Reassessing Sources of Risk Premiums in Currency Markets *

Mikhail Chernov,† Magnus Dahlquist,‡ and Lars A. Lochstoer§

First draft: November, 2023

This draft: March 12, 2025

Abstract

We show that a small set of emerging markets with floating exchange rates expand the investment frontier substantially relative to G10 currencies. The frontier is characterized by an out-of-sample mean-variance efficient portfolio that prices G10- and emerging markets-based trading strategies unconditionally as well as conditionally. Our approach reveals that returns to prominent trading strategies are largely driven by factors that do not command a risk premium. After real-time hedging of such unpriced risks, the Sharpe ratios of these strategies increase substantially, providing new benchmarks for currency pricing models. For instance, the Sharpe ratio of the carry strategy increases from 0.71 to 1.29. The unpriced risks are related to geographically-based currency factors, while the priced risk that drives currency risk premiums is related to aggregate consumption exposure.

JEL classification codes: F31, G12, G15.
Keywords: Currency risk premiums, factor models, mean-variance efficiency.

*We thank Viral Acharya, Federico Baldi-Lanfranchi, Max Croce, Pasquale Della Corte, Lira Mota, as well as participants in seminars and conferences at the Bank for International Settlements, the 2024 BI-SHOF conference, the 2024 CEPR Paris Symposium, CIREQ-CMP Econometrics Conference in Honor of Eric Ghysels, Copenhagen Business School, the 2024 ESSFM meeting in Gerzensee, London Business School, London School of Economics, NYU Stern, University of Notre Dame, University of Lancaster, University of Lausanne / EPFL, the 2024 UNSW Asset Pricing Workshop, University of Texas Austin, and Warwick Business School. We thank Haohang Wu for excellent research assistance. Dahlquist gratefully acknowledges support from the Jan Wallander and Tom Hedelius Foundation. The paper was previously circulated as "An anatomy of currency strategies: The role of emerging markets."

†UCLA, NBER, and CEPR; mikhail.chernov@anderson.ucla.edu.

‡Stockholm School of Economics, and CEPR; magnus.dahlquist@hhs.se.

§UCLA and NBER; lars.lochstoer@anderson.ucla.edu.

1 Introduction

In this paper we revisit the risk-return trade-off implicit in famous currency trading strategies. We are motivated by the fact that the impressive performance of G10-based strategies started to level off in the mid 2000s. At the same time, an expanded universe of currencies appears to maintain high returns to the same style strategies.

Indeed, Figure 1 shows these effects by displaying the cumulative returns for two popular strategies: dollar, i.e., an equal-weighted portfolio of exchange rates vs USD, and cross-sectional carry. Returns on both strategies are computed for G10 only (labeled G) and all currencies with available forward contracts (labeled GEX). We see that, as data on emerging countries becomes available in 1997, carry based on all currencies started decoupling from the G-based carry. Furthermore, G-based carry performance has flattened out during the first decade of the 21st century further contributing to the decoupling. Dollar is exhibiting less dramatic but similar patterns.

These observations lead us to two conjectures. The first one is that incorporating emerging markets is critical to understanding the currency market risk-return trade-off, particularly its conditional dynamics.[1] The second one is based on the fact that risk premiums for certain trading strategies seem to approach zero in the later sample, even though the volatility of the returns from these strategies remains consistent throughout the sample period. Thus, we surmise that the strategies are exposed

[1]We refer to non-G10 currencies as emerging because such currencies are a majority. Formally, such labeling is not correct – some of the non-G10 countries with floating exchange rate would be classified as developed. Even the tightly managed currencies set contain DKK, which is developed. Also, we refer to the remaining currencies as extended emerging-market ones, while some of the countries corresponding to these currencies would be labeled as frontier. The developed-emerging-frontier classification leans heavily on the properties of the stock markets, which we do not consider at all. Therefore, we do not adhere to these labels strictly.

to important drivers of currency return comovements that do not command a risk premium.

Understanding this evidence is important in light of the literature, which argues that currency risk premiums are driven by exposures to the dollar and carry factors, which in turn are strongly related to the two first principal components of currency return volatility (see, e.g., Verdelhan, 2018). We call drivers of currency comovements that do not command a risk premium for "unpriced risk" and seek to conditionally disentangle priced from unpriced risk throughout the sample.

Our first contribution is to show that the continued strong performance of global currency strategies arises from the subset of emerging market economies that are under a floating exchange rate regime, or, more accurately, the least pegged currencies. That is, we show that strategies formed using the full set of currencies do not have significant alpha relative to strategies formed on G10 and emerging floating-regime currencies. We label this set GE (versus G for only G10 and GEX for G10 plus the EXtended set of all emerging-market currencies).

Our second contribution is to build on Chernov, Dahlquist, and Lochstoer (2023) and construct a real-time estimate of the unconditional mean-variance efficient (UMVE) portfolio for G10 and emerging market currencies. Excess returns on this portfolio represent a real-time, tradeable, single factor. This is "the king of all factors" because, in theory, it prices all dynamic admissible trading strategies unconditionally. As a result, the factor correctly prices the full cross-section of currencies and trading strategies associated with them conditionally as well (Hansen and Richard, 1987). Thus, it subsumes all other cross-sectional pricing factors, which have been considered in the literature. Additionally, the stochastic discount factor that condi-

tionally prices excess returns on all currencies and their strategies can be represented as an affine function, with constant coefficients, of the UMVE factor. We show empirically that the estimated UMVE factor does indeed price all currency strategies conditionally and unconditionally, in sample and out of sample (OOS).

In our third contribution, we establish, using the estimated UMVE, that a large fraction of risk associated with currency strategies, some of which appear as pricing factors in prominent currency models, is unpriced. We relate these unpriced risks with long-short strategies capturing geographical regions. The fraction of variation that is priced exhibit strong time variation in the associated risk premiums, which flatten out in the later part of the sample for G10 countries, but not for others; just like realized returns in Figure 1. We demonstrate that the priced risk are significantly correlated with long-run U.S. consumption growth, while the unpriced risks are not.

We find that the risk-return trade-off indeed is substantially affected by the floating-regime emerging market currencies. The conditional maximal Sharpe ratio (MSR) when using only G currencies trends down over the sample, flattening out at a low level during the last 20 years. In contrast, the GE set delivers continued high conditional MSR the during the last 20 years, about twice as high as those for the G set over this period. The annualized sample UMVE Sharpe ratio (SR) for the G set is 1.02 versus 1.34 for the GE set. We emphasize that these UMVE portfolios are constructed in a pure out-of-sample fashion. The incremental SR, captured by the information ratio, for the GE set over the G set is 0.86. The UMVE based on the GE set prices trading strategies constructed using the broader GEX set, as well as trading strategies formed on the GE and G sets. Thus, we focus on the UMVE from the GE set as the single factor that captures priced risk in the currency market.

We implement standard "alpha" regressions of popular trading strategies on the UMVE portfolio. Strikingly, the R^2s in these regressions are low, with the highest being around 30% for the cross-sectional carry strategy. That is, more than two thirds of the return variation in the carry portfolio is due to unpriced risk. For other strategies this fraction is larger. Given our estimates of the conditional covariance matrix of returns, the conditional portfolio weights of each strategy and the UMVE portfolio, we can hedge out this unpriced risk in real time. This leads to strong increases in the strategy SR as the average return is left approximately unaltered but the return variance is substantially reduced. For example, the cross-sectional carry strategy goes from a SR of 0.71 to 1.29 when unpriced risks are hedged out, the dollar factor goes from 0.33 to 0.91, the 12-month cross-sectional momentum strategy from 0.24 to 0.99, and the cross-sectional value goes from 0.65 to 1.29. These are large increases, which suggest that the standard trading strategies used in the literature are not suitable to use as factors for risk pricing due to their contamination from these unpriced risks.

Indeed, we show that the classic dollar-carry model, both unconditional and conditional, cannot explain the cross-section of currency returns. The tricky part about testing this model is that many strategies are spanned by the factors, especially in the conditional setting. Thus, the model mechanically prices prominent, high SR strategies, such as the Dollar Carry strategy, diminishing the power of tests. We exploit the importance of unpriced risks and implement additional tests of the model on the basis of the UMVE and the unpriced components of the strategy returns. The latter are simply the traditional strategies with the priced risks hedged out in real-time using the estimated UMVE. With this approach there is no longer a mechanical connection between strategy returns and candidate factors. Both unconditional and

conditional models are rejected. This happens because the pricing factors in the Dollar Carry model contain large unpriced risks that distort the true risk-return relationship. Because dollar and carry are close to the first two principal components of the variation in currency returns, these results indicate that factors constructed from the main sources of currency co-variation cannot capture the full risk-return trade-off.

Our results lead to the question of what these priced and unpriced risks are economically. We relate the unpriced risks to geographical factors driving currency comovement. In particular, we show that hedging using a Europe factor and a Rest of the World factor, where the returns are the equal-weighted returns of the currencies in that region, goes a long way towards explaining the unpriced risks. Such comovements can be driven by common shocks to the economies of close countries (e.g., Lustig and Richmond, 2020, Richmond, 2019). We show that the priced risk, on the other hand, is significantly related to long-run U.S. consumption growth (while unpriced risk is not related to consumption), indicating that macroeconomic risks indeed are important for understanding currency risk premiums.

Finally, it is natural to worry about the impact of trading costs when considering less liquid currencies in the analysis. In some cases the potential costs, as measured by bid-offer spreads, are so large that one may be concerned that some trading strategies are hindered. There is a large literature on currency transaction costs, which is primarily motivated by the concern that bid-offer spreads based on indicative quotes in standard databases could be too conservative. As a result, many papers, which we review later, consider various estimates of effective proportional trading costs and price impact.

Given that each individual study considers, due to relevant data availability, a limited time frame, a limited set of strategies, or a limited set of currencies, we consider a range of possible transaction costs as a fraction of quoted bid-offer spreads (0%, 25%, 50%, and 100%). While we do not consider price impact, we think that one of the points in this range should be sufficiently close to the combination of effective proportional costs and price impact. We find that transaction costs indeed can be important for our analysis, but that our main conclusions are robust to such frictions.

In particular, we revisit the UMVE analysis accounting for transaction costs. When there are no transaction costs, the UMVE returns and the Gibbons, Ross, and Shanken (1989) (GRS) test statistic can be constructed analytically. In the presence of transaction costs the UMVE portfolio has to be constructed numerically. We adopt the approach of Detzel, Novy-Marx, and Velikov (2023), which was developed for equities. Because currencies from the GE set are capable of spanning GEX-based strategies, which typically have higher transaction costs, when transaction costs are ignored, the costs-based analysis is a robustness check. We find that accounting for transaction costs in the UMVE construction and testing methodology strengthens our conclusion that the GE-based UMVE is sufficient to account for the currency market risk-return trade-off.

Literature. Because the cross-section of equities is orders of magnitudes larger than that of currencies, it is impossible or very difficult to estimate the UMVE portfolio. The equities literature has grappled with this difficulty since at least Roll (1977). The shift towards machine learning methods is primarily driven by this inherent difficulty (see Kelly and Xiu, 2023, for a review). Recently, Daniel, Mota, Rottke, and Santos (2020) and Kozak and Nagel (2023) develop conditions under which a factor model might span the UMVE. Daniel, Mota, Rottke, and Santos (2020) argue

that characteristics are likely to be correlated with unpriced factor risk. As a result, the set of characteristics-based portfolios will not span the UMVE portfolio. Because the UMVE portfolio is unattainable in their case, they cannot construct an optimal hedge and resort to approximations. We can construct the UMVE in the case of exchange rates, and, therefore, the optimal real-time hedges in our setting, and we confirm that hedging out unpriced risks from prominent currency trading strategies leads to substantial performance improvement.

We use the UMVE-construction approach of Chernov, Dahlquist, and Lochstoer (2023), but differ in terms of scope and findings. We construct the UMVE for a wide set of currencies and account for transaction costs. We demonstrate the importance of floating non-G10 countries for capturing the risk-return trade-off in the currency space. We show that the priced components of prominent currency trading strategies unequivocally reject the traditional characteristic-based currency factor models. We further document that the expected returns of the priced components can be explained by a long-run consumption risk factor and that unpriced components relate to geographic factors.

Bansal and Dahlquist (2000) is an early study that considers cross-sectional currency pricing with factor models using both developed and emerging currencies. The subsequent literature considers both emerging currencies and transaction costs, but more in the spirit of robustness checks relative to the main G10 data (e.g., Lustig, Roussanov, and Verdelhan, 2011, Menkhoff, Sarno, Schmeling, and Schrimpf, 2012a). These studies document that the SRs of their carry strategies are higher when both developed and emerging currencies are included (rather than just developed currencies) in high-minus-low portfolios. The differences are smaller when adjusting for transaction costs. The literature has developed both in the direction of more explicit

consideration of the role of emerging markets for currency strategy returns, and the level and impact of transaction costs.

Andrews, Colacito, Croce, and Gavazzoni (2024) document a decline in carry returns for G10 currencies during the post-2008 period and relate this to the compression of short-term interest rate differentials in this period. Nucera, Sarno, and Zinna (2024) demonstrate that adding emerging currencies significantly improves the performance of trading strategies. They consider a comprehensive dataset of currencies. However, they do not consider transaction costs, do not distinguish between different types of emerging currencies, and do not consider conditional pricing. Menkhoff, Sarno, Schmeling, and Schrimpf (2012b) reach similar conclusions for momentum when accounting for transaction costs. Török (2023) considers a broad selection of currencies and the impact of transaction costs, although using an approach different from ours. His main conclusion is that frontier currencies (similar to our GEX minus GE) are instrumental for the impressive performance of carry. We reach the opposite conclusion by showing that GE currencies span GEX strategies, in general, and carry, in particular.

Lyons (2001) raises the concern that bid-offer spreads based on indicative quotes may be overstating the impact of transaction costs. Cespa, Gargano, Riddiough, and Sarno (2022) and Gilmore and Hayashi (2011) estimate effective trading costs. The former paper concludes that these costs are closer to 25% of the indicative bid-offer rates. The literature considers various fractions of the spreads when computing strategy returns net of transaction costs (see, e.g., Kroencke, Schindler, and Schrimpf, 2014, Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b, Menkhoff, Sarno, Schmeling, and Schrimpf, 2017).

Subsequent studies propose to optimize currency trading strategies with respect to transaction costs. Korsaye, Trojani, and Vedolin (2023) solve explicit entropy minimization problem in the presence of general frictions. Their empirical work considers proportional transaction costs such as the ones that we consider. Because of the entropy-based SDF, the pricing cannot be tested conditionally. In their empirical work, they consider subsets of G10 and floating emerging currencies. Orlowski, Sokolovski, and Sverdrup (2021) pursue a similar objective. They address the difficulty of conditional testing of entropy-based SDFs by rolling estimates every 15 years. They use returns on currencies similar to GE net of transaction costs.

Barroso and Santa-Clara (2015) and Filippou, Maurer, Pezzo, and Taylor (2024) consider optimal currency portfolios accounting for transaction costs using data similar to our GE set. Both papers find that optimized portfolios outperform strategies based on a single signal. Neither tests for cross-sectional pricing. Our finding is that both versions of the UMVE portfolio, cost-optimized or not, constructed from a smaller set of countries can explain dynamic trading strategies from a much larger set of currencies.

2 Data

Our objective is to construct the most comprehensive dataset of exchange rates that are traded in both spot and forward markets. Today, the world has around 160 currencies. We consider the WMR FX benchmarks retrieved from Refinitiv Datastream. These spot and forward exchange rates are consistently calculated and used by equity and bond index compilers. The requirement of traded forward prices

narrows down the list to 75 currencies in the WMR database for the period December 31, 1996 to June 30, 2023. Appendix A.1 provides more details about the database and lists the selected currencies with their currency codes. We complement these exchange rates versus the USD with spot exchange rates versus the GBP that go back to January 31, 1990 (to have a longer sample of spot exchange rates for the forecasting of depreciation rates). They are also from the WMR database. We backfill this dataset to 1985 using the data considered in Chernov, Dahlquist, and Lochstoer (2023).

We drop the currencies that were members of the EUR at the launch on January 1, 1999 (ATS, BEF, FIM, FRF, IEP, ITL, NLG, PTE, and ESP). We use the DEM before 1999 and the EUR from 1999 and onwards. We include currencies that joined the EUR after its launch (HRK, CYP, EEK, GRD, LVL, LTL, MTL, SKK, and SIT) up to the date they are fixed. We drop six currencies entirely due to extreme inflation and financial data issues (ARS, EGP, JOD, RUB, TRY, and UAH). We begin with forward exchange rates for IDR, KES, MYR, and PEN in June 2007, January 2012, July 2005, and April 2004, respectively, as they have had capital controls and/or questionable data quality before. In total we have 59 spot and forward exchange rates versus the USD. Appendix A.2 discusses non-deliverable forward contracts.

Having selected the currencies for our analysis, we proceed with classifying them into different sets. Previous literature has focused primarily on the classifiers related to economic development such as G10, developed (closely related but not identical to G10), emerging, or frontier. We shift the focus to liquidity of the spot and forward exchange. Further, exchange arrangements can potentially impact the performance of trading strategies. For example, currencies that are literally pegged to USD will have economically different excess returns (simply equal to the interest rate differential)

as compared to the free-floating ones. In practice, these concepts are closely related to the previously considered classifiers, but, we think, produce a somewhat more nuanced categorizations. We want to stress that currencies and their respective countries often have complicated financial and macroeconomic developments so any classification should not be viewed as immutable. Different teams of researchers might settle on slightly different groupings.

Based on these considerations, our first set coincides with the G10 categorization. Countries in this group enjoy well-developed financial markets and stable macroeconomic environment. The currency markets of these countries are the most liquid in the world. We use the moniker "G" to refer to this set of currencies in the paper.

