theta)),$$

where $g(\cdot)$ is the product of the standardized Skewed Student- t (Sk- t) marginal densities defined in Section 3.1.

Let $p(\Theta)$ be the prior density, the posterior and marginal likelihood are

$$p(\Theta|\mathbf{x}_{1:T}) = \frac{\mathcal{L}_T(\mathbf{x}_{1:T}|\Theta)p(\Theta)}{Z_T}, \quad Z_T = \int \mathcal{L}_T(\mathbf{x}_{1:T}|\Theta)p(\Theta)d\Theta.$$

B.2 Sequential Monte Carlo (SMC) Sampler

The SMC sampler approximates the posterior through a sequence of tempered bridge distributions, following [Del Moral et al. \(2006\)](#),

$$p_n(\Theta|\mathbf{x}_{1:T}) \propto \mathcal{L}(\mathbf{x}_{1:T}|\Theta)^{\phi_n} p(\Theta), \quad 0 = \phi_1 < \phi_2 < \dots < \phi_{N_\phi} = 1. \quad (15)$$

The sampler starts from the prior distribution at $\phi_1 = 0$ and progressively introduces the likelihood until the posterior is reached at $\phi_{N_\phi} = 1$. At each stage, the target distribution is represented by a weighted particle system $\{\Theta_n^i, \bar{v}_n^i\}_{i=1}^M$, where M is the number of particles.

A common deterministic tempering schedule is

$$\phi_n = \left(\frac{n-1}{N_\phi-1} \right)^\lambda,$$

where $\lambda > 0$ controls the speed at which likelihood information is incorporated. The case $\lambda = 1$ gives a linear schedule, while $\lambda > 1$ assigns smaller likelihood increments to the early stages of

the bridge. In the empirical implementation, we use the adaptive tempering rule described in Section B.4.

Each SMC stage consists of three steps: correction, selection, and mutation. The correction step updates particle weights when moving from $p_{n-1}(\cdot)$ to $p_n(\cdot)$; the selection step resamples particles when weight degeneracy becomes severe; and the mutation step rejuvenates the particle system using a Random-Walk Metropolis-Hastings (RWMH) kernel invariant to the current tempered target. The mutation kernel is indexed by the tuning vector

$$\zeta_n = [c_n, \text{vech}(\Sigma_n)],$$

where Σ_n is the proposal covariance matrix and c_n is a scale parameter. In the implementation, Σ_n is estimated from the current particle approximation, while c_n is updated recursively using past rejection rates. SMC methods are particularly convenient in this setting because they provide posterior approximations and marginal likelihood estimates in a single run; see [Li et al. \(2021\)](#) for GARCH models. The full recursion, including the weight update, resampling rule, mutation step, posterior moment approximation, and marginal likelihood estimate, is reported in Algorithm 1.

B.3 Online Sequential Monte Carlo

The recursive portfolio exercise requires updating the posterior whenever the estimation window rolls forward. Re-estimating the model from the prior at each rebalancing date would be computationally costly. We therefore use an Online SMC (OSMC) update in the spirit of [Cai et al. \(2021\)](#), which transforms the particle approximation of the posterior based on the old window into a particle approximation of the posterior based on the new window.

Let $\mathcal{T}_t = \{1, \dots, T\}$ denote the old estimation window and $\mathcal{T}_{t'} = \{2, \dots, T+1\}$ the new window. When the window rolls forward, the first observation leaves the sample and a new observation enters. The online bridge is

$$p_n(\Theta | \mathbf{x}_{\mathcal{T}_{t'}}) \propto \mathcal{L}(\mathbf{x}_{\mathcal{T}_{t'}} | \Theta)^{\phi_n} \mathcal{L}(\mathbf{x}_{\mathcal{T}_t} | \Theta)^{1-\phi_n} p(\Theta), \quad 0 = \phi_1 < \dots < \phi_{N_\phi} = 1. \quad (16)$$

This bridge differs from the initial SMC bridge in (15). There, tempering gradually introduces the likelihood starting from the prior. Here, tempering gradually replaces the likelihood of the old window with the likelihood of the new window. At $\phi_1 = 0$, the target is the posterior based on $\mathcal{T}_t$; at $\phi_{N_\phi} = 1$, the target is the posterior based on $\mathcal{T}_{t'}$.

The correction, selection, and mutation steps are the same as in the standard SMC sampler, but the incremental weights now depend on the likelihood ratio between the new and old windows. The full recursion, including the marginal likelihood update, is reported in Algorithm 2.

Algorithm 1 Sequential Monte Carlo (SMC)

Initialization ($\phi_1 = 0$): Draw the initial particles from the prior distribution $\Theta_1^i \sim p(\Theta)$, for $i = 1, \dots, M$, and assign uniform weights

$$\Theta_1^i \sim p(\Theta), \quad \mathcal{V}_1^i = \tilde{\mathcal{V}}_1^i = \frac{1}{M}, \quad i = 1, \dots, M.$$

for $n = 2$ **to** N_ϕ **do**

1. **Correction:** Adjust the weights of the particles from stage $(n-1)$

$$\tilde{\nu}_n^i = [\mathcal{L}(\mathbf{x}_{1:T} | \Theta_{n-1}^i)]^{\phi_n - \phi_{n-1}}, \quad \tilde{\mathcal{V}}_n^i = \frac{\tilde{\nu}_n^i \mathcal{V}_{n-1}^i}{\frac{1}{M} \sum_{j=1}^M \tilde{\nu}_n^j \mathcal{V}_{n-1}^j}, \quad i = 1, \dots, M.$$

2. **Selection:** Calculate the Effective Sample Size (ESS)

$$ESS_n = \frac{M}{\frac{1}{M} \sum_{i=1}^M (\tilde{\mathcal{V}}_n^i)^2}.$$

if $ESS_n < M/2$ **then**

Resample the particles via systematic resampling and reinitialize the weights

$$\tilde{\mathcal{V}}_n^i = \frac{1}{M}, \quad \hat{\Theta}_n^i \sim \{\Theta_{n-1}^i, \tilde{\mathcal{V}}_n^i\}, \quad i = 1, \dots, M.$$

else

Keep the particles and weights

$$\mathcal{V}_n^i = \tilde{\mathcal{V}}_n^i, \quad \hat{\Theta}_n^i = \Theta_{n-1}^i, \quad i = 1, \dots, M.$$

end if

3. **Mutation:** Propagate each particle $\{\hat{\Theta}_n^i, \mathcal{V}_n^i\}$ via one RWMH step with transition density $\Theta_n^i \sim K_n(\Theta_n | \hat{\Theta}_{n-1}^i; \zeta_n)$ and stationary distribution $p_n(\Theta | \mathbf{x}_{1:T})$.

end for

Compute posterior moments

$$\mathbb{E}_{\pi_n}[h(\Theta)] = \sum_{i=1}^M \bar{\mathcal{V}}_n^i h(\Theta_n^i), \quad \bar{\mathcal{V}}_n^i = \frac{\mathcal{V}_n^i}{\sum_{i=1}^M \mathcal{V}_n^i}.$$