Next, we separate out a set of currencies, which could be described as floating, or least pegged, to which we refer as floating emerging currencies. We use currency classification introduced in Ilzetzki, Reinhart, and Rogoff (2019) (IRR) to do so. We select currencies with predominant IRR scores of 11, 12, or 13 during the period of available forward rate data, indicating that they have had moving bands (allowing for both appreciation and depreciation over time), managed floats, or free floats. Usually, currencies with these scores remain above 10.[2] Thus, we do not change our classification throughout the sample. We refer to the combination of this set with G as "GE" (G10 and Emerging).

The remaining set of currencies is referred to extended emerging currencies. These are currencies with typical IRR scores 2–8, indicating that they have had currency board arrangements, pegs, crawling pegs, bands, and crawling bands. This set also

[2]INR was below 10 during the first half of the sample, so it is at the border of our classification. We place it in the floating set because the currency is lightly managed in the recent half of the sample (score 11, a moving band).

includes TWD, which does not have an IRR score. We refer to the combination of this set with GE as "GEX" when combined with (GE and eXtended Emerging).

In summary, we consider three currency sets:

G: G10 (AUD, CAD, EUR spliced with DEM, JPY, NZD, NOK, SEK, CHF, and GBP).

GE: G10 combined with 12 floating emerging currencies (BRL, CLP, COP, ISK, INR, ILS, MXN, PLN, SGD, ZAR, KRW, and THB).

GEX: G10 combined with 12 floating emerging currencies and 38 extended emerging currencies (BHD, BGN, CNY, HRK, CYP, CZK, DKK, EEK, GHS, GRD, HKD, HUF, IDR, KZT, KES, KWD, LVL, LTL, MYR, MTL, MAD, OMR, PKR, PEN, PHP, QAR, RON, SAR, RSD, SKK, SIT, LKR, TWD, TND, UGX, AED, VND, and ZMW).

We complement the currency data with data on CPI, which are downloaded from the statistical databases of OECD and IMF. We use monthly data over the period January 1976 to June 2023, but for Australia and New Zealand we use quarterly data as monthly data are not available. In the case of quarterly data, the value observed at the end of a quarter is repeated monthly in the next quarter to avoid a look-ahead bias. Taiwan data are from their National Statistics.

Let the USD be the measurement (numeraire) currency. We express all exchange rates in USD per unit of foreign currency. Let S_t^i and F_t^i denote the spot exchange rate and the one-month forward exchange rate of country i, respectively.

The payoff of a forward contract (when buying one unit of the foreign currency) is $S^i_{t+1} - F^i_t$. One common way to scale this payoff to define excess return is to divide by F^i_t :

$$R^{ei}_{t+1} = (S^i_{t+1} - F^i_t)/F^i_t. \tag{1}$$

This definition implies that the amount of foreign currency bought is one "forward" USD. Thus, this is an excess return to a trading strategy regardless of whether covered interest rate parity (CIP) holds or not. Appendix A.3 describes the trading strategies constructed from these returns. The trading strategies are taken from earlier research and, in short, they include the currency "market" (dollar) factor, as well as various carry and momentum factors, and a value factor.[3]

We also consider excess returns based on quoted bid and offer exchange rates. The excess return on the long position, net of transaction costs, is

$$R^{ei}_{\ell,t+1} = (S^i_{b,t+1} - F^i_{o,t})/F^i_{m,t},$$

where b denotes bid, o denotes offer, and m denotes the mid price. Scaling of the position is arbitrary, so we select to keep it the same as in the no-transaction-cost

[3]These trading strategies use USD as both measurement and base currency. Because we work with excess returns, the choice of measurement currency, i.e., the currency in which the investor measures profits, is irrelevant (Daniel, Hodrick, and Lu, 2017, Online Appendix C; Chernov, Dahlquist, and Lochstoer, 2023, Internet Appendix II; Chernov, Haddad, and Itskhoki, 2024, Appendix A, Lemma 3). The choice of base currency (i.e., the currency which is the basis for the positions that a particular portfolio takes) is inconsequential for the cross-sectional strategies that we consider. With respect to the remaining strategies, one could contemplate a different base. Such strategies have been shown to have SR inferior to their counterparts with USD as the base (Lustig, Roussanov, and Verdelhan, 2014, Maggiori, Neiman, and Schreger, 2020).

case, $F^i_{m,t} = F^i_t$. The excess return on the short position, net of transaction costs, is

$$R^{ei}_{s,t+1} = -(S^i_{o,t+1} - F^i_{b,t})/F^i_{m,t}.$$

We compute strategy excess returns, net of transaction costs, using these expressions for individual returns but evaluate the impact of transaction costs for 25%, 50%, and 100% of the quoted bid-offer spreads.

3 Results

We start by providing the motivating evidence. Then we describe how the UMVE portfolio is estimated both with and without accounting for transaction costs. In the rest of the section, we use the estimated UMVE to characterize risk-return trade-off in the currency markets including conditional and unconditional analysis, different groups of currencies, and different levels of transaction costs. We conclude by exploring the role of unpriced risks in the risk-return trade-off and their origins.

3.1 Preliminary evidence

We first plot, in Figure 2, the cumulative returns of four prominent strategies across our three sets of currencies – dollar, carry, momentum, and value for the G, GE, and GEX currency sets. Following Daniel and Moskowitz (2016), we let investors start with \$1 at the beginning of the sample, December 1984. They then each month invest their wealth in the risk-free asset and take positions in currency forwards as dictated by the trading strategy at hand.

The Figure shows that there is substantial heterogeneity in the performance of the trading strategies across time and currency sets. For instance, for all currency sets the dollar strategy appears to have near zero return from the financial crisis and on. The cross-sectional carry strategy based on the GEX set appears to continue to do very well in the latter half, while the carry based on the GE set still experiences decent growth, and the carry based on the G set have little growth in the second half of the sample. Similarly, cross-sectional one-month momentum appears to flatten out much earlier for the G set (around 2000) than for the GE and GEX sets (around 2010).

Value has a relatively steady high return, especially for the GE set. The GEX set actually has the worst performance in this case, reflecting that the value signal, which is based on a notion of mean-reversion in real exchange rates, is a weaker predictor of cross-sectional differences in currency returns for the more managed currencies. In fact, there is broad agreement in the literature that real exchange rate dynamics are substantively different depending on whether nominal exchange rates are pegged or not. In particular, Mussa (1986) demonstrates that floating regimes lead to more volatile real and nominal exchange rates, which could be a reason the RER signal predicts returns better in GE set than GEX.

A natural concern is that some of the outperformance of the GE and GEX sets are a mirage and would disappear when accounting for transaction costs. We start by evaluating the impact of various fractions of the reported bid-offer spreads on the forwards and the spot exchange rates involved in the strategy trades.

Figure 3 displays average bid-offer spreads for spot and forward exchange rates. The

reported spreads are computed as follows (and expressed in basis points, bps):

$$\begin{aligned} BOS_t^i &= (S_{b,t}^i - S_{o,t}^i)/S_{m,t}^i, \\ BOF_t^i &= (F_{b,t}^i - F_{o,t}^i)/F_{m,t}^i. \end{aligned}$$

We see that early in the sample the spreads were elevated for G currencies, starting at 50 bps on average, and converging to around 5 bps by 2002 (Panel A). This pattern is primarily driven by NZD (Panel B). The spot and forward bid-offer spreads are similar in magnitude. To put these numbers into perspective, we consider strategies that experience full turnover every month, thus 5 bps translate into 0.6% lower strategy return per year. This is a substantive negative impact on returns of many strategies.

The trading costs of emerging market currencies with floating exchange rates have come down over time as well albeit later than the G10 currencies (Panel A). The bid-offer spread for forwards is roughly double that for spots. The emerging market currencies from the extended set have experienced an increase in trading costs over time. Starting with the average cost in the 10–20 bps range, these currencies end the sample at a much higher level of around 40 bps. This translates into a 4.8% annual negative impact on strategy excess returns.

The increase in the average GEX bid-offer spread is partly due to the introduction of currencies with higher bid-offer spreads later in the sample. Figure 4 shows the number of currencies in the three sets (G, GE, and GEX) over time.

Using quoted bid-offer prices might not be a realistic representation of costs that a currency market participants could be facing. They could be lower because of

trading efficiencies and special banking relationships. They could be higher because proportional costs that we consider here do not account for the price impact of trading. Rather than getting into the details of these possible effects, for which we have partial data at best, we simply consider different fractions of the quoted bid-offer spreads as a true cost of transacting in these markets. We evaluate scenarios with 25%, 50%, and 100% of the reported spreads.

Table 1 reports mean, volatility, SR, and skewness of the strategy returns across different transaction cost scenarios and across different sets of currencies included in the strategy (G, GE, or GEX). Volatility and skewness are stable across different levels of transactions costs, so changes in SRs are directly linked to changes in average returns. The two quantities are also linked via the t–statistic for significance of average excess returns, which is equal to $SR \times \sqrt{T}$, where $T = 462/12$ years. Thus, SRs with values 0.32 and below imply insignificant average returns.

We display the SRs from Table 1 in Figure 5 to visualize their differences. Focusing on SRs greater than 0.32 and starting with zero transaction costs, we observe a clear pattern of SR increase as we expand the currencies from G to GEX (CS-Mom 12 is the only exception as, despite the improvement, the largest SR is still less than 0.32). Considering the other extreme of 100% transaction costs, most SRs drop below 0.32 and, thus, average returns associated with these strategies are insignificantly different from zero. CS-carry and TS-carry are the two strategies whose returns endure through these high costs. But the benefits of the GEX set disappear as the SRs are similar to the ones from the GE set.

At the intermediate costs (25% and 50%), there is no general discernible advantage to the GEX set over the GE set. A stark exception is CS-Carry, where the SR is

much higher in the case of GEX. TS-Mom 1 is another strategy where considering the GEX set seems to be advantageous.

The observations on transaction costs and their impact on strategy returns pose an important challenge to researchers attempting to understand the risk-return trade-off in currency markets. Ignoring transaction costs might impose unrealistic burden for candidate models to explain the cross-section of strategy returns. Yet, incorporating costs that are too high leaves nothing to be explained and, thus, lowers the burden on models substantially. Keeping this in mind, we continue using different fractions of reported bid-offer spreads as we proceed with evaluation of the risk-return trade-off in these markets.

We note that the return volatility is stable across different assumptions for transaction costs. This suggests that the modeling of the impact of transaction costs on the currency return covariance matrix is not material, and, therefore, we will simply use midpoint prices for estimation of variances and correlations in the subsequent analysis.

While the GEX set appears to be offering superior strategy performance in certain scenarios, ultimately the question is whether the efficient frontier associated with these currencies improves upon the frontier associated with the smaller sets of G or GE. We proceed with this analysis in the next section.

3.2 Estimating the mean-variance efficient portfolio

This section covers our methodology. First, we explain how we estimate the UMVE returns. Second, we describe how we use the UMVE to test pricing of various cur-

rency strategies. The methodology is different depending on whether we ignore transaction costs or account for them.

Without transaction costs

We refer the reader to Chernov, Dahlquist, and Lochstoer (2023) for details, motivation, and additional citations. We seek to correctly price currency risks both conditionally and unconditionally. As pointed out by Hansen and Richard (1987) and Jagannathan (1996) one can achieve this by constructing the UMVE portfolio.

Specifically, suppose we have N basis assets with an $N \times 1$ vector of excess returns R^e_{t+1}. The conditional mean of this vector is $\mu_t = E_t\left(R^e_{t+1}\right)$ and its conditional covariance matrix is $\Sigma_t = V_t\left(R^e_{t+1}\right)$. An admissible trading strategy p in the basis assets has an $N \times 1$ vector of weights w_{pt} that are determined based only on information available up until time t. The resulting excess portfolio return is then $R_{p,t+1} = w_{pt}^\top R^e_{t+1}$.

A conditional mean-variance efficient (CMVE) portfolio is a dynamic trading strategy in these assets that obtains conditional MSR. Because leverage does not affect the conditional SR, any portfolio with weights proportional to $\Sigma_t^{-1}\mu_t$ would be CMVE. The UMVE portfolio is a dynamic trading strategy in the same set of assets that obtains the MSR, both conditionally and unconditionally. Ferson and Siegel (2001) and Jagannathan (1996) show that the UMVE portfolio weights are:

$$w_t^* = \frac{1}{1 + \mu_t^\top \Sigma_t^{-1} \mu_t} \Sigma_t^{-1} \mu_t. \tag{2}$$

That is, the UMVE portfolio is constructed by optimally timing the CMVE portfolio

with asset weights $\Sigma_t^{-1}\mu_t$ using the scalar leverage factor $(1+\mu_t^\top\Sigma_t^{-1}\mu_t)^{-1}$ such that the maximal unconditional SR is obtained.

The UMVE portfolio accounts for all risks conditionally in the sense that the following conditional linear beta pricing relationship holds for any admissible strategy p:

$$E_t(R_{p,t+1}) = \beta_{pt}E_t(R^*_{t+1}), \tag{3}$$

where $\beta_{pt} = Cov_t(R_{p,t+1}, R^*_{t+1})/V_t(R^*_{t+1})$. The UMVE portfolio also implies the unconditional linear beta pricing relationship:

$$E(R_{p,t+1}) = \beta_p E(R^*_{t+1}), \tag{4}$$

where $\beta_p = Cov(R_{p,t+1}, R^*_{t+1})/V(R^*_{t+1})$, for any p. All CMVE portfolios satisfy the conditional linear pricing relationship in Equation (3). The UMVE portfolio is the only CMVE portfolio that satisfies the unconditional linear pricing relationship in Equation (4).

Hansen and Richard (1987) show that one can evaluate all the conditional implications of Equation (3) by testing Equation (4) using all admissible trading strategies as test assets. Because our model's factor represents a return on a traded asset (the UMVE), the model implies that $\alpha_p = 0$ in the time-series regression

$$R_{p,t+1} = \alpha_p + \beta_p R^*_{t+1} + \varepsilon_{p,t+1} \tag{5}$$

for each test asset p (e.g., Cochrane, 2005, Section 12.1). From an economic perspective, this test evaluates whether the UMVE portfolio is the MSR portfolio by evaluating whether the maximal sample Sharpe ratio combination of the UMVE and

the test assets is significantly different from the sample Sharpe ratio of the UMVE:

$$\text{MSR}^2(R^*, R_p) - \text{MSR}^2(R^*) = \alpha_p^\top \Sigma_\varepsilon^{-1} \alpha_p, \tag{6}$$

where Σ_ε is the covariance matrix of the regression residuals.

We largely follow Chernov, Dahlquist, and Lochstoer (2023) to estimate μ_t and Σ_t in an out-of-sample (OOS) fashion. The conditional mean is estimated as a linear function of three characteristics: the forward discount (equivalent to the interest rate differential if CIP holds), the five-year change in the real exchange rate (a measure of value or mean-reversion), the one-year change in the nominal exchange rate (a measure of momentum or trend). We modify the procedure for the covariance matrix to ensure that it applies to a larger and unbalanced panel of currencies. See Appendix A.4 for further details. Denote the excess return on this portfolio $R^*_{t+1} = w_t^{*\top} R^e_{t+1}$.

To appreciate the value added of this portfolio vis-a-vis the traditional characteristic-based factors, consider conditions which ensure that a CMVE is spanned by characteristics-based factors. If the conditional mean is linear in characteristics, the returns on these factors should be

$$R^F_{t+1} = A_t^\top C_t^\top \Sigma_t^{-1} R^e_{t+1}, \tag{7}$$

where C_t is an $N \times J$ matrix of characteristics, and A_t is an arbitrary $J \times J$ transformation matrix.[4] These factors are different from the characteristic-based portfolios considered in the literature, which use characteristics similar to ours. For example, the seminal dollar-carry model uses a constant and the forward discount as char-

[4] If $A_t = (C_t^\top \Sigma_t^{-1} C_t)^{-1}$ one obtains 'GLS factors' considered by Daniel, Mota, Rottke, and Santos (2020). $A_t = I_J$ leads to 'MVE' factors discussed by Kozak and Nagel (2023).

acteristics ($J = 2$). The corresponding dollar and carry factors span the CMVE if $\Sigma_t \propto I_N$ in Equation (7), or, more generally, if characteristics span covariances of excess returns with characteristic-based factors. These requirements are unlikely to hold in the data.

Thus, even if a researcher prefers a factor-based approach, she needs to estimate the covariance matrix Σ_t to map the characteristics into factors. Further, the factors would only span a CMVE portfolio, which is insufficient for the tests of the unconditional linear beta pricing model. To achieve the UMVE it is necessary to estimate loadings of expected excess returns on the characteristics. This step is also useful for verifying the implicit assumption in Equation (7) that characteristics indeed describe expected excess currency returns.

It is not clear a-priori that our estimates of μ_t and Σ_t, which are constructed using information available only up until time t, are correct. That is, there is no guarantee that the resulting UMVE portfolio would have the ex ante MSR, which would lead to non-zero α_p. Therefore, we validate our estimates of μ_t, Σ_t, and the resulting UMVE weights by performing standard GRS joint tests of $\alpha_p = 0$ across a set of most promising trading strategies p proposed in the literature.

Affleck-Graves and McDonald (1989) and Zhou (1993) point out that the GRS test tends to overreject the null hypothesis when excess returns are not normally distributed. Currency returns are definitely non-normally distributed, as documented in Table 1 and elsewhere in the literature (see, e.g., Chernov, Graveline, and Zviadadze, 2018, Farhi, Fraiberger, Gabaix, Ranciere, and Verdelhan, 2009). Thus, we report bootstrapped p–values along with the traditional GRS ones. Our bootstrap procedure is nested within the one developed for non-zero transaction costs, so we

postpone its description until the next section.

With transaction costs

The closed-form expression for the UMVE portfolio weights in Equation (2) no longer applies if there are transaction costs. The reason is that absolute values of returns on the long and short positions in the same asset are no longer identical. Thus, the UMVE portfolio has to be constructed numerically, and we have to use bootstrap to test whether it prices various trading strategies correctly. We apply the approach of Detzel, Novy-Marx, and Velikov (2023) to the former problem, and develop our own test for the latter problem.

Because long and short positions no longer produce the same returns, up to the sign, we consider a $2N$-dimensional vector of excess returns, $\tilde{R}^e_{t+1}$. It contains excess returns corresponding to long positions in individual currencies, net of transaction costs, $R^e_{\ell,t+1}$ and excess returns corresponding to short positions, net of transaction costs, $R^e_{s,t+1}$. We denote the corresponding $2N \times 1$ vector of expected excess returns by $\tilde{\mu}_t$ and the $2N \times 2N$ covariance matrix of excess returns by $\tilde{\Sigma}_t$. The $2N \times 1$ vector of portfolio weights $\tilde{w}^*_t$ that maximizes the conditional SR, solves the following problem:

$$\text{MSR}(\tilde{R}^*_{t+1}) = \max_{\tilde{w}^*_t} \frac{\tilde{w}^{*\top}_t \tilde{\mu}_t}{\left(\tilde{w}^{*\top}_t \tilde{\Sigma}_t \tilde{w}^*_t\right)^{1/2}}, \tag{8}$$

which we solve numerically at each time t. Li, DeMiguel, and Martin-Utrera (2023) demonstrate that the MSR remains a sufficient statistic for for the investment opportunity set in the presence of proportional costs (their Proposition 2).

Generically, a solution to this problem would yield CMVE portfolio weights. As is

the case with the no-transaction-costs case, we are interested in testing the portfolio using unconditional moments and, thus, require the UMVE portfolio. To ensure that we obtain the UMVE weights, we scale the candidate $\tilde{w}_t^*$ by $\left(1+\tilde{\mu}_t^\top \tilde{\Sigma}_t^{-1}\tilde{\mu}_t\right)^{-1}$ following Equation (2).

In the case with transaction costs, one cannot implement and reasonably interpret a regression along the lines of Equation (5). For example, if a test asset p happens to have a negative beta with respect to the UMVE, one has to construct a short version of the UMVE, net of trading costs, but this is impossible to determine before running the regression. Thus, we implement a different test that captures the economic interpretation of GRS in Equation (6). Thus, we test whether

$$\mathrm{MSR}^2(\tilde{R}^*, R_p) - \mathrm{MSR}^2(\tilde{R}^*)$$

is significantly different from zero. Here, the SRs are constructed as the ex post MSR combination of the assets taking into account the transaction costs.[5] Economically, this is the same test as the GRS test, which also tests whether the ex post MSR combination of the test assets and the factor(s) is greater than that of the factor(s). In particular, we impose the null hypothesis by subtracting from each strategy the alpha in a regression of the net of transaction costs return on the strategy on the net of transaction costs UMVE return. This ensures that all assets other than the

[5] An emerging literature studies the effect of netting across the combination of different factors and the resulting potential savings in transactions costs (Baldi-Lanfranchi, 2024, DeMiguel, Martín-Utrera, Nogales, and Uppal, 2020). The economic objective is to construct an optimal portfolio accounting for transactions costs. Our objective is to evaluate whether one can improve SR by optimizing over test asset returns and the asset that already solves the optimal portolio problem with transaction costs. In our framework, such optimal portfolio is the UMVE portfolio that solves the problem in Equation (8). Thus, we treat test assets (currency strategy) as a monolith when constructing $\mathrm{MSR}^2(\tilde{R}^*, R_p)$ and do not consider potential netting across them. If we were to do so, we would end up with the original problem.

UMVE will have a zero weight in the MSR computation also in the transaction cost case. We then draw from these net-of-alpha returns to get the distribution of our test statistic under the null hypothesis that the test assets cannot increase the SR relative to the UMVE factor.[6]

The benefits of the UMVE-based approach

A typical approach in the currency literature is similar to that of Fama and French (1993), which was developed for equities. That is, researchers select characteristics thought to be related to expected return – interest rate differential (carry), prior recent returns (momentum), or the real exchange rate (value). Then they form factors based on these characteristics by going long currencies with a high value of the characteristic and short currencies with a low value of the characteristic. As forcefully argued by Daniel, Mota, Rottke, and Santos (2020) in the context of equities, such construction may contaminate factors with "unpriced risks."

Consider a simple example to appreciate the importance of this issue and how it can be rectified via the UMVE-based analysis. Assume two currencies i and j with excess returns:

$$R_t^i = \frac{1}{2}F_t + \beta_i U_t, \qquad R_t^j = -\frac{1}{2}F_t + \beta_j U_t,$$

where F_t and U_t are traded excess return factors with unit variance. Further, assume that F_t is "priced" with $E_t\left(F_{t+1}\right) = 1$, and U_t is "unpriced" $E_t\left(U_{t+1}\right) = 0$.

[6]There is a number of differences in our implementation as compared to that of Detzel, Novy-Marx, and Velikov (2023) besides different asset classes. Detzel, Novy-Marx, and Velikov (2023) maximize unconditional (in-sample) SR only. Further, because they are comparing different factor models' asset returns, their bootstrap tests are designed to select the best factor model, and thus have a different design. Resampling is a correct procedure if the excess returns are i.i.d. If they are not, our procedure biases towards overrejection.

Suppose a researcher picks a characteristic with values equal to the asset returns' loadings on F_t. The standard characteristic-based portfolio sort then constructs the factor as long the high characteristic asset and short the low characteristic asset:

$$R_t^i - R_t^j = F_t + (\beta_i - \beta_j)\, U_t.$$

The long-short portfolio return is exposed to the unpriced risk if $\beta_i \neq \beta_j$. The variance of the unpriced factor will then contribute to the variance of the long-short portfolio without increasing the expected return of the long-short portfolio. As a result, the SR of the long-short portfolio will be lower than for a portfolio without exposure to the unpriced risk.

Let's instead consider a factor based on the MVE portfolio weights. The mean and covariance of the asset returns are:

$$\mu_t = \begin{bmatrix} \frac{1}{2} \\ -\frac{1}{2} \end{bmatrix}, \quad \Sigma_t = \begin{bmatrix} \frac{1}{4} + \beta_i^2 & -\frac{1}{4} + \beta_i\beta_j \\ -\frac{1}{4} + \beta_i\beta_j & \frac{1}{4} + \beta_j^2 \end{bmatrix}.$$

Then,

$$\Sigma_t^{-1}\mu_t = \begin{bmatrix} 2\frac{\beta_j}{\beta_i+\beta_j} \\ -2\frac{\beta_i}{\beta_i+\beta_j} \end{bmatrix}.$$

Applying these optimal weights to the currency excess returns, we have:

$$2\frac{\beta_j}{\beta_i + \beta_j} R_t^i - 2\frac{\beta_i}{\beta_i + \beta_j} R_t^j = F_t.$$

That is, the unpriced risk is hedged out by judicious weighting of the assets. Hence, the MVE portfolio has only priced risk, and therefore the highest SR.

Any candidate explanation of a strategy performance must be related to the priced, rather than unpriced, component of that strategy returns. Therefore, UMVE-based analysis delivers the relevant reference measurements directly.

3.3 Pricing performance

In this section we describe asset pricing tests of the estimated UMVE. We vary the considered set of currencies, level of transaction costs, and whether we are testing unconditional or conditional model implications. As an alternative to the UMVE, we also consider a standard factor model used in the literature as the model for cross-sectional currency pricing.

Unconditional model implications

As a starting point, we consider the UMVE portfolio constructed from G currencies alone (in the absence of transaction costs). First we check if this UMVE portfolio can price strategies constructed from G currencies only. This is the asset pricing exercise performed by Chernov, Dahlquist, and Lochstoer (2023), so our main purpose here is to confirm that the conclusions of that study still hold in an updated sample. Appendix A.5 demonstrates that this is indeed the case. We note that the SR of this UMVE is 1.02. Second, we demonstrate in the same Appendix that this UMVE cannot price strategies constructed from GE currencies. Thus, we do not consider the UMVE constructed from G currencies going forward.

Moving on to the UMVE portfolio constructed from GE currencies, we evaluate if it can price strategies constructed from GE currencies (in the absence of transaction

costs). Panel A of Table 2 displays test results for individual strategies. None of the alphas are significant, and the adjusted R^2s are quite low indicating that the strategies are exposed to substantial unpriced risks. The presence of the unpriced risk does not indicate that the UMVE is misspecified. Indeed, because all alphas are insignificant, the model does price the cross-section of test assets correctly. These conclusions echo those of Chernov, Dahlquist, and Lochstoer (2023) for G currencies. The leftmost column of Table 3 shows that the bootstrap p–value for the joint test that the alphas equal zero is 0.653 (0.514 for the standard GRS-test), indicating that the model prices the strategies at a comfortable level. The SR of the UMVE portfolio is 1.341, much higher than those of the individual strategies and for the UMVE constructed from G currencies only. The ex post MSR from combining the test assets with the UMVE portfolio is 1.431. Thus, there is a relatively small increase, consistent with the high p–value.

Next, we check if the same UMVE portfolio can price the more challenging set of strategies constructed from GEX currencies. Panel B of Table 2 displays test results for individual strategies. With the exception of TS-Mom 1, all alphas are insignificant, suggesting that the portfolio can handle a much larger set of currencies. The bottom part of the leftmost column of Table 3 shows that the model is not rejected on the basis of the bootstrap p–value of 0.118 in the joint test that all alphas are zero. The GRS p–value of 0.05 indicates a marginal rejection. As we have observed earlier, transaction costs are much larger for the 'X' currencies in the GEX set. Thus, one might anticipate that these p-values will be higher at realistic levels of bid-offer spreads.

As a warm-up, we test whether the GE-based UMVE can price the GE-based portfolios accounting for transactions costs. Because transaction costs make returns less

challenging to explain, one would expect success given that the model could do it without costs with a comfortable p–value above 0.5. However, this is not a priori obvious as the UMVE portfolio also suffers from transaction costs. The three rightmost columns of Table 3 display test results for fractions of bid-offer spread at 25%, 50%, and 100%, respectively. We see that as trading costs increase, the UMVE SR drops to 1.203, 1.157, and 0.974 respectively. However, if anything, it becomes easier and easier to explain strategy returns with bootstrapped p–values ranging from 0.573 for 25% to 0.927 for 100%.

Testing whether the GE-based UMVE can price GEX-based portfolios is more interesting because the zero-cost p–values are lower. The last rows of the rightmost three columns of Table 3 report the results. Here we fail to reject the model even when costs are at 25% of bid-offer spreads (the bootstrapped p–value is 0.200). Recalling the evidence in Table 1, the strategy performance in GEX sample is still formidable at 25% costs. The SR are often the largest for carry and one-month momentum strategies. Thus, the failure to reject the UMVE constructed from a much smaller sample of GE currencies is economically significant.

Conditional model implications

Since the UMVE portfolio should price assets both conditionally and unconditionally, we can easily implement a test of the model's conditional implications. This test has the additional advantage of testing the model entirely OOS. We consider only the zero transaction cost case for these tests.

In particular, to implement the OOS test of the UMVE portfolio, we follow Chernov, Dahlquist, and Lochstoer (2023) and exploit the fact that the conditional linear beta

pricing model (3) holds. Therefore, if we remove the priced component of a portfolio return, the residual

$$R_{u,t+1} = R_{p,t+1} - \beta_{pt} R^*_{t+1}, \tag{9}$$

referred to as the unpriced component of a portfolio, should have zero alpha. The conditional beta is equal to

$$\beta_{pt} = \frac{w_{pt}^\top \Sigma_t w_t^*}{w_t^{*\top} \Sigma_t w_t^*}. \tag{10}$$

Because of all its ingredients are known in real time, the beta and unpriced component can be computed in real time as well. Thus, testing if $E(R_{u,t+1}) = 0$ (zero alpha) is an unconditional test of the model's conditional implications and also amounts to an OOS test of the model.

Table 4 reports the test results applied to $R_{u,t+1}$ when test portfolios are constructed from GE currencies. None of the alphas are significant at the 5% level. That is, the returns that obtain from hedging out the priced component, $\beta_{pt} R^*_{t+1}$, indeed have statistically zero average return (alpha) and SRs (SR-hedged, reported in the rightmost column). The p–value for the test that alphas are jointly zero is 0.379, thus failing to reject the model.

The return $R_{u,t+1}$ represents unpriced risks in each strategy's return. Figure 6 shows the SR of each strategy, as well as the corresponding SRs of the strategy's unpriced component, $R_{u,t+1}$, and the portfolio where unpriced risks are hedged out, $R_{p,t+1} - R_{u,t+1}$, which is equal to $\beta_{pt} R^*_{t+1}$. Note that the construction of $\beta_{pt} R^*_{t+1}$ is done in real time and thus both unpriced and priced components of portfolios represent returns

tradeable on tradeable portfolios.

All the unpriced components have SRs close to zero, while the priced components of portfolios generally have much higher SRs than their original counterparts. For instance, the cross-sectional carry strategy goes from a SR of 0.71 to 1.29 when unpriced risks are hedged out, the dollar factor goes from 0.33 to 0.91, the 12-month cross-sectional momentum strategy from 0.24 to 0.99, and the cross-sectional value goes from 0.65 to 1.29. These large differences suggest that the the standard strategies employed in the literature are not suitable to use as factors for risk-based pricing due to their contamination from these substantial unpriced risks.

It is important to remember that although the hedged strategies are more attractive than their original counterpart, the UMVE portfolio is the best alternative to any of the original strategies. This is because the UMVE has the largest SR. Indeed, even though some of the hedged strategies have high SRs (CS-Carry with SR=1.29, TS-Carry with SR=1.31, and CS-Value with SR=1.29), the UMVE's SR of 1.34 exceeds all of them.

As an additional illustration of the model's conditional pricing implications, Figure 7 displays the average return and average predicted return, $\beta_{pt}E_t(R^*_{t+1})$, for four of the trading strategies. The top panel gives the results from the first half of the sample, while the bottom panel gives the results for the second half of the sample. The average realized and predicted returns are much higher in the first half of the sample than the last half, with the exception of the cross-sectional carry strategy. Thus, the model can account for the trends in the risk premiums over the sample, as also seen in Figure 2.

3.4 The role of emerging markets

Our evidence suggest that there is a large economic benefit to expanding the set of G currencies to GE. The SR of the corresponding UMVE increases from 1.02 to 1.34. Further, the latter UMVE can price not only strategies constructed from GE currencies but also those from 59 currencies in GEX. This is a remarkable result as the UMVE portfolio in theory prices only strategies based on currencies used for the UMVE portfolio construction. Thus, the implication of the finding is that GEX-based strategies are not exposed to new priced risks as compared to the GE-based ones. There is little advantage to trading currencies from the extended emerging set. We elaborate on these findings by considering their implications from different angles.

We start by evaluating the role of currencies from emerging markets by revisiting the motivational Figure 2 through the lens of the UMVE portfolios. Specifically, Figure 8 compares cumulative returns on the priced components of the four strategies using the UMVE portfolio based on G and GE to hedge out the unpriced component. We compare these two versions of the priced component to GEX-based full strategies as they appear to be the most profitable (before transaction costs).

The relative performance of unhedged GEX-based and hedged G-based strategies is mixed. The former dominates in the cases of dollar and carry, while the latter dominates in the cases of momentum and value. This mixed performance suggests that emerging currencies contain interesting return potential. The hedged GE-based strategies fully dominates both hedged G-based and unhedged GEX-based strategies by a wide margin.

As we noted earlier, the UMVE portfolio itself is in fact the best alternative to the

individual strategies. We plot cumulative returns of the GE-based UMVE in all four panels of Figure 8. We see that the UMVE dominates all strategies but carry and especially value do come close to the efficient portfolio.

Figure 9 illustrates the impact of transaction costs (for the case of 50% of quoted bid-offer spread as an example). We report the same four original strategies for all three sets: G, GE, and GEX. Similar to our observations about SR in Figure 5, we see that transaction costs lead to deterioration in the GEX-based performance, although momentum and especially carry continue to outperform their counterparts based on smaller currency sets. We also transaction-cost-optimized GE-based UMVE as per Equation (8). The most important take-away from the figure is that the UMVE is not impacted by transaction costs as much as the strategies. In fact, the gap between the UMVE and carry or value has increased as compared to Figure 8.

Taken together, this evidence illustrates the significant economic impact of combining currencies into the efficient portfolio. We can also conclude that the main advantage of considering emerging economies lies in the small subset of the least managed currencies.

Figure 10 illustrates the improved conditional risk-return trade-off after adding the floating-regime emerging currencies. It displays the conditional MSR for the UMVE based on G currencies versus that based on GE currencies. The former has a strong downward trend over the sample, while the latter does not to the same extent.

The documented high MSR are driven to a large extent by the strength of the carry strategy for the emerging currencies. Figure 11 visualizes this. Panel A shows that the conditional mean of the CS-Carry for GE currencies is much higher than that of the G currencies when data on all currencies are available. Panel B shows that the

volatility of the CS-Carry is roughly the same across the G and GE currencies (if anything the GE volatility is slightly lower). Taken together, while the conditional mean of G carry falls during the mid 2000s the volatility stays roughly at the same level. This is not the case for GE carry where the mean stays at roughly the same level. Consistent with Andrews, Colacito, Croce, and Gavazzoni (2024), these differences can be understood in the context of the post-GFC monetary policy in the G economies which lead to overall compression in interest rates and their differences – the key driver of the carry signal.