Marginal likelihood approximation

$$\hat{p}(\mathbf{x}_{1:T}) = \prod_{n=2}^{N_\phi} \left(\frac{1}{M} \sum_{i=1}^M \tilde{\nu}_n^i \mathcal{V}_{n-1}^i \right).$$

Algorithm 2 Online Sequential Monte Carlo (OSMC)

Initialization: ($\phi_1 = 0$): Draw the initial particles from the posterior distribution based on the previous window $p_{n,t}(\Theta | \mathbf{x}_{\mathcal{T}_t})$, for $i = 1, \dots, M$, and assign uniform weights

$$\Theta_1^i \sim p_{n,t}(\Theta | \mathbf{x}_{\mathcal{T}_t}), \quad \mathcal{V}_1^i = \tilde{\mathcal{V}}_1^i = \frac{1}{M}, \quad i = 1, \dots, M.$$

for $n = 2$ **to** N_ϕ **do**

1. **Correction:** Adjust the weights of the particles from stage $(n-1)$

$$\tilde{\nu}_n^i = \left[\frac{\mathcal{L}(\mathbf{x}_{\mathcal{T}_{t'}} | \Theta_{n-1}^i)}{\mathcal{L}(\mathbf{x}_{\mathcal{T}_t} | \Theta_{n-1}^i)} \right]^{\phi_n - \phi_{n-1}}, \quad \tilde{\mathcal{V}}_n^i = \frac{\tilde{\nu}_n^i \mathcal{V}_{n-1}^i}{\frac{1}{M} \sum_{j=1}^M \tilde{\nu}_n^j \mathcal{V}_{n-1}^j}, \quad i = 1, \dots, M.$$

2. **Selection:** (equal to Algorithm 1)

3. **Mutation:** Propagate each particle $\{\hat{\Theta}_n^i, \mathcal{V}_n^i\}$ via one RWMH step with transition density $\Theta_n^i \sim K_n(\Theta_n | \hat{\Theta}_{n-1}^i; \zeta_n)$ and stationary distribution $\tilde{\pi}_n(\Theta)$.

end for

Compute posterior moments: (equal to Algorithm 1)

Marginal likelihood approximation:

$$\hat{p}(\mathbf{x}_{\mathcal{T}_{t'}}) = \hat{p}(\mathbf{x}_{\mathcal{T}_{t'}} | \mathbf{x}_{\mathcal{T}_t}) \cdot \hat{p}(\mathbf{x}_{\mathcal{T}_t}) = \prod_{n=2}^{N_\phi} \left(\frac{1}{M} \sum_{i=1}^M \tilde{\nu}_n^i \mathcal{V}_{n-1}^i \right) \cdot \hat{p}(\mathbf{x}_{\mathcal{T}_t}),$$

B.4 Adaptive Tempering Schedule

The number of tempering steps required by the online update depends on how different the old and new posteriors are. When the rolling-window update is small, a large increase in ϕ_n can be used without inducing severe weight degeneracy. When the new observation substantially changes the likelihood, smaller increments are needed. We therefore choose the tempering schedule adaptively by controlling the Effective Sample Size (*ESS*).

To describe the rule, consider the OSMC bridge in (16). Given the current value ϕ_{n-1} , define for any candidate $\phi \in (\phi_{n-1}, 1]$

$$\tilde{v}^i(\phi) = \left[\frac{\mathcal{L}_{t'}(\theta_{n-1}^i)}{\mathcal{L}_t(\theta_{n-1}^i)} \right]^{\phi - \phi_{n-1}}, \quad \tilde{\mathcal{V}}^i(\phi) = \frac{\tilde{v}^i(\phi) \mathcal{V}_{n-1}^i}{\frac{1}{M} \sum_{j=1}^M \tilde{v}^j(\phi) \mathcal{V}_{n-1}^j}, \quad \widehat{ESS}(\phi) = \frac{M}{\frac{1}{M} \sum_{i=1}^M \left(\tilde{\mathcal{V}}^i(\phi) \right)^2}.$$

The next tempering value is chosen so that the predicted ESS after reweighting is a fixed fraction of the previous *ESS*

$$f(\phi) = \widehat{ESS}(\phi) - \alpha \widehat{ESS}_{n-1} = 0,$$

where $\alpha \in (0, 1)$ controls the targeted reduction in *ESS*. A value of α close to one produces smaller tempering increments and therefore a more conservative transition between bridge distributions.

Algorithm 3 Adaptive Tempering Schedule

- (1) If $f(1) \geq 0$, set $\phi_n = 1$.
 - (2) If $f(1) < 0$, choose $\phi_n \in (\phi_{n-1}, 1)$ such that $f(\phi_n) = 0$.
-

In practice, the root may be very close to the lower bound when the likelihood ratio is highly variable. To make the search more stable, we use a grid of candidate upper bounds and solve the root problem locally.

Algorithm 4 Adaptive Tempering Schedule with Incremental Upper Bounds

- (1) Create a grid of candidate values $\{\phi_j\}_{j=1}^{N_\phi}$ on $[\phi_{n-1}, 1]$.
 - (2) Find the first grid point ϕ_{j+1} such that $f(\phi_{j+1}) < 0$.
 - (3) If no such point exists, set $\phi_n = 1$.
 - (4) Otherwise, choose $\phi_n \in (\phi_j, \phi_{j+1})$ such that $f(\phi_n) = 0$.
-

The same adaptive rule can be applied to the standard SMC sampler by replacing the likelihood ratio in the incremental weights with the likelihood increment implied by the bridge in (15). The resulting tempering schedule is monotone by construction and terminates when $\phi_n = 1$.

Finally, we check the estimation performance of SMC and OSMC in Table 9. The Table reports the posterior mean and 90% credible interval, $[5\%, 95\%]$, of the model parameters at the initial batch estimation and at the final online update, for the nine currencies included in the asset allocation exercise. Comparing the initial and final columns shows that the parameters are not static. Several volatility and shape parameters shift meaningfully over the sample, which is the recursive-updating counterpart of the time-varying conditional moments exploited by the MVSK strategy in the main text.

Table 9: Estimated parameters: initial batch and final online update

		CHF	EUR	JPY	GBP	AUD	CAD	NOK	SEK	SGD
ω	Initial	0.013 [0.009, 0.018]	0.009 [0.007, 0.012]	0.024 [0.017, 0.031]	0.006 [0.003, 0.010]	0.020 [0.009, 0.033]	0.005 [0.002, 0.009]	0.014 [0.010, 0.019]	0.021 [0.014, 0.029]	0.003 [0.002, 0.005]
	Final	0.000 [†] [0.000, 0.000]	0.000 [†] [0.000, 0.000]	0.004 [0.003, 0.006]	0.008 [0.004, 0.012]	0.007 [0.006, 0.010]	0.003 [0.001, 0.004]	0.013 [0.009, 0.017]	0.007 [0.004, 0.010]	0.002 [0.001, 0.002]
α	Initial	0.047 [0.038, 0.056]	0.048 [0.041, 0.055]	0.078 [0.061, 0.094]	0.037 [0.028, 0.046]	0.047 [0.037, 0.059]	0.054 [0.043, 0.065]	0.061 [0.051, 0.074]	0.054 [0.044, 0.066]	0.055 [0.042, 0.068]
	Final	0.018 [0.013, 0.023]	0.026 [0.021, 0.031]	0.052 [0.040, 0.065]	0.050 [0.037, 0.064]	0.031 [0.025, 0.037]	0.036 [0.029, 0.043]	0.059 [0.048, 0.071]	0.034 [0.026, 0.043]	0.041 [0.033, 0.049]
β	Initial	0.929 [0.914, 0.942]	0.932 [0.921, 0.942]	0.871 [0.846, 0.895]	0.945 [0.929, 0.958]	0.923 [0.896, 0.946]	0.924 [0.900, 0.943]	0.914 [0.898, 0.926]	0.906 [0.886, 0.925]	0.912 [0.891, 0.934]
	Final	0.981 [0.976, 0.986]	0.973 [0.968, 0.977]	0.934 [0.917, 0.950]	0.925 [0.900, 0.946]	0.948 [0.937, 0.958]	0.949 [0.937, 0.961]	0.918 [0.901, 0.934]	0.949 [0.934, 0.962]	0.937 [0.922, 0.949]
γ_0	Initial	-2.076 [-2.865, -1.336]	-1.473 [-2.155, -0.788]	-1.402 [-2.075, -0.620]	-0.246 [-2.627, 1.735]	-2.616 [-3.374, -1.902]	2.391 [1.286, 3.531]	-0.334 [-1.905, 1.254]	-2.495 [-3.391, -1.673]	-2.336 [-3.262, -1.460]
	Final	-2.800 [-4.100, -1.617]	-2.015 [-3.081, -0.958]	-3.139 [-4.712, -1.666]	-3.738 [-5.174, -2.321]	-1.755 [-2.711, -0.808]	-1.765 [-2.881, -0.647]	-2.379 [-3.367, -1.409]	-2.594 [-3.549, -1.638]	-2.399 [-3.759, -1.223]
γ_1^+	Initial	-0.062 [-1.093, 1.025]	-0.283 [-1.543, 0.628]	-2.009 [-2.883, -1.196]	0.385 [-0.566, 1.420]	-0.011 [-1.212, 0.925]	0.168 [-1.000, 1.306]	-0.748 [-1.444, -0.090]	0.029 [-0.688, 0.679]	0.278 [-0.819, 1.362]
	Final	-0.311 [-1.300, 0.538]	0.141 [-1.016, 1.213]	-1.269 [-3.362, 0.710]	1.806 [0.640, 3.365]	0.023 [-0.780, 0.836]	0.267 [-0.881, 1.315]	1.397 [0.272, 2.635]	1.122 [-0.279, 2.763]	0.465 [-0.431, 1.425]
γ_1^-	Initial	-0.098 [-0.616, 0.449]	0.609 [-0.275, 1.588]	-0.724 [-1.405, -0.083]	-2.174 [-4.280, -0.352]	0.049 [-0.955, 0.998]	-1.125 [-2.665, 0.341]	-1.366 [-2.393, -0.410]	1.277 [0.311, 2.417]	-0.374 [-2.022, 1.013]
	Final	0.089 [-0.763, 0.915]	0.242 [-0.580, 1.124]	2.052 [0.739, 3.416]	0.429 [-1.526, 2.202]	0.852 [-0.289, 2.387]	1.255 [0.049, 2.559]	0.504 [-0.575, 1.578]	1.061 [-0.110, 2.469]	-0.305 [-1.770, 1.020]
γ_2	Initial	0.109 [-0.092, 0.293]	0.342 [0.151, 0.539]	-0.181 [-0.338, 0.003]	-0.367 [-0.667, -0.032]	0.158 [-0.036, 0.379]	0.297 [0.070, 0.521]	-0.065 [-0.316, 0.178]	-0.058 [-0.199, 0.074]	0.224 [0.024, 0.433]
	Final	-0.072 [-0.450, 0.316]	0.010 [-0.379, 0.448]	0.389 [0.029, 0.723]	0.227 [-0.063, 0.514]	-0.051 [-0.411, 0.329]	0.015 [-0.320, 0.381]	-0.114 [-0.512, 0.243]	0.172 [-0.144, 0.482]	-0.035 [-0.382, 0.322]
κ_0	Initial	0.302 [0.192, 0.428]	-0.092 [-0.217, 0.037]	0.078 [-0.088, 0.247]	-0.227 [-0.353, -0.093]	0.013 [-0.113, 0.143]	0.043 [-0.111, 0.201]	0.072 [-0.048, 0.192]	0.109 [0.007, 0.219]	0.059 [-0.048, 0.174]
	Final	-0.136 [-0.275, 0.001]	-0.164 [-0.338, 0.031]	0.119 [0.014, 0.238]	-0.206 [-0.329, -0.085]	0.011 [-0.107, 0.141]	0.059 [-0.125, 0.243]	-0.065 [-0.210, 0.057]	0.072 [-0.059, 0.194]	0.082 [-0.042, 0.208]
κ_1^+	Initial	0.109 [-0.061, 0.278]	-0.047 [-0.191, 0.086]	-0.012 [-0.177, 0.159]	0.027 [-0.143, 0.199]	0.225 [0.087, 0.370]	0.028 [-0.137, 0.184]	0.271 [0.160, 0.377]	0.060 [-0.067, 0.189]	0.233 [0.076, 0.401]
	Final	-0.130 [-0.281, 0.018]	-0.165 [-0.271, -0.063]	0.051 [-0.080, 0.187]	-0.121 [-0.284, 0.067]	0.042 [-0.103, 0.181]	0.044 [-0.143, 0.242]	0.135 [0.010, 0.259]	0.048 [-0.119, 0.192]	0.120 [-0.023, 0.246]
κ_1^-	Initial	-0.011 [-0.094, 0.081]	0.062 [-0.075, 0.207]	0.086 [-0.060, 0.230]	0.305 [0.108, 0.501]	-0.023 [-0.178, 0.125]	0.195 [-0.008, 0.394]	-0.399 [-0.575, -0.219]	-0.117 [-0.264, 0.020]	0.164 [0.048, 0.276]
	Final	0.138 [0.009, 0.273]	-0.072 [-0.177, 0.030]	0.038 [-0.170, 0.245]	0.253 [0.105, 0.412]	0.048 [-0.113, 0.203]	0.122 [-0.050, 0.291]	0.106 [-0.055, 0.261]	0.105 [-0.025, 0.227]	-0.051 [-0.205, 0.097]
κ_2	Initial	-0.257 [-0.430, -0.063]	0.169 [-0.034, 0.343]	-0.096 [-0.242, 0.068]	0.029 [-0.123, 0.180]	-0.064 [-0.211, 0.087]	-0.213 [-0.407, -0.013]	0.079 [-0.136, 0.276]	0.129 [-0.055, 0.314]	0.045 [-0.114, 0.193]
	Final	-0.191 [-0.571, 0.178]	-0.707 [-0.844, -0.524]	0.002 [-0.371, 0.394]	-0.025 [-0.313, 0.287]	0.108 [-0.317, 0.497]	-0.210 [-0.546, 0.178]	-0.045 [-0.366, 0.333]	-0.018 [-0.405, 0.358]	-0.035 [-0.408, 0.348]
<i>DCC Parameters</i>										
δ_1	Initial	0.013 [0.012, 0.015]								
	Final	0.011 [0.010, 0.012]								
δ_2	Initial	0.978 [0.975, 0.981]								
	Final	0.977 [0.973, 0.980]								