3.5 Interpreting the evidence

In this section, we show that our results are robust to currency perspective and highlight how our results indicate large common sources of unpriced risks in currency returns. In terms of the currency perspective, it may be unintuitive for some to think of currency risks as unpriced since currency depreciation rates show up in the SDF when one changes currency numeraire. For instance, when going from a USD SDF to an AUD SDF, the AUD versus USD depreciation rate shows up in the AUD SDF and is thus priced from the AUD perspective. We show next that the unpriced risks we document remain large when viewed from other currency perspectives, despite the fact that depreciation rates show up in the SDF when one changes numeraire. We follow Backus, Foresi, and Telmer (2001) and Lustig, Roussanov, and Verdelhan (2011) and use an affine no-arbitrage model to work through this line of reasoning. As a part of such a model, we switch to logs from levels. For clarity of exposition, we ignore conditional means of variables. As a result, depreciation returns and currency excess returns coincide.

We assume that the innovation in the log depreciation rate of country i relative to USD is

$$\widetilde{\Delta s}_t^i \equiv \Delta s_t^i - E_{t-1}(\Delta s_t^i) = \beta_{it-1} f_t + \delta_i^\top u_t + \sigma_i \varepsilon_{it}.$$

Here, f_t represents innovation in the UMVE portfolio expressed in terms of USD:

$$\widetilde{m}_t \equiv m_t - E_{t-1}(m_t) = -b f_t,$$

and u_t represents common unpriced sources of variation in currencies, while ε_{it} are cross-sectionally uncorrelated currency-specific innovations. Without loss of generality, we assume that elements of u_t are orthogonal to each other.

We can then express the SDF in any other currency via

$$\begin{aligned} \widetilde{m}_t^i &= \widetilde{m}_t + \widetilde{\Delta s}_t^i = -b f_t - \widetilde{\Delta s}_t^i \\ &= -(b + \beta_{it-1}) f_t - \delta_i^\top u_t - \sigma_i \varepsilon_{it}. \end{aligned}$$

When one changes the currency perspective, the SDF "gains" an additional local term.

To gauge the size of the local SDF term, we consider the evidence. The volatility of the depreciation rate is about 0.1 per annum. Thus, the volatility of the component of the depreciation rate that is orthogonal to the UMVE, $V^{1/2}(\delta_i^\top u_t + \sigma_i \varepsilon_{it}) \leq 0.1$. Our evidence in Table 3 implies that the volatility of the SDF, as proxied by the SR of the UMVE, $V^{1/2}(\widetilde{m}_t)$, is greater than 1. Thus, the local SDF component we gain when changing currency perspective is small relative to the common term arising from the USD perspective. In fact, it would take about 400 years of data to establish

that the maximal SR this extra term can give rise to is significantly different from zero.[7]

Next, we explain how our evidence implies large unpriced common components in currency returns. Consider a trading strategy, p, with weight in currency i at time t equal to ω_{it}. Shocks to the returns to this strategy are then

$$\begin{aligned}\widetilde{r}_{p,t+1} &= \sum_{i=1}^{n} \omega_{it}\beta_{it} f_{t+1} + \sum_{i=1}^{n} \omega_{it}\delta_i^\top u_{t+1} + \sum_{i=1}^{n} \omega_{it}\sigma_i \varepsilon_{i,t+1} \\ &\approx \beta_{pt} f_{t+1} + \delta_{pt}^\top u_{t+1},\end{aligned}$$

where $\beta_{pt} = \sum_{i=1}^{n} \omega_{it}\beta_{it}$, $\delta_{pt} = \sum_{i=1}^{n} \omega_{it}\delta_i$, and we assume that idiosyncratic currency shocks are approximately diversified away. Panel A of table 2 implies that $V(\delta_{pt}^\top u_{t+1})/V(\widetilde{r}_{p,t+1}) \geq 0.7$. That is, there are large, common unpriced risks u_t shared across currencies that appear in all the strategy portfolios. A natural question is what these unpriced risks represent, which we discuss after analyzing the implications for these unpriced factors for existing currency risk models.

3.6 Implications of substantial unpriced risks

This section delves deeper into the implications of our evidence. The presence of large unpriced components in famous strategies suggests that hedging out these components would lead to new test assets that are substantially more informative about the pricing channels of currency risks. We explore this line of reasoning by considering

[7] The standard error for a sample mean, μ, equals $\sigma/\sqrt{T}$. The t-statistic for testing whether this sample mean is different from zero is then $\sqrt{T} \times \mu/\sigma = \sqrt{T} \times SR$. Since the annualized volatility of the extra term in the SDF is about 0.1, the maximal annualized SR it can give rise to is also 0.1. Thus, we need $T = 400$ to obtain a t-statistic of 2.

two benchmarks: the dollar-carry factor model and the consumption-based model. Large unpriced components also raise a question of their economic origins. That is, the currency strategies we consider are well-diversified, especially in the GE and GEX sets, so the unpriced components must be driven by factors that are important for currency covariation but not risk premiums. It is then informative, from a theory viewpoint, to better understand what types of economic shocks lead to such factors. We study this question in the last subsection.

Evaluating the dollar-carry model

The seminal dollar-carry model of Lustig, Roussanov, and Verdelhan (2011) may serve as a natural alternative to the analysis considered here. However, this model is typically rejected using test assets beyond the various flavors of carry. Liu, Maurer, Vedolin, and Zhang (2023) hypothesize that the documented rejections could be due to the unconditional nature of the model (constant betas). Thus, they advocate a conditional version (time-varying betas). The unconditonal and conditional versions coincide only if the model factors — dollar and carry — span the UMVE portfolio unconditonally. In this section we test the unconditional and conditional versions of the dollar-carry model on the basis of the GE-based portfolios.

The test of the unconditional model is implemented via "alpha" regressions leading to the GRS and bootstrap tests just like elsewhere in the paper. Table 5 reports both individual alphas and the p–values for the joint test. The model matches the returns to dollar and cross-sectional carry perfectly. This is mechanical given that the dollar and carry are both on the left- and right-hand sides of the regression. The model is struggling with dollar carry, both types of time-series momentum, and value. The

p–value for the joint test is basically 0, strongly rejecting the model.

The idea of the test of the conditional model is similar to that of the UMVE test we have considered earlier in Equation (9). Under the null of the conditional two-factor dollar-carry model, the excess returns on the residual portfolio

$$R_{u,t+1} = R_{p,t+1} - \beta_t^{F\top} R_{t+1}^F$$

should be equal to zero, on average; here β_t^F is a bivariate vector of conditional dollar and carry betas, and R_{t+1}^F is a vector of corresponding factor returns. Conditional betas can be computed in real time via Equation (10).

Table 5 gives the alphas of the individual strategies (see the column labeled OOS α) and the joint test. As is the case with the unconditional model, dollar and cross-sectional carry should be fitted perfectly, and they are. Furthermore, dollar carry is also spanned conditionally. Thus, a zero alpha for this strategy is mechanical as well. The model has trouble matching the time-series 12-month momentum and cannot price the cross-sectional value strategy. For that reason, the p–value is 0.017 and the model is rejected in the joint test.

Both tests indicate that test portfolios other than the ones related to dollar and carry contain important information that clashes with the factors used in either version of the model. One concern that might arise is that the model is rejected on the basis of just a few portfolios and, thus, might serve as an attractive alternative to the more involved UMVE approach. First, we note that testing the conditional version of the model requires estimates of all trading strategies' conditional factor betas. A natural way to implement this is to compute the conditional covariance matrix of currency returns, which is a key ingredient to the UMVE portfolio. Second, one way

to address the concern about too few portfolios that are not mechanically related to the model factors is to include the new information from the UMVE and the importance of hedging out unpriced risks from test strategies. Thus, we turn to the implications arising from these findings.

To this end, we test if either the unconditional or conditional dollar-carry model can price the UMVE portfolio. This is a model test in the spirit of Barillas and Shanken (2017). The results are reported in the last row of Table 5. The alphas are in both cases highly significant with t-statistics greater than 6. The SR of the UMVE portfolio after hedging out the dollar and carry factors in real-time is 1.06. In other words, since the original UMVE SR is 1.34, the dollar-carry model still leaves a lot of risk unexplained.

Further, in Table 6 we test the dollar-carry model on the unpriced components of the trading strategies. These have, as we have shown earlier, statistically zero average returns and SRs and thus should be easy targets for pricing. Nevertheless, the dollar-carry model is rejected both in its unconditional and conditional versions, with many individual alphas being strongly significant – especially those related to Dollar and Carry strategies.

The reason for this is two-fold. First, since there are large unpriced components in the Dollar and Carry pricing factors, these factors are significantly positively correlated with the unpriced components of many of the test assets. Second, the average returns to the test assets are all close to zero, while the average returns to the pricing factors are large and positive. These two facts combined lead to several strongly negative and significant alphas, as the regressions match the averages of the left- and right-hand sides.

Thus, while a model with the standard dollar and carry factors on the right-hand side can account for the original dollar, carry, and dollar carry trading strategies, it cannot account for the the priced risk, as measured by the UMVE portfolio, and it cannot account for the unpriced components of the trading strategies either. A corollary to this is that the two first principal components, which are closely related to the standard dollar and carry factors, do not fully capture the determination of currency risk premiums.

Consumption risk and currency pricing

We evaluate both short- and long-run consumption growth as the factor determining currency risk premiums, motivated by the models of Lustig and Verdelhan (2007) and Zviadadze (2017). Burnside (2011) questions some of the findings on the grounds of imprecisely estimated betas. Under the null hypothesis of the consumption model, betas should be estimated more accurately if test assets are stripped of unpriced risks. This is because the unpriced risk is unrelated to consumption growth under the null and thus only injects noise in the estimation. Of course, this is an empirical question, and we therefore test whether consumption is related to currency risk premiums using both the original trading strategies, as well as their priced and unpriced components. We enrich the earlier studies by considering the nine most profitable strategy returns as opposed to portfolios that are constructed along a single characteristic (the interest rate differential for carry).

We measure short-run consumption as contemporaneous quarterly, per capita real services and nondurables aggregate U.S. consumption. Long-run consumption growth is the sum of short-run consumption growth over 12 quarters, consistent

with Parker and Julliard (2005). Since consumption data is quarterly, we consider quarterly strategy returns in this analysis.

The evidence regarding the short-run consumption is mixed, and we report it in Appendix A.6. We report the evidence regarding the long-run consumption in Table 7. In particular, we regress strategy returns, and their priced and unpriced components, on long-run consumption growth with the focus on the significance of the consumption betas. When we use returns on the original strategies as test assets, only CS-Value has a significant exposure to consumption growth. In contrast, when we consider the priced components of the same strategies, where unpriced risk is hedged out, 6 out of 9 assets have significant betas. Finally, we verify, consistent with the null of the consumption model, that there is no significant exposure of the unpriced components to consumption.

Adjusted R^2 from time-series regressions are largely negative when we use original strategy returns or their unpriced components. In contrast, there is only one borderline negative adjusted R^2 for priced components. This evidence further supports the existence of relationship between long-run consumption growth and priced components of the strategy returns.

Detecting the connection between currency risk premiums and consumption risk is challenging due to substantial unpriced risk present in popular trading strategies, despite their high SRs. Hedging out this unpriced risk significantly improves the precision of the pricing signal and allows us to detect the economically appealing relationship between the consumption risk and currency risk premiums.

Implications for the unpriced risks

In this subsection we study the large unpriced risks that appear in common currency trading strategies. We propose a more intuitive description of what these risks actually represents by relating them to geographically-based currency factors. For brevity, in this subsection we focus our attention on the cross-sectional carry since this is the strongest strategy for emerging markets and, indeed, across all markets for the second half of the sample.

In the equity markets, Daniel, Mota, Rottke, and Santos (2020) argue that the industry risks are unpriced risks affecting standard equity expected-return factors. We evaluate whether a similar phenomenon could arise in currency markets based on geographic proximity (see, Lustig and Richmond, 2020, and Richmond, 2019).

To this end, we construct a Europe factor, which is long an equal-weighted portfolio of European countries, and a Rest of the World (ROW) factor, which is long an equal-weighted portfolio of the remainder of the GE countries.[8] Next, at each time t we conditionally project the returns of the UMVE-based unpirced component onto the returns of these two portfolios. That is, we use the conditional covariance matrix to construct the conditional betas of the unpriced component of each strategy return on the two geographic portfolio returns. Finally, we construct "Geo-hedged" portfolio returns by subtracting the geographically based hedging portfolio from the original strategy returns:

$$R^{Geo}_{p,t+1} = R_{p,t+1} - \beta^{Eur}_{p,t} R^{Eur}_{p,t+1} - \beta^{ROW}_{p,t} R^{ROW}_{p,t+1}.$$

[8] A Europe factor has been proposed in prior research to account for currency comovements (see, Aloosh and Bekaert, 2022, and Greenaway-McGrevy, Mark, Sul, and Wu, 2018).

We note that the time t betas are, as before, constructed only using information available up until time t. The interpretation of the regression coefficients is that one goes long a notional of $\$\beta^{Eur}$ of the Europe portfolio and $\$\beta^{ROW}$ of the Rest-of-the-World portfolio to construct the geographically-based hedge for a \$1 long notional of the strategy portfolio.

The projected betas are shown in Figure 12. The hedging portfolio is generally short Europe and long ROW, with the exception of the European currency crisis in the early 1990s. The betas are relatively stable and of reasonable magnitude.

Figure 13A shows the SR of the original carry strategy and the SRs of the UMVE-hedged version of the carry strategy. The former is 0.71 whereas the latter is 1.29. That is, there is a substantial increase in the SR. Panel B of the figure also shows the information ratio (capturing the marginal increase in SR) of the carry hedged for unpriced risks, which is higher than the SR for the original carry.

Figure 13 also shows the Sharpe and information ratios of the carry when unpriced risk is hedged out using the geographic factors as outlined above. In this case, the SR increases substantially relative to the original carry, although not by quite as much as with the optimal UMVE-based hedge. Nevertheless, it appears that these simple geographic factors go a long way towards explaining what the unpriced risks are. It is natural to draw a parallel to the unpriced industry risks in the equity market (Daniel, Mota, Rottke, and Santos, 2020) – shocks and flows to a geographic area causes currency comovements but arguably are not priced sources of risk in the currency market. Because the carry trade also loads on these shocks, the risk-return trade-off improves when these shocks are hedged out.

Finally, the figure displays the results of expanding this analysis to the remaining

eight strategies. We see that geographic hedging leads to a substantial improvements in Sharpe ratios, as measured by information ratios, across all strategies. Geographic information ratios are particularly close to the optimal ones for dollar and the three carry strategies. The cases of both CS-Mom strategies and value leave the most room for improvement, that is, for additional sources of unpriced risks. Overall, geographic factors play an important role in capturing unpriced variation in the strategy returns.

Future theoretical models should explain why geographic risks are not priced, that the standard carry trade loads on these unpriced risks and that there is a substantial increase in the SR of the carry trade once hedging out these unpriced risks. Finally, the models should explain that currency risk premiums in the G10 economies have trended down, while the carry premium and interest rate differentials the emerging markets remain high.

4 Conclusion

In this study, we explore the risk-return trade-off in the currency market. We consider common trading strategies when expanding the focus from G10 currencies to including emerging-market currencies. While the extant literature argues for improved performance when expanding the set of currencies, we find that this only enhances carry strategies, especially when accounting for transaction costs. Moreover, the benefit of including emerging economies extends only to a handful economies with floating currencies.

We construct a real-time mean-variance efficient portfolio from G10 and floating-regime emerging-market currencies. This portfolio prices trading strategies for all

types of currencies and characterizes risk premiums at each point in time. It yields risk premium dynamics consistent with both declines in average returns of G10 trading strategies over the sample and continued high carry returns of the emerging market trading strategies. Furthermore, it makes it possible for us to conditionally decompose returns into priced and unpriced components.

We show that trading strategies, including dollar and carry, contain significant amounts of unpriced risks (that increase the return variance but do not command risk premiums). By hedging out the unpriced risks, we properly characterize the risk-return trade-off in the currency market and provide new benchmarks for models of currency risk premiums. We relate the unpriced risks to currency comovements arising from geographical factors, while the priced risks in the currency market are related to long-run consumption exposures.

References

Affleck-Graves, John, and Bill McDonald, 1989, Nonnormalities and tests of asset pricing theories, *Journal of Finance* 44, 889–908.

Aloosh, Arash, and Geert Bekaert, 2022, Currency factors, *Management Science* 68, 3975–4753.

Andrews, Spencer, Ric Colacito, Mariano Max Massimiliano Croce, and Federico Gavazzoni, 2024, Concealed carry, *Journal of Financial Economics* 159, 103874.

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pedersen, 2013, Value and momentum everywhere, *Journal of Finance* 68, 929–985.

Backus, David, Silverio Foresi, and Chris Telmer, 2001, Affine term structure models and the forward premium anomaly, *Journal of Finance* 56, 279–304.

Baldi-Lanfranchi, Federico, 2024, Transaction-cost-aware factors, Working paper.

Bansal, Ravi, and Magnus Dahlquist, 2000, The forward premium puzzle: Different tales from developed and emerging economies, *Journal of International Economics* 51, 115–144.

Barillas, Fancisco, and Jay Shanken, 2017, Which alpha?, *Review of Financial Studies* 30, 1316–1338.

Barroso, Pedro, and Pedro Santa-Clara, 2015, Beyond the carry trade: Optimal currency portfolios, *Journal of Financial and Quantitative Analysis* 50, 1037–1056.

Burnside, Craig, 2011, The cross section of foreign currency risk premia and consumption growth risk: Comment, *American Economic Review* 101, 3456–3476.

Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011a, Do peso problems explain the returns to the carry trade?, *Review of Financial Studies* 24, 853–891.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, 2011b, Carry trade and momentum in currency markets, *Annual Review of Financial Economics* 3, 511–535.

Cespa, Giovanni, Antonio Gargano, Steven J. Riddiough, and Lucio Sarno, 2022, Foreign exchange volume, *Review of Financial Studies* 35, 2386–2427.