Notes: Posterior mean and [5%, 95%] credible interval of each parameter at the initial batch estimation (30-Dec-2005) and at the final online update (19-Sep-2025), for the nine currencies included in the asset allocation exercise: CHF, EUR, JPY, GBP, AUD, CAD, NOK, SEK, and SGD (against USD). Parameters are reported for returns expressed in percentage points. ω , α , and β are the currency-specific GARCH intercept, ARCH, and persistence parameters; δ_1 and δ_2 are parameters common to all currencies. γ_0 - γ_2 and κ_0 - κ_2 , the currency-specific loadings governing the dynamics of conditional skewness and kurtosis, respectively. [†]Rounds to 0.000 at three decimals; posterior mean is on the order of 10^{-6} (CHF) and 10^{-5} (EUR).

C Moments of the Portfolio Returns

This appendix reports the formulas used to construct the conditional moments and co-moments that enter the mean-variance-skewness-kurtosis (MVS-K) portfolio problem. The vector of exchange-rate returns is denoted by $\mathbf{x}_{t+1} = (x_{t+1}^1, \dots, x_{t+1}^N)'$, currency excess returns are $\mathbf{y}_{t+1} = \mathbf{d}_t + \mathbf{A}_t \mathbf{x}_{t+1}$, where $\mathbf{d}_t = (r_t^1 - r_t^d, \dots, r_t^N - r_t^d)'$ and $\mathbf{A}_t = \text{diag}(1 + r_t^1, \dots, 1 + r_t^N)$. The portfolio weights chosen at time t are $\mathbf{w}_t = (w_{1,t}, \dots, w_{N,t})'$, and the portfolio excess return is $y_{t+1}^p = \mathbf{w}_t' \mathbf{y}_{t+1}$. Throughout the appendix, $E_t[\cdot] = E[\cdot \mid \mathcal{F}_t]$.

C.1 Moments of the Skewed Student- t Innovations

The econometric model writes exchange-rate returns as

$$\mathbf{x}_{t+1} = \mathbf{B}_{t+1} \mathbf{z}_{t+1}, \quad \mathbf{B}_{t+1} \mathbf{B}_{t+1}' = \mathbf{H}_{t+1},$$

where $\mathbf{H}_{t+1}$ is the conditional covariance matrix of $\mathbf{x}_{t+1}$, and $\mathbf{z}_{t+1} = (z_{1,t+1}, \dots, z_{N,t+1})'$ has conditionally independent components. Each innovation $z_{i,t+1}$ follows the standardized skewed Student- t (Sk- t) density used in the main text, with degrees of freedom $\eta_{i,t+1}$ and asymmetry parameter $\lambda_{i,t+1}$. This parameterization imposes

$$E_t[z_{i,t+1}] = 0, \quad E_t[z_{i,t+1}^2] = 1, \quad i = 1, \dots, N.$$

For completeness, we provide the closed-form expressions for the raw moments of $z_{i,t+1}$. Let

$$c_{i,t+1} = \frac{\Gamma[(\eta_{i,t+1} + 1)/2]}{\sqrt{\pi(\eta_{i,t+1} - 2)} \Gamma(\eta_{i,t+1}/2)},$$

$$a_{i,t+1} = 4\lambda_{i,t+1} c_{i,t+1} \frac{\eta_{i,t+1} - 2}{\eta_{i,t+1} - 1}, \quad b_{i,t+1} = (1 + 3\lambda_{i,t+1}^2 - a_{i,t+1}^2)^{1/2}.$$

For any integer $q \geq 0$ such that $q < \eta_{i,t+1}$, define the positive half-moment of the corresponding symmetric standardized Sk- t kernel as

$$I_q(\eta_{i,t+1}) = \frac{(\eta_{i,t+1} - 2)^{q/2}}{2\sqrt{\pi}} \frac{\Gamma((q+1)/2) \Gamma((\eta_{i,t+1} - q)/2)}{\Gamma(\eta_{i,t+1}/2)}.$$

Then the m -th raw moment of $z_{i,t+1}$, for $m < \eta_{i,t+1}$, is

$$\begin{aligned} M_{i,t+1}^{(m)} &= \mathbb{E}_t[z_{i,t+1}^m] \\ &= b_{i,t+1}^{-m} \sum_{q=0}^m \binom{m}{q} (-a_{i,t+1})^{m-q} I_q(\eta_{i,t+1}) [(1 + \lambda_{i,t+1})^{q+1} + (-1)^q (1 - \lambda_{i,t+1})^{q+1}]. \end{aligned}$$

The constants $a_{i,t+1}$ and $b_{i,t+1}$ imply $M_{i,t+1}^{(1)} = 0$ and $M_{i,t+1}^{(2)} = 1$. The conditional third and fourth central moments of the standardized innovation are therefore

$$\begin{aligned} \mu_{i,t+1}^{(3)} &= M_{i,t+1}^{(3)} - 3M_{i,t+1}^{(1)}M_{i,t+1}^{(2)} + 2\left(M_{i,t+1}^{(1)}\right)^3, \\ \mu_{i,t+1}^{(4)} &= M_{i,t+1}^{(4)} - 4M_{i,t+1}^{(1)}M_{i,t+1}^{(3)} + 6M_{i,t+1}^{(2)}\left(M_{i,t+1}^{(1)}\right)^2 - 3\left(M_{i,t+1}^{(1)}\right)^4. \end{aligned}$$

Since innovations are standardized, these simplify to $\mu_{i,t+1}^{(3)} = M_{i,t+1}^{(3)}$ and $\mu_{i,t+1}^{(4)} = M_{i,t+1}^{(4)}$. The fourth moment exists when $\eta_{i,t+1} > 4$, which is imposed in the empirical implementation.

C.2 Co-Moments of Exchange-Rate Returns

Let $B_{ir,t+1}$ denote the element (i, r) of $\mathbf{B}_{t+1}$. Let also $S_{ijk,t+1}$ and $K_{ijk\ell,t+1}$ respectively the element (i, j, k) and (i, j, k, ℓ) of the tensors of third and fourth co-moments of $\mathbf{x}_{t+1}$. Because $\mathbf{x}_{t+1}$ is zero mean, the conditional covariance matrix is

$$\mathbf{H}_{t+1} = \mathbf{B}_{t+1}\mathbf{B}_{t+1}',$$

where $\mathbf{B}_{t+1}$ is a matrix square root factorization, which in the empirical application is obtained by spectral decomposition. Fixing the principal square root makes the mapping from the conditional covariance matrix and the marginal shape parameters to the higher-order co-moments uniquely defined.

The third-order co-moment tensor of $\mathbf{x}_{t+1}$ has elements

$$S_{ijk,t+1} = \mathbb{E}_t[x_{t+1}^i x_{t+1}^j x_{t+1}^k] = \sum_{r=1}^N B_{ir,t+1} B_{jr,t+1} B_{kr,t+1} \mu_{r,t+1}^{(3)},$$

while the fourth-order co-moment tensor has elements

$$\begin{aligned}
K_{ijkl,t+1} &= \mathbb{E}_t[x_{t+1}^i x_{t+1}^j x_{t+1}^k x_{t+1}^\ell] \\
&= \mathbf{H}_{ij,t+1} \mathbf{H}_{kl,t+1} + \mathbf{H}_{ik,t+1} \mathbf{H}_{jl,t+1} + \mathbf{H}_{il,t+1} \mathbf{H}_{jk,t+1} \\
&\quad + \sum_{r=1}^N B_{ir,t+1} B_{jr,t+1} B_{kr,t+1} B_{\ell r,t+1} \left(\mu_{r,t+1}^{(4)} - 3 \right).
\end{aligned} \tag{17}$$

Equation (17) separates the Gaussian covariance contribution from the excess-kurtosis contribution of the standardized Sk- t innovations. If all innovations were conditionally Gaussian, then $\mu_{r,t+1}^{(4)} = 3$ and the last term would vanish. For portfolio calculations, we use the matrix representations

$$\begin{aligned}
\mathbb{E}_t[\mathbf{x}_{t+1}] &= 0, \quad \mathbb{E}_t[\mathbf{x}_{t+1} \mathbf{x}_{t+1}'] = \mathbf{H}_{t+1}, \quad \mathbb{E}_t[\mathbf{x}_{t+1} \mathbf{x}_{t+1}' \otimes \mathbf{x}_{t+1}] = \mathbf{S}_{t+1}', \\
\mathbb{E}_t[\mathbf{x}_{t+1} \mathbf{x}_{t+1}' \otimes \mathbf{x}_{t+1} \otimes \mathbf{x}_{t+1}] &= \mathbf{K}_{t+1}',
\end{aligned}$$

that is, $\mathbf{S}_{t+1} \in \mathbb{R}^{N \times N^2}$ and $\mathbf{K}_{t+1} \in \mathbb{R}^{N \times N^3}$, are obtained by stacking the elements $S_{ijk,t+1}$ and $K_{ijkl,t+1}$ over the last indices.

D Parameter Uncertainty

This appendix describes how posterior parameter uncertainty is incorporated into the predictive moments used in the portfolio problem. The objective is twofold. First, we show how the predictive moments entering the fourth-order utility approximation are obtained by integrating over the posterior distribution of the model parameters (Harvey et al., 2010; Jondeau and Rockinger, 2012). Second, we clarify the role of estimation error in the conditional mean. In general, uncertainty about expected returns propagates into all noncentral moments of portfolio returns (DeMiguel et al., 2009). In our setting, however, the conditional mean of currency excess returns is pinned down by observed interest rate differentials and is therefore not a function of the estimated parameters (Ackermann et al., 2017). This provides a useful environment to isolate the contribution of parameter uncertainty in the conditional variance, co-skewness, and co-kurtosis. The notation follows the model in Section 3.1.

Let N denote the number of currencies and let Θ collect all parameters of the multivariate DCC-GARCH model with time-varying skewed Student- t innovations. At rebalancing date t , the information set is denoted by $\mathcal{F}_t$, and the Bayesian estimator delivers a particle approximation

to the posterior distribution,

$$p(\boldsymbol{\Theta}|\mathcal{F}_t) \approx \left\{ \boldsymbol{\Theta}_t^{(j)}, \bar{\mathcal{W}}_t^{(j)} \right\}_{j=1}^M, \quad \bar{\mathcal{W}}_t^{(j)} \geq 0, \quad \sum_{j=1}^M \bar{\mathcal{W}}_t^{(j)} = 1,$$

where M is the number of posterior particles, and $\bar{\mathcal{W}}_t^{(j)}$ are normalized importance weights. For a given parameter vector $\boldsymbol{\Theta}^{(j)}$, $j = 1, \dots, M$, the model specifies the conditional distribution of exchange-rate returns as

$$\mathbf{x}_{t+1} = \mathbf{H}_{t+1}(\boldsymbol{\Theta}^{(j)})^{1/2} \mathbf{z}_{t+1}, \quad \mathbf{z}_{t+1} \sim \text{MSk-}t\left(\boldsymbol{\eta}_{t+1}(\boldsymbol{\Theta}^{(j)}), \boldsymbol{\lambda}_{t+1}(\boldsymbol{\Theta}^{(j)})\right),$$

where $\mathbf{z}_{t+1}$ has zero mean and identity covariance. Currency excess returns are obtained through the deterministic carry transformation

$$\mathbf{y}_{t+1} = \mathbf{d}_t + \mathbf{A}_t \mathbf{x}_{t+1}, \quad \mathbf{d}_t = (r_t^1 - r_t^d, \dots, r_t^N - r_t^d)', \quad \mathbf{A}_t = \text{diag}(1 + r_t^1, \dots, 1 + r_t^N).$$

Both $\mathbf{d}_t$ and $\mathbf{A}_t$ are observed at time t . Under the short-horizon no-predictability restriction for exchange-rate changes, $\mathbb{E}_t[\mathbf{x}_{t+1}] = \mathbf{0}$, the conditional mean of currency excess returns is

$$\mathbb{E}_t[\mathbf{y}_{t+1}] = \mathbf{d}_t.$$

Thus, the conditional mean entering the portfolio problem is not a function of $\boldsymbol{\Theta}$. Parameter uncertainty affects the allocation only through the conditional second, third and fourth moment of $\mathbf{y}_{t+1}$.