Chernov, Mikhail, Magnus Dahlquist, and Lars Lochstoer, 2023, Pricing currency risks, *Journal of Finance* 78, 693–730.

Chernov, Mikhail, Jeremy Graveline, and Irina Zviadadze, 2018, Crash risk in currency returns, *Journal of Financial and Quantative Analisys* 53, 131–170.

Chernov, Mikhail, Valentin Haddad, and Oleg Itskhoki, 2024, What do financial markets say about the exchange rate?, Working Paper.

Cochrane, John, 2005, *Asset Pricing*, revised edition (Princeton University Press, Princeton, NJ).

Daniel, Kent, Robert J. Hodrick, and Zhongjin Lu, 2017, The carry trade: Risks and drawdowns, *Critical Finance Review* 6, 211–262.

Daniel, Kent, and Tobias Moskowitz, 2016, Momentum crashes, *Journal of Financial Economics* 122, 221–247.

Daniel, Kent, Lira Mota, Simon Rottke, and Tano Santos, 2020, The cross-section of risk and return, *Review of Financial Studies* 33, 1927–1979.

De Santis, Giorgio, Robert Litterman, Adrien Vesval, and Kurt Winkelmann, 2003, Covariance matrix estimation, in R. Litterman, ed., *Modern Investment Management: An Equilibrium Approach* (Wiley, London).

DeMiguel, Victor, Alberto Martín-Utrera, Francisco J. Nogales, and Raman Uppal, 2020, A transaction-cost perspective on the multitude of firm characteristics, *Review of Financial Studies* 33, 2180–2222.

Detzel, Andrew, Robert Novy-Marx, and Mihail Velikov, 2023, Model comparison with transaction costs, *Journal of Finance* 78, 1743–1775.

Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Farhi, Emmanuel, Samuel Paul Fraiberger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan, 2009, Crash risk in currency markets, Working Paper.

Ferson, Wayne, and Andrew Siegel, 2001, The efficient use of conditioning information in portfolios, *Journal of Finance* 56, 967–982.

Filippou, Ilias, Thomas A. Maurer, Luca Pezzo, and Mark P. Taylor, 2024, Importance of transaction costs for asset allocation in foreign exchange markets, Working Paper.

Frazzini, Andrea, and Lasse Pedersen, 2014, Betting against beta, *Journal of Financial Economics* 111, 1–25.

Gibbons, Michael, Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 59, 1121–1152.

Gilmore, Stephen, and Fumio Hayashi, 2011, Emerging market currency excess returns, *American Economic Journal: Macroeconomics* 3, 85–111.

Greenaway-McGrevy, Ryan, Nelson C. Mark, Donggyu Sul, and Juh-Lin Wu, 2018, Identifying exchange rate common factors, *International Economic Review* 59, 2193–2218.

Hansen, Lars, and Scott Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Ilzetzki, Ethan, Carmen M. Reinhart, and Kenneth S. Rogoff, 2019, Exchange arrangements entering the twenty-first century: Which anchor will hold?, *Quarterly Journal of Economics* 134, 599–646.

Jagannathan, Ravi, 1996, Relation between the slopes of the conditional and unconditional mean-standard deviation frontiers of asset returns, in S. Saito, K. Sawaki, and K. Kubota, eds., *Modern Portfolio Theory and its Applications: Inquires into Asset Valuation Problems* (Center for Academic Societies, Osaka, Japan).

Jorda, Oscar, and Alan M. Taylor, 2012, The carry trade and fundamentals: Nothing to fear but FEER itself, *Journal of International Economics* 88, 74–90.

Kelly, Bryan, and Dacheng Xiu, 2023, Financial machine learning, *Foundations and Trends® in Finance* 13, 205–363.

Koijen, Ralph S.J., Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt, 2018, Carry, *Journal of Financial Economics* 127, 197–225.

Korsaye, Sofonias Alemu, Fabio Trojani, and Andrea Vedolin, 2023, The global factor structure of exchange rates, *Journal of Financial Economics* 148, 21–46.

Kozak, Serhiy, and Stefan Nagel, 2023, When do cross-sectional asset pricing factors span the stochastic discount factor?, Working Paper.

Kroencke, Tim A, Felix Schindler, and Andreas Schrimpf, 2014, International diversification benefits with foreign exchange investment styles, *Review of Finance* 18, 1847–1883.

Li, Sicong, Victor DeMiguel, and Alberto Martin-Utrera, 2023, Comparing factor models with price-impact costs, Working Paper.

Liu, Sining, Thomas Maurer, Andrea Vedolin, and Yaoyuan Zhang, 2023, Dollar and carry redux, Working Paper.

Lustig, Hanno, and Robert J. Richmond, 2020, Gravity in the exchange rate factor structure, *Review of Financial Studies* 33, 3492–3540.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

Lyons, Richard, 2001, *The Microstructure Approach to Exchange Rates*, volume 333 (Citeseer).

Ma, Guonan, Corrinne Ho, and Robert N. McCauley, 2004, The markets for non-deliverable forwards in Asian currencies, *BIS Quarterly Review, June* .

Maggiori, Matteo, Brent Neiman, and Jesse Schreger, 2020, International currencies and capital allocation, *Journal of Political Economy* 126, 2019–2066.

Maurer, Thomas, Thuy-Duong To, and Ngoc-Khanh Tran, 2022, Pricing implications of covariances and spreads in currency markets, *Review of Asset Pricing Studies* 12, 336–388.

McCauley, Robert N., and Chang Shu, 2016, Non-deliverable forwards: Impact of currency internationalisation and derivatives reform, *BIS Quarterly Review, December* .

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012a, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012b, Currency momentum strategies, *Journal of Financial Economics* 106, 660–684.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2017, Currency value, *Review of Financial Studies* 30, 416–441.

Moskowitz, Tobias J., Yao Hua Ooi, and Lasse Heje Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.

Mussa, Michael, 1986, Nominal exchange rate regimes and the behavior of real exchange rates: Evidence and implications, in *Carnegie-Rochester Conference series on public policy*, volume 25, 117–214, Elsevier.

Nucera, Federico, Lucio Sarno, and Gabriele Zinna, 2024, Currency risk premiums redux, *Review of Financial Studies* 37, 356–408.

Orlowski, Piotr, Valeri Sokolovski, and Erik Sverdrup, 2021, Benchmark currency stochastic discount factors, Working Paper.

Parker, Jonathan, and Christian Julliard, 2005, Consumption risk and the cross-section of expected returns, *Journal of Political Economy* 113, 185–222.

Richmond, Robert J., 2019, Trade network centrality and currency risk premia, *Journal of Finance* 74, 1315–1361.

Roll, Richard, 1977, A critique of the asset pricing theory's tests part I: On past and potential testability of the theory, *Journal of Financial Economics* 4, 129–176.

Schmittmann, Jochen, and Chua Han Teng, 2020, Offshore currency markets: Non-deliverable forwards (NDFs) in Asia, IMF Working Paper.

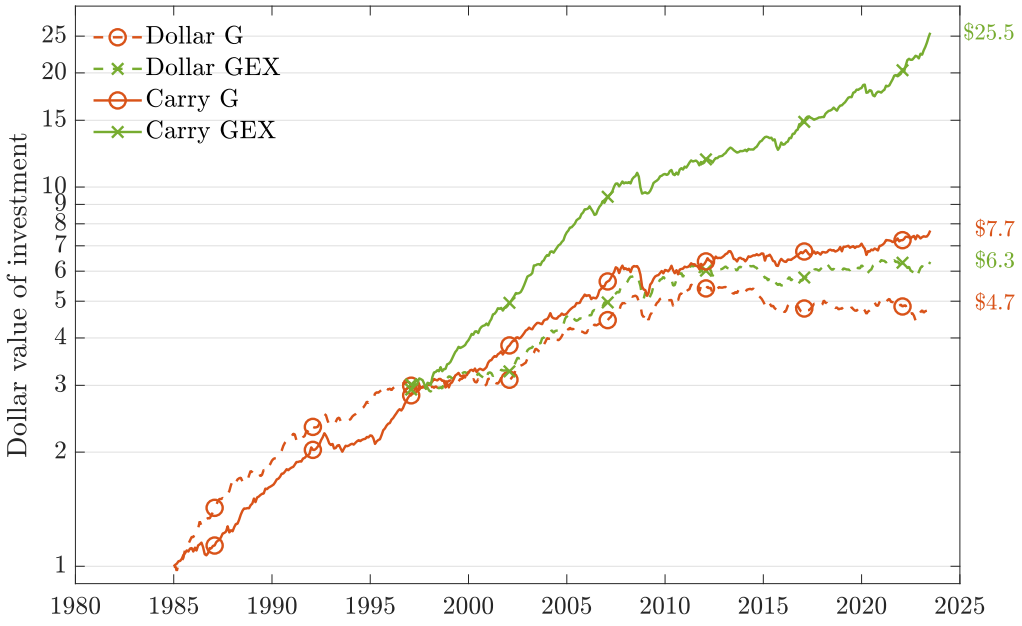
Török, Ákos, 2023, Exotic currencies and the frontier premium in foreign exchange markets, Working paper.

Verdelhan, Adrien, 2018, The share of systematic variation in bilateral exchange rates, *Journal of Finance* 73, 375–418.

Zhou, Guofu, 1993, Asset-pricing tests under alternative distributions, *Journal of Finance* 48, 1927–1942.

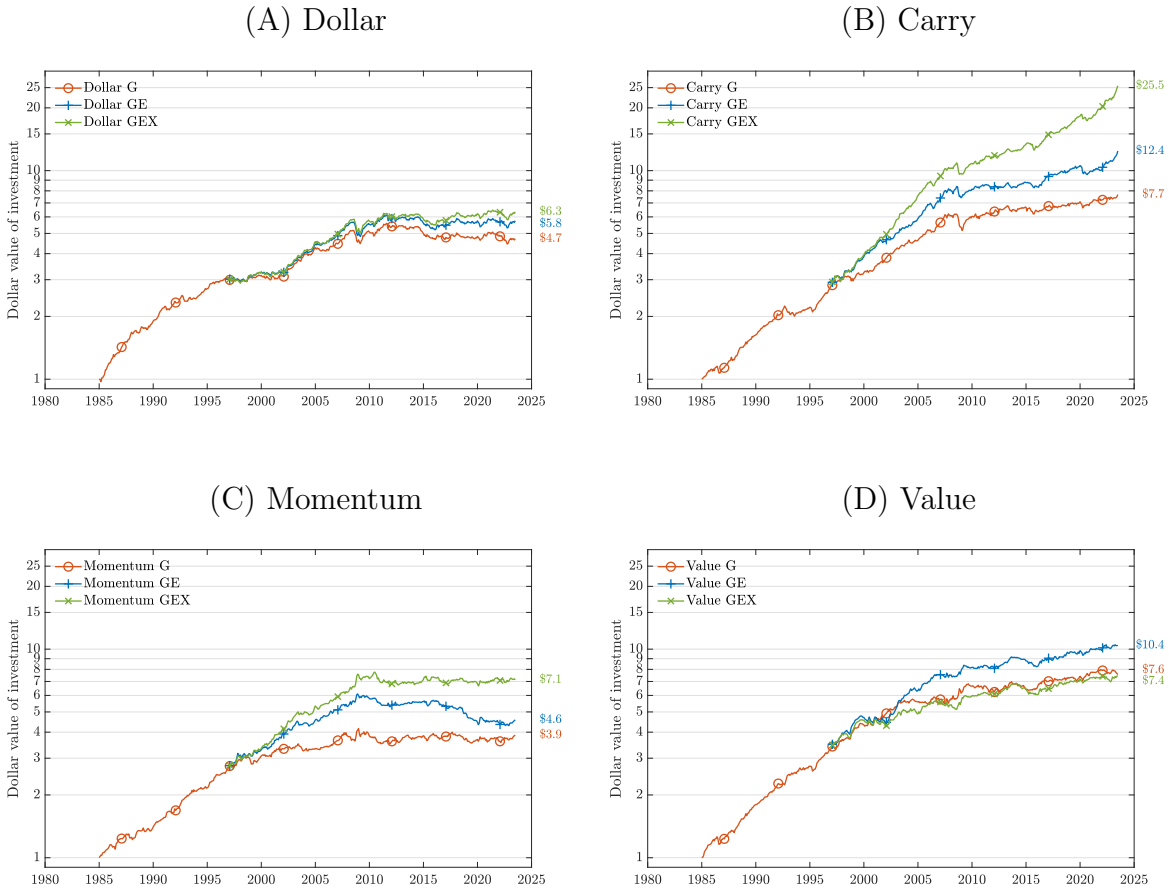
Zviadadze, Irina, 2017, Term-structure of consumption risk premia in the cross-section of currency returns, *Journal of Finance* 72, 1529–1566.

Figure 1
Investments in dollar and carry currency strategies



The figure shows dollar values of investments in the dollar and cross-sectional carry strategies (log-scale on the vertical axis). Dollar and carry investments are for both G10 currencies only (labeled G) and all currencies in our sample (labeled GEX). The cumulative gross return of an investment between t and T is given by $\Pi_{s=t+1}^{T}(1 + R_{f,s} + R_{p,s})$, where R_f is the simple return of a risk-free asset and R_p is the excess return of a currency strategy p. The excess return is ex post standardized to have an annualized standard deviation of 5% in the periods 1985–1996 and 1997–2023. The right of the figure shows the final dollar value for each of the investments, given a \$1 investment in end-December 1984. The sample is monthly from 1985 to 2023.

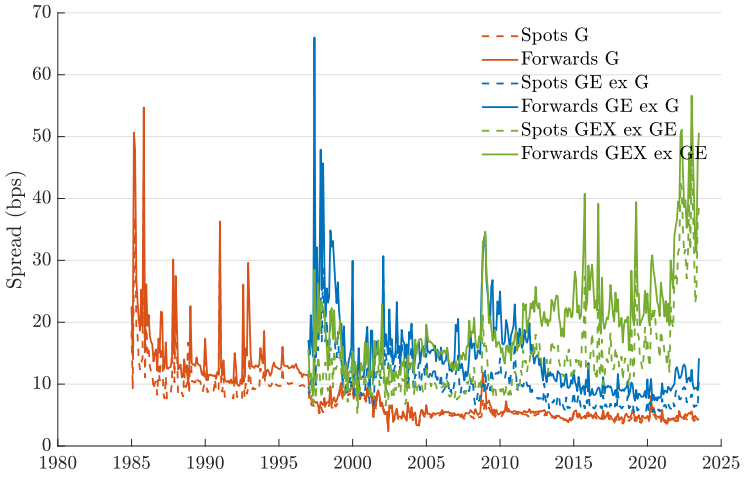
Figure 2
Investments in currency strategies



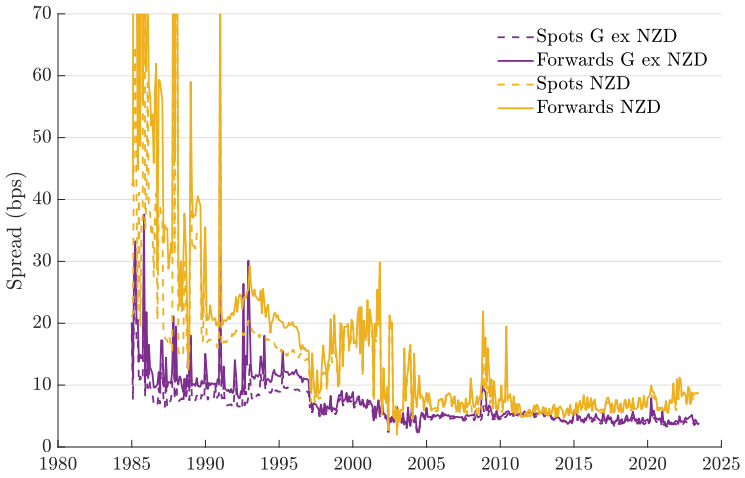
The figure shows dollar values of investments in the dollar, and cross-sectional carry, one-month momentum, and value strategies (log-scale on the vertical axis). The three lines in each panel correspond to strategies constructed from G10 currencies only (G, red line with circles), from G10 and floating-regime emerging-market currencies (GE, blue line with plus signs), and G10 and all emerging-market currencies in our sample (GEX, green line with crosses). The sample is monthly from 1985 to 2023. See also the caption of Figure 1.