D.1 Conditional Portfolio Moments for a Fixed Parameter Draw

Let $\mathbf{H}_{t+1}(\boldsymbol{\Theta}^{(j)})$, $\mathbf{S}_{t+1}(\boldsymbol{\Theta}^{(j)})$, and $\mathbf{K}_{t+1}(\boldsymbol{\Theta}^{(j)})$ denote the conditional second, third and fourth moments of $\mathbf{x}_{t+1}$ implied by the model for a fixed parameter value $\boldsymbol{\Theta}^{(j)}$. We use the matrix representation $\mathbf{S}_{t+1} \in \mathbb{R}^{N \times N^2}$ and $\mathbf{K}_{t+1} \in \mathbb{R}^{N \times N^3}$. Since $\mathbf{d}_t$ is deterministic at time t , the central co-moments of currency excess returns are

$$\begin{aligned} \mathbf{H}_{t+1}^y(\boldsymbol{\Theta}^{(j)}) &= \mathbf{A}_t \mathbf{H}_{t+1}(\boldsymbol{\Theta}^{(j)}) \mathbf{A}_t, \\ \mathbf{S}_{t+1}^y(\boldsymbol{\Theta}^{(j)}) &= \mathbf{A}_t \mathbf{S}_{t+1}(\boldsymbol{\Theta}^{(j)}) (\mathbf{A}_t \otimes \mathbf{A}_t), \\ \mathbf{K}_{t+1}^y(\boldsymbol{\Theta}^{(j)}) &= \mathbf{A}_t \mathbf{K}_{t+1}(\boldsymbol{\Theta}^{(j)}) (\mathbf{A}_t \otimes \mathbf{A}_t \otimes \mathbf{A}_t). \end{aligned}$$

For a candidate portfolio weights vector ($\mathbf{w}_t$), the portfolio excess return is $y_{t+1}^p = \mathbf{w}_t' \mathbf{y}_{t+1}$, and conditional on $\Theta^{(j)}$, the central portfolio moments are

$$\begin{aligned}\mu_{p,t+1} &= \mathbf{w}_t' \mathbf{d}_t, & s_{p,t+1}(\Theta^{(j)}) &= \mathbf{w}_t' \mathbf{S}_{t+1}^y(\Theta^{(j)}) (\mathbf{w}_t \otimes \mathbf{w}_t), \\ h_{p,t+1}(\Theta^{(j)}) &= \mathbf{w}_t' \mathbf{H}_{t+1}^y(\Theta^{(j)}) \mathbf{w}_t, & \kappa_{p,t+1}(\Theta^{(j)}) &= \mathbf{w}_t' \mathbf{K}_{t+1}^y(\Theta^{(j)}) (\mathbf{w}_t \otimes \mathbf{w}_t \otimes \mathbf{w}_t).\end{aligned}$$

Here $s_{p,t+1}(\Theta^{(j)})$ and $\kappa_{p,t+1}(\Theta^{(j)})$ are central third and fourth moments. The corresponding conditional noncentral moments are

$$\begin{aligned}m_{p,t+1}^{(1)}(\Theta^{(j)}) &= \mu_{p,t+1}, \\ m_{p,t+1}^{(2)}(\Theta^{(j)}) &= h_{p,t+1}(\Theta^{(j)}) + \mu_{p,t+1}^2, \\ m_{p,t+1}^{(3)}(\Theta^{(j)}) &= s_{p,t+1}(\Theta^{(j)}) + 3 h_{p,t+1}(\Theta^{(j)}) \mu_{p,t+1} + \mu_{p,t+1}^3, \\ m_{p,t+1}^{(4)}(\Theta^{(j)}) &= \kappa_{p,t+1}(\Theta^{(j)}) + 4 s_{p,t+1}(\Theta^{(j)}) \mu_{p,t+1} + \\ &\quad + 6 h_{p,t+1}(\Theta^{(j)}) \mu_{p,t+1}^2 + \mu_{p,t+1}^4.\end{aligned}$$

D.2 Averaging Predictive Noncentral Moments

The Bayesian portfolio problem maximizes expected utility where the expectation is taken jointly over future returns and parameter uncertainty

$$\max_{\mathbf{w}_t} \int \mathbb{E}_t \left[U^{[4]}(W_{t+1}) | \Theta \right] p(\Theta | \mathcal{F}_t) d\Theta.$$

Using the fourth-order Taylor approximation around wealth $W_{t+1} = 1$, we obtain

$$\int \mathbb{E}_t \left[U^{[4]}(W_{t+1}) | \Theta \right] p(\Theta | \mathcal{F}_t) d\Theta = \sum_{k=0}^4 \frac{1}{k!} \left. \frac{\partial^k U(W_{t+1})}{(\partial W_{t+1})^k} \right|_{W_{t+1}=1} \underbrace{\int m_{p,t+1}^{(k)}(\Theta) p(\Theta | \mathcal{F}_t) d\Theta}_{\bar{m}_{p,t+1}^{(k)}},$$

where $m_{p,t+1}^{(0)}(\Theta) = 1$. Hence, the relevant predictive objects are posterior averages of noncentral moments, not posterior averages of standardized co-moments. The predictive noncentral moments are approximated using the M posterior draws

$$\bar{m}_{p,t+1}^{(k)} \approx \frac{1}{M} \sum_{j=1}^M m_{p,t+1}^{(k)}(\Theta^{(j)}), \quad k = 1, \dots, 4.$$

It is worth noting that, at the last step of the SMC, we perform a systematic resampling, and therefore we can use the simple average over the posterior draws. The same reasoning applies to the mean-variance strategy, which truncates the Taylor expansion of the utility function at the second order.

E Statistical Tests

This appendix details the inference procedures underlying the significance stars reported in Table 4. All tests are one-sided, with the null hypothesis that the mean-variance-skewness-kurtosis (MVSK) strategy does not improve upon the mean-variance (MV) benchmark and the alternative that it does improve. They are computed on the paired out-of-sample realizations of the two strategies over the common evaluation period. Because the relevant object is always a difference between two performance measures computed on the same dates, every test is implemented on paired data (utility flows or paired returns for the two strategies), so as to preserve the dependence between the two performance estimates rather than treating them as independent samples.

Performance fee (certainty-equivalent) test. Let $u_t = U(1+y_t^p)$ denote the realized period- t CRRA utility flow of a strategy, as in (5), and let $\Delta u_t = u_t^{MVSK} - u_t^{MV}$. The null hypothesis is $\mathbb{E}[\Delta u_t] \leq 0$ against $\mathbb{E}[\Delta u_t] > 0$. We regress Δu_t on a constant and compute a HAC standard error for the intercept using a Bartlett (Newey-West) kernel with automatic bandwidth selection via the AR(1) plug-in rule (Andrews, 1991) and a small-sample degrees-of-freedom correction. The test statistic is the HAC t -statistic for the intercept, and the one-sided p -value is obtained from the standard normal distribution. The procedure is analogous to the Diebold-Mariano test, with the realized utility differential serving as the loss differential (Diebold and Mariano, 1995). It is the basis for the performance fee (Φ) column of Table 4.

Sharpe ratio test. We test $H_0 : SR^{MVSK} \leq SR^{MV}$ using the framework of Ledoit and Wolf (2008). The test statistic studentizes the difference in sample Sharpe ratios by a prewhitened HAC standard error computed with the quadratic-spectral kernel and automatic bandwidth selection, delivering an asymptotic p -value from the standard normal distribution. As a robustness check on the asymptotic approximation, we also compute a paired circular block bootstrap p -value (Politis and Romano, 1992): for 4,999 bootstrap replications and block lengths $\{5, 10, 15\}$, we resample blocks of paired dates (common to both strategies), recompute the studentized statistic on each resampled pair, and obtain the bootstrap p -value as the fraction of bootstrap statistics at least as large as the one computed on the original data.