Figure 3
Average bid-offer spreads of exchange rates

(A) Bid-offer spreads

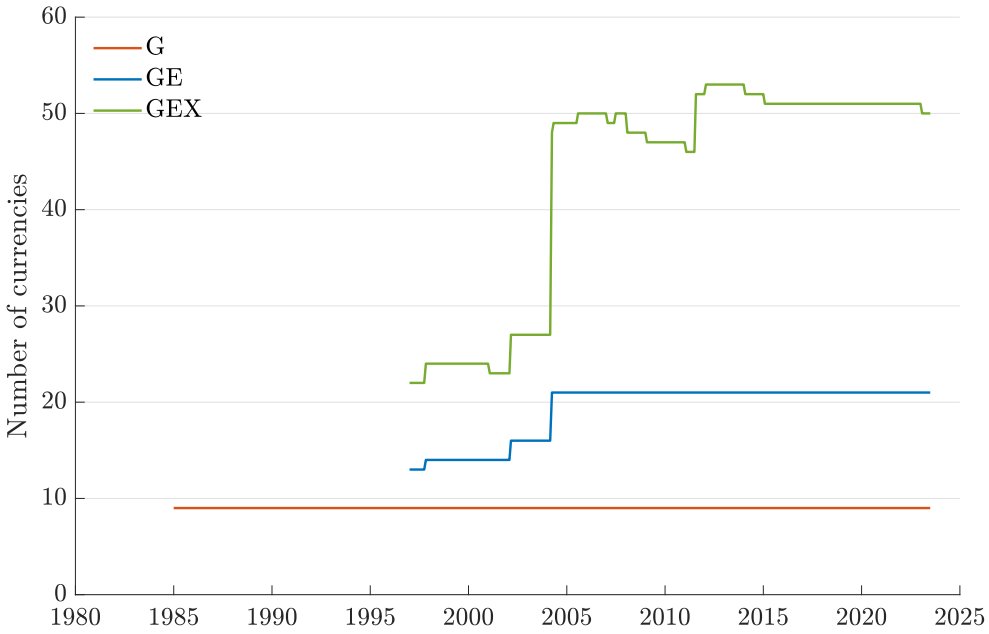


(B) Bid-offer spreads for G10 subgroups



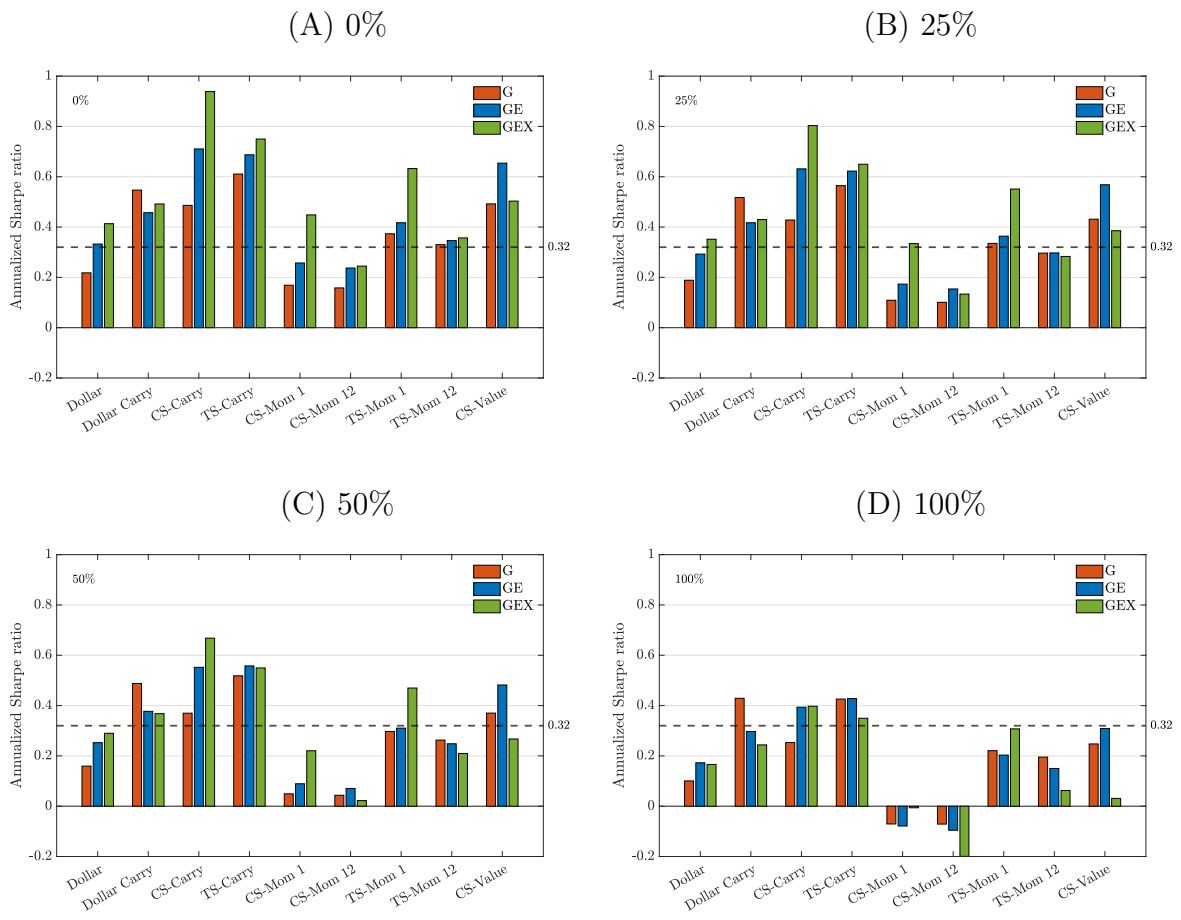
Panel A shows the average bid-offer spreads for spot (dashed lines) and forward (solid lines) markets for the three currency sets: G10 (G), floating emerging (GE ex G), and extended emerging (GEX ex GE). Panel B shows the impact of NZD costs on the G10 set by breaking it up into NZD only and G10 ex NZD. This panel is truncated at 70 bps to facilitate comparison with Panel A. The sample is monthly from 1985 to 2023.

Figure 4
Number of currencies



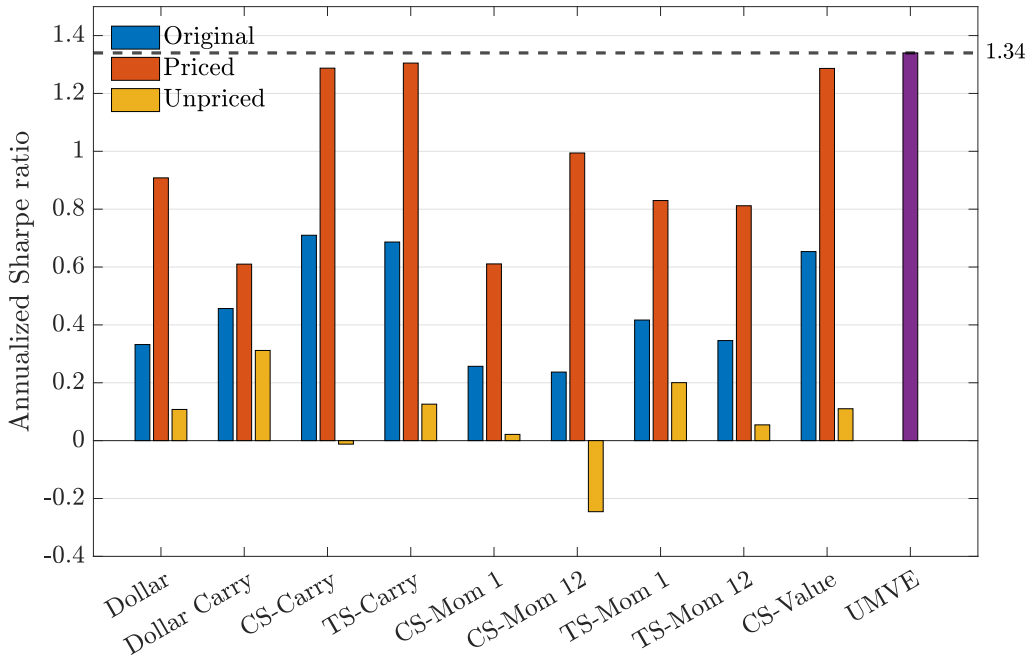
The figure shows the number of currencies within each currency set over time: G10 (G), G10 plus floating-regime emerging markets (GE), and G10 plus all emerging market currencies in our sample (GEX).

Figure 5
Sharpe ratio of strategies with different transaction costs



The figure shows annualized sample Sharpe ratios for the nine trading strategies computed for the different levels of transaction costs, given next to each panel header as a fraction of quoted bid-offers. The sample is monthly from 1985 to 2023.

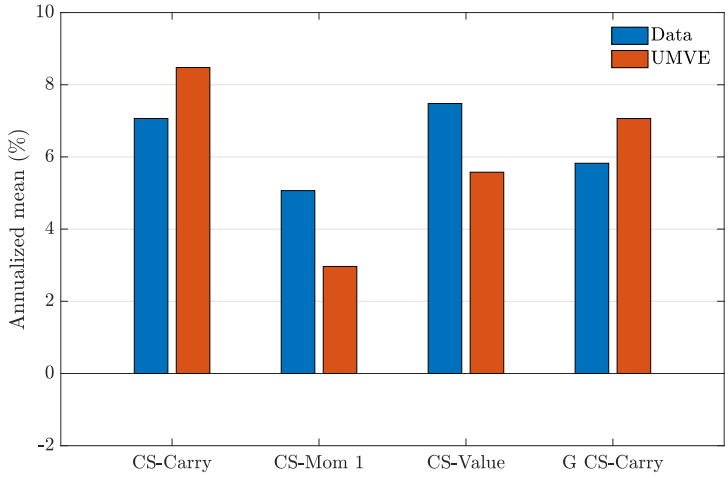
Figure 6
Sharpe ratio of priced and unpriced components of strategies



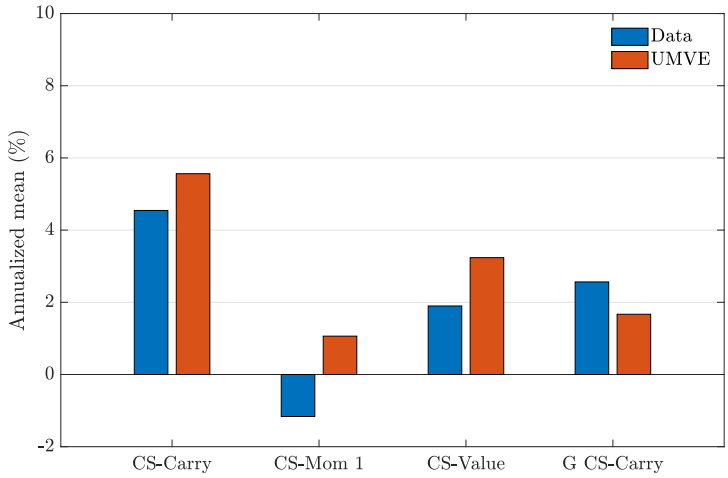
The figure shows annualized sample Sharpe ratios for three portfolios associated with each strategy. The "original" refers to the baseline version of the strategy, "priced" refers to the portfolio consisting of the original strategy return minus its unpriced component, and "unpirced" refers to the latter, which is constructed in real time. The purple bar shows the Sharpe ratio for the UMVE portfolio. The sample is monthly from 1985 to 2023.

Figure 7
Mean realized and expected returns in subsamples

(A) 1985–2004

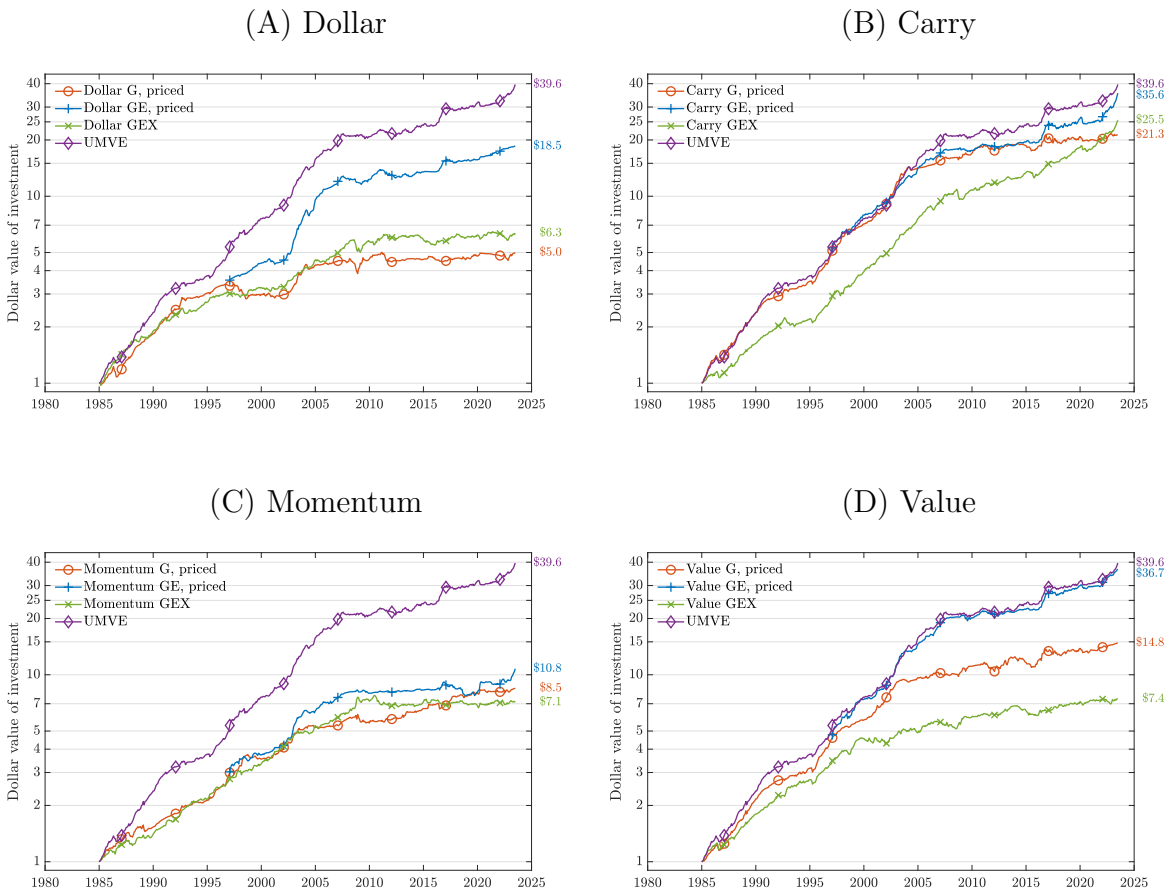


(B) 2005–2023



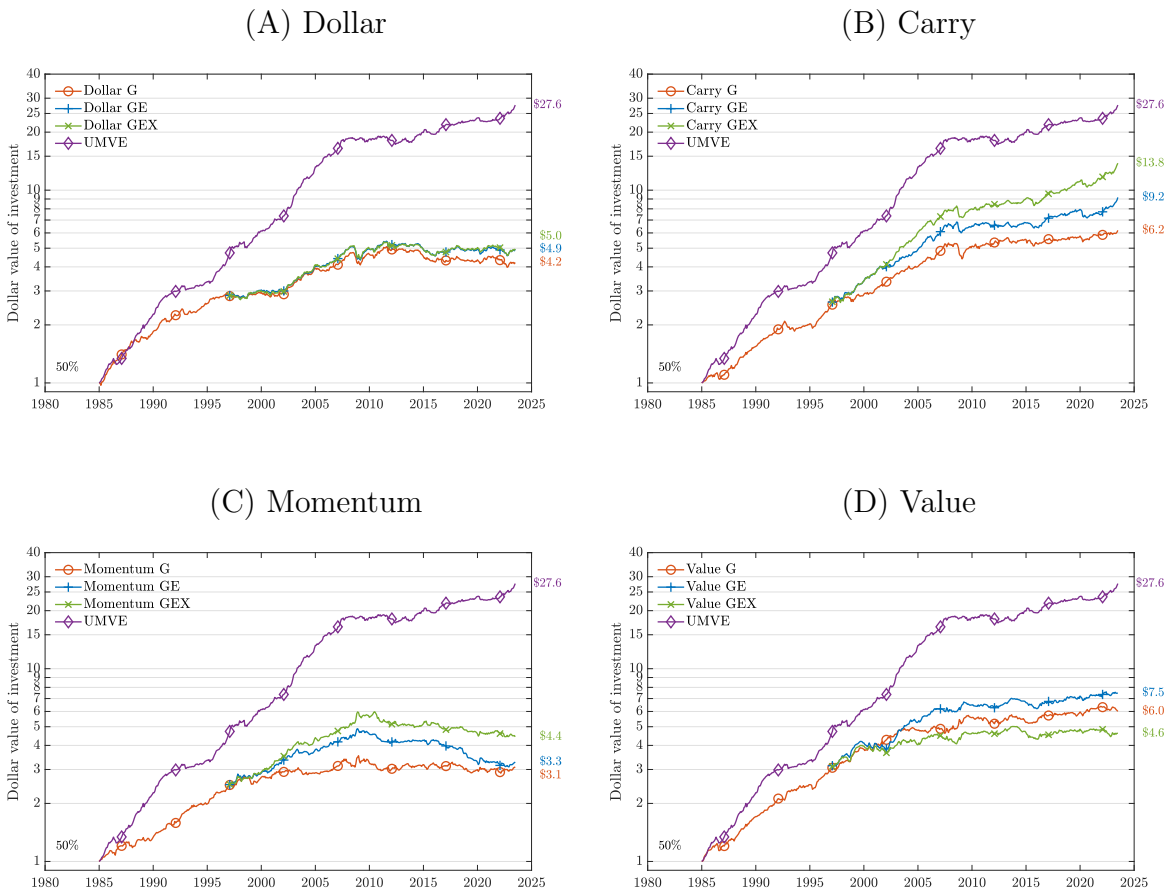
The figures shows sample average annualized returns for four strategies as given on the horizontal axis, as well as the average annualized conditional risk premium as given by our real-time UMVE construction. The upper panel shows these quantities for the first half of the sample, while the lower panel shows the results for the second half of the sample.

Figure 8
Investments in currency strategies vs the UMVE



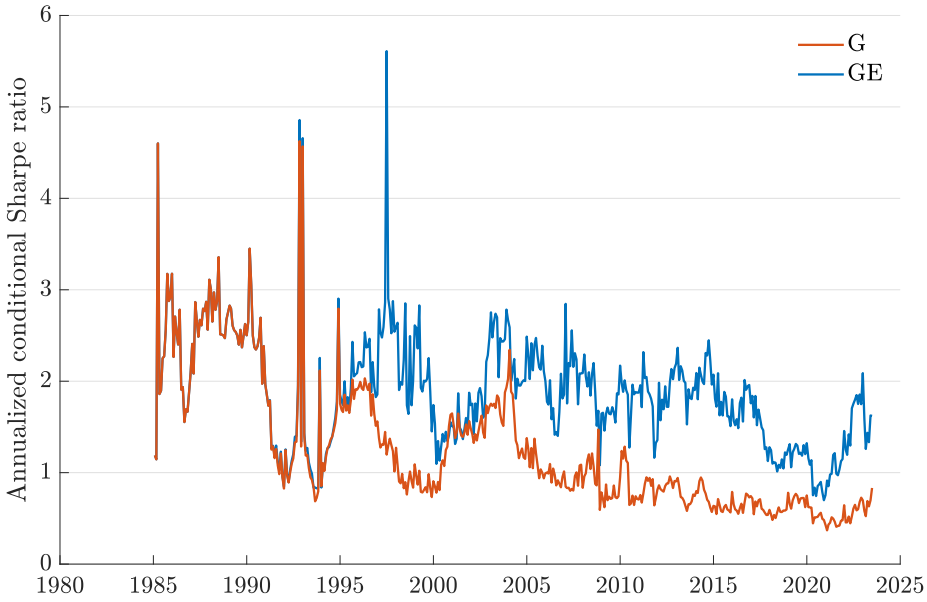
The figure shows dollar values of investments in the dollar, and cross-sectional carry, one-month momentum, value strategies for GEX currencies, their priced components using the G- and GE-based UMVEs, and the GE-based UMVE (log-scale on the vertical axis). The sample is monthly from 1985 to 2023. See also the caption of Figure 1.

Figure 9
Investments in currency strategies vs the UMVE, after costs



The figure shows dollar values of investments in the dollar, and cross-sectional carry, one-month momentum, value strategies for the three different sets of currencies, and the GE-based UMVE, after transaction costs equal to 50% of the quoted bid-offer spread (log-scale on the vertical axis). The sample is monthly from 1985 to 2023. See also the caption of Figure 1.

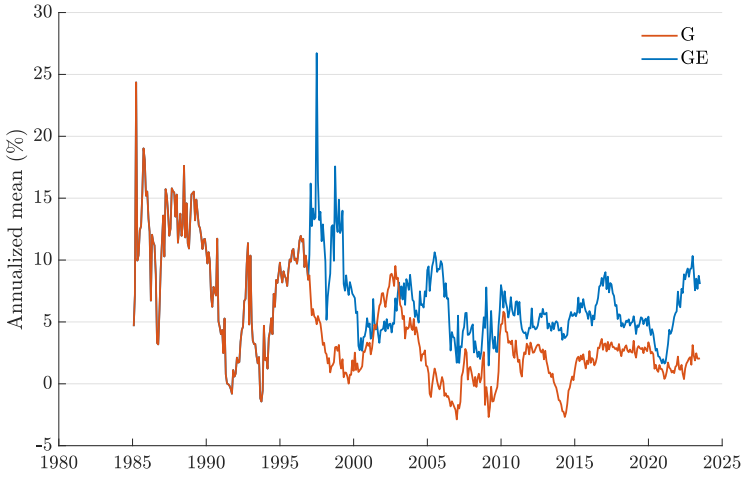
Figure 10
Conditional price of risk for G and GE UMVEs



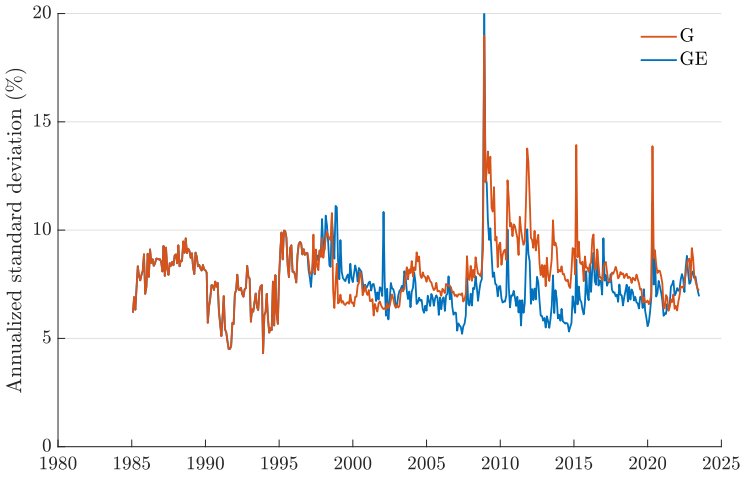
The figure shows the conditional annualized Sharpe ratio of the UMVE portfolio constructed from G10 (G) and G10 plus floating-regime emerging markets (GE) currencies. The sample is monthly from 1985 to 2023.

Figure 11
The conditional mean and standard deviation of carry

(A) Means

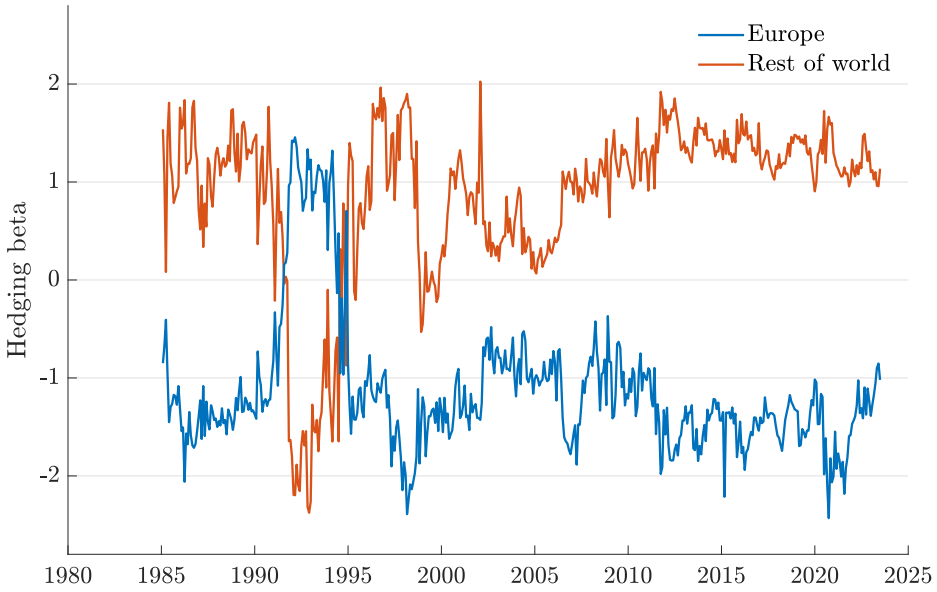


(B) Standard deviations



The figure shows the conditional mean and standard deviation of cross-sectional carry for the G and GE currencies. They are computed using our modeling of conditional means and covariance matrix of excess returns for individual currencies.

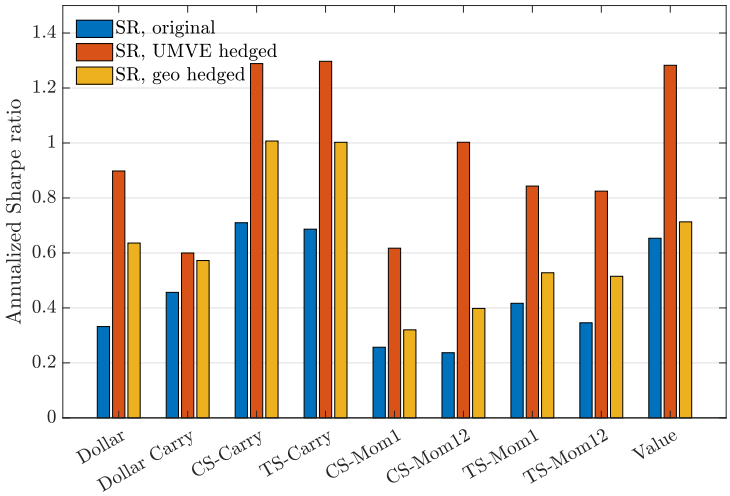
Figure 12
Geographic hedging exposures



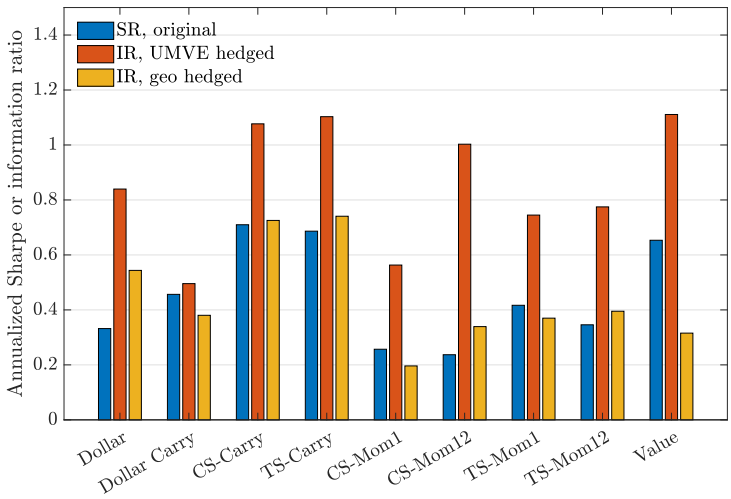
The figure shows the conditional beta on two geographic factors for the portfolio that hedges out unpriced risks from the cross-sectional carry factor. The geographic factors are European currencies and Rest of world currencies. The currency set used in the figure is the G10 plus floating-regime emerging markets (GE) set. The sample is monthly from 1985 to 2023.

Figure 13
Sharpe and information ratios for hedging strategies

(A) Sharpe ratios



(B) Information ratios



The figures gives sample annualized sample Sharpe ratios (panel A) and information ratios (panel B) for different strategies and versions of these strategies where unpriced risks, constructed optimally via UMVE or via geographic factors, are hedged out. The sample is monthly from 1985 to 2023.

Table 1: Summary statistics

		TC = 0%			TC = 25%			TC = 50%			TC = 100%		
Strategy	Stat	G	GE	GEX	G	GE	GEX	G	GE	GEX	G	GE	GEX
Dollar	mean	1.75	2.54	2.68	1.51	2.23	2.28	1.28	1.92	1.88	0.80	1.31	1.07
	std	8.04	7.64	6.49	8.03	7.63	6.49	8.02	7.63	6.48	8.01	7.62	6.48
	SR	0.22	0.33	0.41	0.19	0.29	0.35	0.16	0.25	0.30	0.10	0.17	0.17
	skew	–0.07	–0.28	–0.12	–0.07	–0.28	–0.12	–0.07	–0.28	–0.13	–0.08	–0.29	–0.14
Dollar Carry	mean	4.34	3.47	3.19	4.11	3.17	2.78	3.87	2.86	2.38	3.40	2.25	1.57
	std	7.95	7.61	6.47	7.95	7.60	6.47	7.94	7.60	6.47	7.93	7.59	6.47
	SR	0.55	0.46	0.49	0.52	0.42	0.43	0.49	0.38	0.37	0.43	0.30	0.24
	skew	–0.05	–0.07	0.01	–0.06	–0.07	0.00	–0.06	–0.07	–0.01	–0.06	–0.08	–0.02
CS-Carry	mean	4.19	5.80	6.52	3.69	5.16	5.58	3.19	4.51	4.64	2.18	3.21	2.76
	std	8.64	8.18	6.95	8.63	8.18	6.95	8.63	8.18	6.95	8.64	8.19	6.96
	SR	0.49	0.71	0.94	0.43	0.63	0.80	0.37	0.55	0.67	0.25	0.39	0.40
	skew	–0.75	–0.76	–0.98	–0.76	–0.77	–0.98	–0.76	–0.78	–0.98	–0.78	–0.80	–0.97
TS-Carry	mean	3.12	3.24	3.01	2.88	2.94	2.61	2.65	2.63	2.21	2.17	2.02	1.41
	std	5.11	4.72	4.02	5.11	4.72	4.02	5.11	4.72	4.02	5.11	4.72	4.03
	SR	0.61	0.69	0.75	0.56	0.62	0.65	0.52	0.56	0.55	0.43	0.43	0.35
	skew	–0.53	–0.79	–0.77	–0.53	–0.78	–0.77	–0.53	–0.78	–0.76	–0.52	–0.79	–0.75
CS-Mom 1	mean	1.36	1.95	3.35	0.87	1.31	2.50	0.39	0.68	1.65	–0.57	–0.60	–0.05
	std	8.04	7.60	7.49	8.04	7.60	7.50	8.03	7.60	7.51	8.03	7.60	7.53
	SR	0.17	0.26	0.45	0.11	0.17	0.33	0.05	0.09	0.22	–0.07	–0.08	–0.01
	skew	0.14	0.23	0.20	0.12	0.22	0.20	0.11	0.21	0.20	0.09	0.18	0.19
CS-Mom 12	mean	1.29	1.80	1.88	0.82	1.17	1.03	0.35	0.54	0.17	–0.58	–0.73	-1.54
	std	8.14	7.58	7.69	8.14	7.58	7.69	8.14	7.58	7.51	8.14	7.59	7.71
	SR	0.16	0.24	0.24	0.10	0.15	0.13	0.04	0.07	0.02	–0.07	–0.10	–0.20
	skew	–0.48	–0.49	–0.54	–0.49	–0.51	–0.54	–0.50	–0.52	0.20	–0.52	–0.54	–0.55
TS-Mom 1	mean	2.31	2.38	3.12	2.07	2.08	2.72	1.84	1.77	2.32	1.36	1.16	1.52
	std	6.19	5.72	4.94	6.19	5.71	4.95	6.19	5.71	4.95	6.18	5.71	4.95
	SR	0.37	0.42	0.63	0.33	0.36	0.55	0.30	0.31	0.47	0.22	0.20	0.31
	skew	0.53	0.57	0.45	0.52	0.56	0.44	0.52	0.55	0.43	0.50	0.53	0.40
TS-Mom 12	mean	2.26	2.12	1.90	2.03	1.82	1.51	1.80	1.52	1.11	1.34	0.92	0.33
	std	6.86	6.12	5.32	6.85	6.12	5.32	6.85	6.12	5.32	6.85	6.11	5.33
	SR	0.33	0.35	0.36	0.30	0.30	0.28	0.26	0.25	0.21	0.20	0.15	0.06
	skew	–0.20	–0.54	–0.50	–0.20	–0.55	–0.51	–0.21	–0.55	–0.51	–0.22	–0.57	–0.53
CS-Value	mean	3.66	4.68	3.43	3.20	4.07	2.63	2.75	3.45	1.82	1.83	2.21	0.21
	std	7.45	7.18	6.83	7.44	7.17	6.82	7.43	7.16	6.82	7.42	7.15	6.82
	SR	0.49	0.65	0.50	0.43	0.57	0.38	0.37	0.48	0.27	0.25	0.31	0.03
	skew	–0.04	–0.16	–0.24	–0.04	–0.17	–0.24	–0.05	–0.18	–0.24	–0.06	–0.19	–0.24

We report basic summary statistics for nine high Sharpe ratio trading strategies. The sample is monthly from 1985 to 2023. We group the statistics by the level of transaction costs (TC) as a fraction of bid-offer spread in the market and by the different set of included currencies: set G means G10 currencies, set GE means a combination of G10 and floating emerging currencies, set GEX is a combination of GE and remaining emerging currencies in our sample.

Table 2: Testing the GE-UMVE

Panel A GE

Strategy	SR	$E(R^e)$	t–stat	α	t–stat	β	t–stat	R^2_{adj}
Dollar	0.332	2.54	2.06	1.04	0.76	0.15	2.67	0.019
Dollar Carry	0.457	3.47	2.84	1.87	1.35	0.16	2.82	0.023
CS-Carry	0.710	5.80	4.41	–0.35	–0.26	0.60	11.38	0.314
TS-Carry	0.686	3.24	4.26	0.56	0.69	0.26	9.97	0.178
CS-Mom 1	0.257	1.95	1.60	0.95	0.57	0.10	1.27	0.008
CS-Mom 12	0.237	1.80	1.47	–1.58	–1.19	0.33	5.37	0.109
TS-Mom 1	0.417	2.38	2.59	1.38	1.28	0.10	1.98	0.015
TS-Mom 12	0.346	2.12	2.15	0.31	0.27	0.18	3.62	0.046
CS-Value	0.653	4.69	4.06	1.64	1.20	0.30	5.07	0.098

Panel B GEX

Strategy	SR	$E(R^e)$	t–stat	α	t–stat	β	t–stat	R^2_{adj}
Dollar	0.413	2.68	2.57	1.59	1.39	0.13	2.23	0.014
Dollar Carry	0.491	3.18	3.05	1.68	1.46	0.17	3.06	0.028
CS-Carry	0.938	6.52	5.82	1.56	1.31	0.57	9.82	0.281
TS-Carry	0.749	3.01	4.65	0.86	1.24	0.25	8.72	0.157
CS-Mom 1	0.448	3.35	2.78	2.66	1.70	0.08	0.92	0.003
CS-Mom 12	0.245	1.88	1.52	–1.04	–0.75	0.34	4.53	0.078
TS-Mom 1	0.632	3.12	3.92	2.38	2.59	0.09	1.67	0.011
TS-Mom 12	0.357	1.90	2.21	0.39	0.40	0.17	3.41	0.043
CS-Value	0.503	3.43	3.12	0.83	0.63	0.30	4.44	0.079

The table shows the annualized Sharpe ratio, average excess return, and t–statistic of the average excess returns to each trading strategy, along with its "alpha," "beta," and R^2 with respect to the UMVE portfolio. The latter is constructed using the GE set of currencies for both panels. The t–statistics are heteroskedasticity-adjusted. Panel A shows results for strategies constructed from GE set of currencies. Panel B considers strategies constructed from the GEX set of currencies. The sample is monthly from 1985 to 2003.