Modified Sharpe ratio test. The modified (Cornish–Fisher, VaR-based) Sharpe ratio test of [Ardia and Boudt \(2015\)](#) is applied at the 5% loss threshold. We use the `msharpeTesting` function of the `PeerPerformance` R package, in both its asymptotic HAC form and its studentized circular block bootstrap form, with 4,999 bootstrap replications and block lengths $\{5, 10, 15\}$. The package returns a two-sided p -value, which we convert to the one-sided p -value relevant for $H_0 : mSR^{MVSK} \leq mSR^{MV}$ using the sign of the test statistic.

Adjusted Sharpe, Omega, and Sortino ratio tests. No closed-form asymptotic distribution is available for the difference in adjusted Sharpe ratios, Omega ratios, or Sortino ratios, so inference relies exclusively on a paired circular block bootstrap applied directly to the two return series ([Politis and Romano, 1992](#)). For 4,999 bootstrap replications and block lengths $\{5, 10, 15\}$, we draw circular blocks of common dates, apply them to both the MVSK and MV return series so as to preserve their joint dependence structure, and recompute each statistic’s difference, $\hat{\Delta}^*$, on the resampled pair. The bootstrap distribution is recentered to impose the null, $\hat{\Delta}_c^* = \hat{\Delta}^* - \hat{\Delta}$, and the (continuity-corrected) one-sided p -value is $\left(\#\{\hat{\Delta}_c^* \geq \hat{\Delta}\} + 1\right) / (4,999 + 1)$, where $\hat{\Delta}$ is the statistic computed on the original paired data.

Reporting convention. All bootstrap-based tests are computed for block lengths $b = 5, 10$, and 15 to check robustness to the choice of block length. Table 4 reports stars based on the asymptotic HAC p -value where one is available (performance fee, Sharpe ratio, and modified Sharpe ratio) and on the block-bootstrap p -value with $b = 10$ otherwise (adjusted Sharpe, Omega, and Sortino ratios). Given the relatively long out-of-sample period of 1,029 weekly observations, the resulting inference is stable across the alternative block lengths considered: the significance pattern is unchanged for $b = 5$ and $b = 15$. For brevity, these additional bootstrap results are not reported. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively, throughout.

F Additional Empirical Results

F.1 Robustness to Higher Transaction Costs

This appendix repeats the net-of-cost comparison of Section 4.4 after doubling both transaction cost components relative to the baseline. The proportional spot cost rises from 1 to 2 basis point and the proportional swap cost from 0.1 to 0.2 basis point, i.e. $c_i^{\text{spot}} = 0.0002$ and $c_i^{\text{swap}} = 0.00002$ for all currencies. This is a conservative check in the sense of Section 4.4. Although it does not reach the 5–16 basis point range of full quoted spreads reported by [Darvas \(2009\)](#) for the pre-2010

period, it doubles the baseline cost used throughout the paper while leaving the portfolio rules themselves unchanged.

Table 10: Return distribution and performance diagnostics net of transaction costs (double transaction costs): MVSK and Currency Benchmarks

γ	Strategy	μ_p	σ_p	SK_p	KU_p	SR_p	aSR _p	Ω_p -ratio	S_p -ratio
$\mathcal{W}^U$ (Unconstrained)									
	L4S4*	5.075	30.851	-0.702	10.602	0.164	0.160	1.070	0.223
5	MVSK	6.228	25.435	0.980	11.242	0.245	0.250	1.107	0.383
15	MVSK	2.063	8.485	0.977	11.636	0.243	0.248	1.106	0.380
$\mathcal{W}^G$ (Unit gross exposure)									
	L4S4	0.508	3.085	-0.702	10.602	0.164	0.160	1.070	0.223
5	MVSK	1.260	5.976	0.185	9.120	0.211	0.210	1.086	0.302
15	MVSK	0.276	3.446	-0.161	6.865	0.080	0.080	1.031	0.112
$\mathcal{W}^{DN}$ (Dollar-neutral)									
	L4S4	0.508	3.085	-0.702	10.602	0.164	0.160	1.070	0.223
5	MVSK	1.164	4.156	-0.122	6.813	0.280	0.275	1.111	0.398
15	MVSK	0.600	3.013	-0.365	5.300	0.199	0.196	1.077	0.278
$\mathcal{W}^L$ (Long-only)									
	EW	-0.416	7.371	-0.059	5.624	-0.056	-0.056	0.979	-0.079
5	MVSK	0.256	7.248	-0.182	5.172	0.035	0.035	1.014	0.049
15	MVSK	-0.083	6.331	-0.089	5.332	-0.013	-0.013	0.995	-0.018

Notes: This table reproduces Table 8 using transaction costs twice as large as in the baseline specification in Section 4.4: 2 basis point for spot transactions and 0.2 basis point for swap transactions. EW denotes the equally weighted currency portfolio. L4S4 denotes the long-four/short-four carry portfolio, which takes equal-weighted long positions in the four highest-yielding currencies and short positions in the four lowest-yielding currencies. L4S4* denotes L4S4 scaled to a gross exposure of 10. Benchmark statistics do not depend on risk aversion and are therefore reported once per mandate. MVSK denotes the mean-variance-skewness-kurtosis portfolio. Portfolio Mean (μ_p) and Volatility (σ_p) are annualized and expressed in percentage points. Skewness (SK_p) and Kurtosis (KU_p) are reported at portfolio return frequency to avoid vanishing effects. Sharpe ratio (SR_p), adjusted Sharpe ratio (aSR_p), Omega ratio (Ω_p -ratio) and Sortino ratio (S_p -ratio) are annualized and reported in ratio units.

Doubling transaction costs leaves the qualitative ranking of Table 8 unchanged. Mean-Variance-Skewness-Kurtosis (MVSK) still strictly dominates its benchmark for the unconstrained, dollar-neutral, and long-only investors at both risk-aversion levels. For $\mathcal{W}^U$ the net Sharpe ratio of MVSK, 0.245 and 0.243, remains comfortably above the doubled-cost long-four/short-four leverage 10 carry portfolio (L4S4*) benchmark, 0.164; for $\mathcal{W}^{DN}$, MVSK's net Sharpe ratio of 0.280 and 0.199 remains above the L4S4 benchmark of 0.164 at both risk-aversion levels; and for $\mathcal{W}^L$, MVSK's net Sharpe ratio is positive at $\gamma = 5$ while EW's is negative, exactly as in the baseline-cost case, although at $\gamma = 15$ the higher cost is now enough to turn the net Sharpe ratio of MVSK marginally negative as well. As in the baseline case, the exception is $\mathcal{W}^G$. MVSK still edges out L4S4 at $\gamma = 5$ (0.211 versus 0.164) but falls behind at $\gamma = 15$ (0.080 versus 0.164), the same qualitative pattern documented in Section 4.4, only somewhat more pronounced because of the higher turnover-driven drag. Overall, doubling both cost components does not overturn any

of the rankings established in the main text; it only compresses the margins by which MVSK outperforms, consistent with the transaction-cost mechanism described in Section 4.4.

F.2 Alternative Gross-Exposure Bounds, Net of Transaction Costs

This appendix examines whether the currency-benchmark comparisons in Sections 4.4 and 4.3 are sensitive to the amount of gross exposure the investor is permitted to take. We consider three additional upper bounds on gross exposure,

$$\sum_{i=1}^N |w_i| \leq \kappa, \quad \kappa \in \{1, 5, 20\},$$

with net exposure left unconstrained. For each bound, we compare the MVSK strategy with the long-four/short-four carry benchmark (L4S4), using the baseline transaction-cost specification of Section 4.4 (1 basis point for spot transactions and 0.1 basis point for swap transactions).

Table 11: Return distribution and performance diagnostics net of transaction costs: MVSK under alternative gross-exposure limits and Currency Benchmark

γ	Strategy	μ_p	σ_p	SK_p	KU_p	SR_p	aSR _p	Ω_p -ratio	$\mathcal{S}_p$ -ratio
	L4S4	0.908	3.086	-0.709	10.625	0.294	0.276	1.129	0.403
<i>Gross exposure ≤ 1, net exposure unconstrained</i>									
5	MVSK	1.946	5.977	0.181	9.111	0.326	0.320	1.136	0.472
15	MVSK	0.939	3.447	-0.154	6.871	0.272	0.267	1.110	0.388
<i>Gross exposure ≤ 5, net exposure unconstrained</i>									
5	MVSK	3.636	13.417	-0.395	6.202	0.271	0.264	1.108	0.381
15	MVSK	3.132	7.207	0.703	12.972	0.435	0.423	1.187	0.665
<i>Gross exposure ≤ 20, net exposure unconstrained</i>									
5	MVSK	11.178	23.069	0.755	11.429	0.485	0.474	1.213	0.757
15	MVSK	4.989	8.505	1.036	11.639	0.586	0.573	1.277	0.962

Notes: Results are net of the baseline transaction costs set in Section 4.4 (1 basis point for spot transactions and 0.1 basis point for swap transactions) under three alternative upper bounds on gross exposure, $\sum_{i=1}^N |w_i| \leq \kappa$, with $\kappa \in \{1, 5, 20\}$ and net exposure left unconstrained. L4S4 denotes the long-four/short-four carry portfolio, which takes equal-weighted long positions in the four highest-yielding currencies and short positions in the four lowest-yielding currencies. MVSK denotes the mean-variance-skewness-kurtosis portfolio. Portfolio Mean (μ_p) and Volatility (σ_p) are annualized and expressed in percentage points. Skewness (SK_p) and Kurtosis (KU_p) are reported at portfolio return frequency to avoid vanishing effects. Sharpe ratio (SR_p), adjusted Sharpe ratio (aSR_p), Omega ratio (Ω_p -ratio) and Sortino ratio ($\mathcal{S}_p$ -ratio) are annualized and reported in ratio units.

The robustness results show that the net performance of MVSK relative to the L4S4 benchmark depends on the amount of gross exposure the investor is permitted to take. With a unit gross-exposure bound, MVSK achieves a higher Sharpe ratio than L4S4 for $\gamma = 5$ (0.326 versus 0.294), but a slightly lower Sharpe ratio for $\gamma = 15$ (0.272 versus 0.294). At $\kappa = 5$, the ranking reverses across risk-aversion levels: the MVSK Sharpe ratio is 0.271 for $\gamma = 5$, slightly below

L4S4, but rises to 0.435 for $\gamma = 15$, well above the benchmark. With a looser bound of $\kappa = 20$, MVSK outperforms L4S4 at both risk-aversion levels, with Sharpe ratios of 0.485 and 0.586, respectively, compared with 0.294 for L4S4.

The results therefore indicate that the benchmark comparison is not specific to a unit-exposure implementation. Allowing greater gross exposure can materially increase the realized risk-adjusted performance of MVSK, particularly for the more risk-averse investor. At the same time, the relationship is not monotonic across all risk-aversion levels: the $\gamma = 5$, $\kappa = 5$ case provides a small reversal of the ranking. Overall, the results support the robustness of the main benchmark comparison across a range of gross-exposure bounds, while also showing that the economic magnitude of the comparison depends on the amount of exposure the investor is permitted to take.

RECENT PUBLICATIONS BY *CEIS Tor Vergata*

Exchange-Rate Regimes and the Behaviour of Exporters

Cosimo Petracchi, Luca Riva and Marco S. Petterson

CEIS Research Paper, 627 August 2026

Pricing options on illiquid assets using liquid market benchmarks: an application to energy markets

Federico Aluigi, Lucia Caramellino, Paolo Pigato and Edoardo Scrima

CEIS Research Paper, 626 August 2026

Image Considerations in Second-Hand Adoption: A Vignette Experiment

Luca Congiu, Luisa Lorè and Mariangela Zoli

CEIS Research Paper, 625 July 2026

Communication Failure: The Hidden Face of The Tragedy of the Commons

Philippe Colo and Guillaume Pommey

CEIS Research Paper, 624 July 2026

Female Mayors and Public Procurement

Erminia Florio and Giancarlo Spagnolo

CEIS Research Paper, 623 July 2026

Optimal Structure of Penalties with Judgment-Proof Injurers

Guillaume Pommey

CEIS Research Paper, 622 June 2026

Flexibility Versus Security in Agency Contracts with Moral Hazard

Lorenzo Bozzoli and Guillaume Pommey

CEIS Research Paper, 621 June 2026

The Misallocation Costs of Inflation: A Sufficient Statistics Approach

Klaus Adam, Andrey Alexandrov and Henning Weber

CEIS Research Paper, 620 April 2026

Sensitivity of the Euro OIS Term Structure to ECB Policy Rate Surprises

Stefano Herzel and Marco Nicolosi

CEIS Research Paper, 619 January 2026

Pink queue and Double Standard: what Markov Chains uncover in academic career progressions

Marianna Brunetti and Annalisa Fabretti

CEIS Research Paper, 618 December 2025

DISTRIBUTION

Our publications are available online at www.ceistorvergata.it

DISCLAIMER

The opinions expressed in these publications are the authors' alone and therefore do not necessarily reflect the opinions of the supporters, staff, or boards of CEIS Tor Vergata.

COPYRIGHT

Copyright © 2026 by authors. All rights reserved. No part of this publication may be reproduced in any manner whatsoever without written permission except in the case of brief passages quoted in critical articles and reviews.

MEDIA INQUIRIES AND INFORMATION

For media inquiries, please contact Barbara Piazzzi at +39 06 72595584 or by e-mail at barbara.piazzzi@ceis.uniroma2.it. Our web site, www.ceistorvergata.it, contains more information about Center's events, publications, and staff.

DEVELOPMENT AND SUPPORT

For information about contributing to CEIS Tor Vergata, please contact at +39 06 72595587 or by e-mail at segr.ceis@economia.uniroma2.it