Table 3: Unconditional GRS-style tests with transaction costs

	Transaction costs			
	0%	25%	50%	100%
MSR(R^*)	1.341	1.203	1.157	0.974
GE				
MSR(R^*, R_p)	1.431	1.260	1.190	0.984
p–value	0.653	0.573	0.721	0.927
	(0.514)			
GEX				
MSR(R^*, R_p)	1.524	1.297	1.200	0.985
p–value	0.118	0.200	0.572	0.904
	(0.050)			

This table shows unconditional tests of the UMVE portfolio formed on G10 and floating-regime emerging markets (GE) countries implemented via bootstrap. "Transaction costs" refers to the fraction of the reported bid-offer spread that is used when computing returns. We consider fractions of 0%, 25%, 50% and 100%. The first row reports the sample (maximal) SR of the UMVE portfolio, $MSR(R^*)$. The middle part reports the maximal ex post sample SR obtained from combining strategies formed on the GE countries with the UMVE portfolio. The p–value reported corresponds to a bootstrap of the null hypethesis that the ex ante SR of (R^*, R_p) is the same as that of the ex ante SR of R^*, similar in spirit to the classic GRS test. The bottom part of the table corresponds to the case where the test assets are formed using the GEX sample (G10, floating-regime emerging, and the rest of emerging). The p–values reported in parentheses for zero transaction-cost case are computed using the GRS test. The sample is monthly from 1985 to 2023.

Table 4: Conditional tests of GE UMVE

Strategy	SR	$E(R^e)$	OOS α	t–stat	SR-hedged
Dollar	0.33	2.54	0.79	0.67	0.11
Dollar Carry	0.46	3.47	2.28	1.94	0.31
CS-Carry	0.71	5.80	−0.08	−0.08	−0.01
TS-Carry	0.69	3.24	0.52	0.78	0.13
CS-Mom 1	0.26	1.95	0.14	0.13	0.02
CS-Mom 12	0.24	1.80	−1.66	−1.53	−0.25
TS-Mom 1	0.42	2.38	1.07	1.24	0.20
TS-Mom 12	0.35	2.12	0.33	0.34	0.05
CS-Value	0.65	4.69	0.72	0.68	0.11
p–value		0.379	(0.379)		

This table shows tests of the conditional implications of the model where the UMVE is constructed from the GE set of currencies. See main text for details. "SR" refers to the annualized Sharpe ratio of the original strategy. "SR-hedged" refers to the out-of-sample annualized Sharpe ratio of the unpriced component of a strategy, that is, when the priced risk is hedged out conditionally in real-time — the Sharpe ratio of $R_{p,t+1} - \beta_{p,t} R^*_{t+1}$. The "OOS alpha" refers to the annualized average return to the strategy when the priced risk is hedged out conditionally in real-time ("OOS" stands for out-of-sample), and the t-statistics are adjusted for heteroskedasticity. The p–value for the joint test that OOS α's equal zero is computed by bootstrap (p–value reported in parentheses is computed using the GRS statistic). The sample is monthly from 1985 to 2023.

Table 5: Tests of dollar-carry model

			Unconditional model			Conditional model		
Strategy	SR	$E(R^e)$	α	t–stat	R^2_{adj}	OOS α	t–stat	SR-hedged
Dollar	0.33	2.54	0		1.000	0		0
Dollar Carry	0.46	3.47	2.74	2.04	0.091	0		0
CS-Carry	0.71	5.80	0		1.000	0		0
TS-Carry	0.69	3.24	0.59	1.13	0.626	0.02	0.11	0.02
CS-Mom 1	0.26	1.95	2.32	1.67	0.008	0.17	0.23	0.04
CS-Mom 12	0.24	1.80	1.39	1.10	0.016	0.75	0.95	0.15
TS-Mom 1	0.42	2.38	2.77	2.65	0.006	0.17	0.68	0.11
TS-Mom 12	0.35	2.12	2.23	2.13	0.001	0.42	1.66	0.27
CS-Value	0.65	4.69	3.48	2.70	0.053	2.49	3.05	0.49
p–value				0.001	(0.000)		0.014	(0.017)
Barillas-Shanken test:								
UMVE	1.34	10.25	7.12	6.66	0.31	5.64	6.58	1.06

This table shows test of both the unconditional and conditional dollar-carry models using the strategy returns as test assets. See main text for details, as well as the captions of tables 2 and 4. The p–value for the joint test is computed by bootstrap (p–value reported in parentheses is computed using the GRS statistic). The first (second) set corresponds to the unconditional (conditional) model. "SR-hedged" refers to the out-of-sample annualized Sharpe ratio when the model factors are hedged out conditionally in real-time — the Sharpe ratio of $R_{p,t+1} - \beta_t^{F\top} R^F_{t+1}$. The Barillas-Shanken style tests reported at the bottom of the table are testing whether the dollar-carry model prices the UMVE conditionally and unconditionally. The sample is monthly from 1985 to 2023.

Table 6: Tests of dollar-carry model on the strategies' unpriced components

	Unpriced components			Unconditional model		Conditional model	
Strategy	SR	$E(R^e)$	t–stat	α	t–stat	OOS α	t–stat
Dollar	0.11	0.79	0.67	–1.10	–3.71	–0.81	–4.16
Dollar Carry	0.31	2.28	1.94	1.97	1.49	–0.65	–3.31
CS-Carry	–0.12	–0.08	–0.08	–3.97	–6.29	–3.00	–6.36
TS-Carry	0.13	0.52	0.78	–1.23	–2.13	–1.39	–5.58
CS-Mom 1	0.02	0.14	0.13	0.89	0.78	–0.83	–1.12
CS-Mom 12	–0.25	–1.66	–1.53	–1.05	–0.93	–1.62	–1.96
TS-Mom 1	0.20	1.07	1.24	1.80	1.86	–0.39	–1.42
TS-Mom 12	0.05	0.33	0.34	1.11	1.09	–0.59	–1.60
CS-Value	0.11	0.72	0.68	0.63	0.55	0.41	0.53
p–value				0.013	(0.000)	0.000	(0.000)

This table shows test of both the conditional and unconditional dollar-carry models the *unpriced* components of the strategy returns as test assets. See main text for details, as well as the captions of tables 2 and 4. The p–value for the joint test is computed by bootstrap (p–value reported in parentheses is computed using the GRS statistic). The first (second) set corresponds to the unconditional (conditional) model. The sample is monthly from 1985 to 2023.

Table 7: Long-run consumption betas

	Original strategy			Priced component			Unpriced component		
Strategy	beta	t–stat	R^2_{adj}	beta	t–stat	R^2_{adj}	beta	t–stat	R^2_{adj}
Dollar	0.101	0.926	–0.002	0.053	1.838	0.012	0.048	0.459	–0.006
Dollar Carry	0.170	1.404	0.009	0.033	1.116	–0.000	0.136	1.233	0.004
CS-Carry	0.088	0.720	–0.004	0.244	3.366	0.075	–0.155	–1.456	0.009
TS-Carry	0.051	0.714	–0.004	0.101	3.256	0.060	–0.050	–0.877	–0.002
CS-Mom 1	0.092	0.845	–0.002	0.147	3.222	0.058	–0.056	–0.626	–0.005
CS-Mom 12	0.011	0.105	–0.007	0.143	2.480	0.041	–0.132	–1.455	0.005
TS-Mom 1	0.054	0.682	–0.004	0.075	3.030	0.045	–0.021	–0.297	–0.006
TS-Mom 12	0.084	0.892	–0.001	0.069	1.810	0.018	0.016	0.172	–0.007
CS-Value	0.228	2.224	0.022	0.127	2.662	0.038	0.102	1.022	0.000
$p-$value (equal betas)		0.069			0.005			0.641	

In this table we regress the original nine strategy returns, as well as their priced and unpriced components, on 3-year consumption growth based. The priced and unpriced components are constructed in real-time in an out-of-sample fashion – see the main text for further details. The data frequency is quarterly, from 1985 to 2023. The consumption betas are computed as $\beta_{i,c} = Cov(R^{ei}_{t,t+1}, \Delta c_{t,t+12})/Var(\Delta c_{t,t+12})$, in the spirit of Parker and Julliard (2005). In the bottom row of the table, we report p–values for testing if all betas are equal within a return set (original, priced, unpriced). Standard errors are computed via GMM adjusting for autocorrelation.

A Appendix

A.1 Details of the dataset

We consider daily spot and one-month forward exchange rates for 75 currencies versus the USD in the WM Refinitiv database (retrieved from Refinitiv Eikon) for the period December 31, 1996 to June 30, 2023. The closing bid and offer spot exchange rates are fixings around 4pm in London. Mid rates are calculated as the mean of bid and offer rates.

We complement the currency data with consumer price index (CPI) data, retrieved from the statistical databases of OECD (https://stats.oecd.org/) and IMF (https://www.imf.org/en/Data). For OECD, we use the data under "General Statistics" and then "Key Short-Term Economic Indicators"; for IMF, we use the data under "National Accounts and Price Statistics" and then "Consumer Price Index". For both OECD and IMF we retrieve "Consumer Price Index, All items." We use monthly data over the period January 1976 to June 2023, but for Australia and New Zealand we use quarterly data (with repeated monthly values) as monthly data are not available. In the case of quarterly data, the value observed at the end of a quarter is repeated monthly in the next quarter to avoid a look-ahead bias. Taiwan data are from the National Statistics website (https://eng.stat.gov.tw/cp.aspx?n=2327).

The countries (with currency ISO codes) are: Argentina (ARS), Australia (AUD), Austria (ATS), Bahrain (BHD), Belgium (BEF), Brazil (BRL), Bulgaria (BGN), Canada (CAD), Chile (CLP), China (CNY), Colombia (COP), Croatia (HRK), Cyprus (CYP), Czech Republic (CZK), Denmark (DKK), Egypt (EGP), Estonia (EEK), Eurozone (EUR), Finland (FIM), France (FRF), Germany (DEM), Ghana (GHS), Greece (GRD), Hong Kong (HKD), Hungary (HUF), Iceland (ISK), India (INR), Indonesia (IDR), Ireland (IEP), Israel (ILS), Italy (ITL), Japan (JPY), Jordan (JOD), Kazakhstan (KZT), Kenya (KES), Kuwait (KWD), Latvia (LVL), Lithuania (LTL), Malaysia (MYR), Malta (MTL), Mexico (MXN), Morocco (MAD), Netherlands (NLG), New Zeeland (NZD), Norway (NOK), Oman (OMR), Pakistan (PKR), Peru (PEN), Philippines (PHP), Poland (PLN), Portugal (PTE), Qatari (QAR), Romania (RON), Russia (RUB), Saudi Arabia (SAR), Serbia (RSD), Singapore (SGD), Slovakia (SKK), Slovenia (SIT), South Africa (ZAR), South Korea (KRW), Spain (ESP), Sri Lanka (LKR), Sweden (SEK), Switzerland (CHF), Taiwan (TWD), Thailand (THB), Tunisia (TND), Turkey (TRY), Uganda (UGX), Ukraine

(UAH), United Arab Emirates (AED), Vietnam (VND), Zambia (ZMW), United Kingdom (GBP).

The main text describes how we go from 75 to 59 currencies and divide the currencies into three currency sets (labeled G, GE, and GEX).

We complement the above currency data with daily spot and forward exchange rates for the G10 currencies (AUD, CAD, DEM & EUR, JPY, NZD, NOK, SEK, CHF, GBP) from January 1, 1976 to December 31, 1996 (used in Chernov, Dahlquist, and Lochstoer, 2023). We use WMR and Thompson Reuters exchange rates versus the GBP up to October 1983 or December 1984, and Barclays Bank International (BBI) exchange rates versus the USD from October 1983 or December 1984 (when available earliest). We use Financial Times exchange rates for the JPY versus the USD from January 31, 1976 to June 30, 1978, obtained from David Hsieh. Forward exchange rates for AUD and NZD are available from December 1984, and thus January 1985 is the common starting month for currency excess returns.

The monthly dataset keeps the last day of every month in the daily dataset.

When we cumulative returns in figures we use the risk-free rate from the Kenneth R. French Data Library
(https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

A.2 Non-deliverable forward contracts

Many observers note that trading in regular, deliverable, forward contracts (DFs) is restricted for foreign investors. Therefore, realistic strategy returns should be computed using the information from offshore non-deliverable forward contracts (NDFs). This appendix provides some background on NDFs and argues that the need for using their prices is not as clear-cut as it may seem at a first glance.

NDFs are forward currency contracts which settle in USD by payments that are equivalent to the difference between the spot rate at the maturity date and the agreed forward rate. In contrast to DFs, NDFs trade outside the direct jurisdiction of the sovereigns corresponding to the contracted currencies (i.e., offshore).

According to Ma, Ho, and McCauley (2004) and Schmittmann and Teng (2020), NDF markets have developed in response to restrictions that constrained access

to onshore markets. Sovereign policy regarding NDF markets varies a lot. Some countries, for example, Korea and Taiwan, maintain few restrictions on onshore financial institutions participation in the NDF market. In contrast, Malaysia has explicitly forced international participants into onshore DF market. Thus, policy restrictions mean more foreign trading in MYR DF rather than less. CNY has over time shifted towards mostly trading in offshore deliverable contracts resulting in a relatively modest role of NDF.

Asia accounts for the most traded NDF currencies worldwide, with the largest volumes in the KRW, INR, and TWD (Ma, Ho, and McCauley, 2004, McCauley and Shu, 2016, Schmittmann and Teng, 2020). We have NDF data for eight currencies (CNY, IDR, INR, KRW, KZT, MYR, PHP, TWD). The consensus in the literature is that the KRW offshore and onshore markets are closely integrated with TWD a close second (Ma, Ho, and McCauley, 2004, Schmittmann and Teng, 2020). Currency markets in other countries exhibit a greater degree of segmentation.

Greater degree of segmentation does not necessarily imply that the NDF market automatically represents more accurate valuation, or better reflects prices available to foreign investors. The argument that foreign investors have access to NDFs only is not necessarily correct. For example, besides the aforementioned Malaysia, McCauley and Shu (2016) show that there is trading by foreigners in DF on CNY, INR, KRW, and TWD. Indeed, restrictions on DF trading by foreigners does not mean an outright ban. Although regulations vary between countries, such restrictions entail requirements to register and to have a certain size of assets. This suggests that the largest international dealers might still be able to participate in DF markets.

Further, wider bid-offer spreads in the NDF market (see Figure A1) implies potentially poor price discovery and execution. Lastly, reliable NDF data are available for much shorter period, starting in late 2009 or early 2010, depending on a specific currency. With only 10 years of data, it is difficult to use NDF information in a reliable manner.

A.3 Trading strategies

Below we describe nine leading trading strategies, which were proposed in the literature.

The portfolio excess return of a trading strategy is:

$$R_{p,t+1} = \sum_{i=1}^{N} w_{pt}^{i} R_{t+1}^{ei}, \tag{A.1}$$

where w_{pt}^{i} is the portfolio's weight in currency i at time t and N is the number of currencies. The weight of a given portfolio p can be based on a signal, z_{pt}^{i}, and chosen such that the portfolio has an exposure to the USD or not. We use so-called rank and sign weights based on these signals. See Asness, Moskowitz, and Pedersen (2013), Koijen, Moskowitz, Pedersen, and Vrugt (2018), and Moskowitz, Ooi, and Pedersen (2012) for the use and further discussions of such weights.

We use rank weights for cross-sectional (CS) strategies. The rank of a currency is based on the signal and the weight is based on the rank according to:

$$w_{pt}^{i} = \kappa \left(\text{rank}(z_{pt}^{i}) - N^{-1} \sum_{i=1}^{N} \text{rank}(z_{pt}^{i}) \right), \tag{A.2}$$

where the scaling constant κ makes the portfolio one USD long and one USD short (and hence USD neutral). For example, for the G10 with nine currencies versus the USD, the possible weight values are $+0.4$, $+0.3$, $+0.2$, $+0.1$, 0.0, -0.1, -0.2, -0.3, and -0.4. Note that the weights depend on the currency ranks, the long and short positions sum to $+1$ and -1, respectively, and the net exposure to the USD is zero. This extends straightforwardly when we consider GE and GEX currencies.

We use sign weights for time-series (TS) strategies. The weights are then $+1$ or -1, depending on the sign of the signal, and the net exposure to the USD can be positive or negative. We further scale these sign weights with N to get a portfolio volatility similar to the ones of the cross-sectional strategies. However, this scaling does not affect the inference of a strategy's risk-adjusted performance.

Lastly, the dollar strategy differs from both CS and TS approaches as it is an equal-weighted average of the individual currency returns (Lustig, Roussanov, and Verdelhan, 2011).The dollar strategy can be seen as an equal-weighted market portfolio of currencies. It simply goes long all currencies versus the USD.

We consider three carry strategies. The dollar carry strategy uses the average forward discount (across all currencies) as a signal. Specifically, it goes long (short) all currencies versus the USD when the average forward discount is positive (negative)

(Lustig, Roussanov, and Verdelhan, 2014). Hence, the dollar carry strategy is a conditional version of the dollar strategy above: when the average forward discount is positive, it goes long the dollar strategy; when the forward discount is negative, it goes short the dollar strategy.

The CS-carry strategy uses an individual currency's forward discount as a signal and the ranking weights as described above. Currencies with relatively high forward discounts have positive weights and currencies with relatively low forward discounts have negative weights (similar to Lustig, Roussanov, and Verdelhan, 2011, who construct a high-minus-low carry portfolio rather than using the rank weights). Recall that the CS strategies are USD neutral.

The TS-carry uses the sign of the individual currency's forward discount as a signal. It goes long (short) currencies with a positive (negative) discount (Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011a, Daniel, Hodrick, and Lu, 2017). At each point in time, a varying number of currencies may have a positive or negative forward discount, so there is a time-varying exposure to the USD.

We consider two CS momentum strategies, which use the currency's performance as a signal. The CS-mom 1 strategy uses the performance in the most recent month as a signal (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012b, Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2011a) and the CS-mom 12 strategy uses the performance in the most recent year skipping the most recent month as a signal (Asness, Moskowitz, and Pedersen, 2013). Specifically, weights are rank-based as described above.

We also consider two TS momentum strategies, which use the sign of the currency's recent performance as a signal. The TS-mom 1 strategy uses the currency's last month performance (Burnside, Eichenbaum, and Rebelo, 2011b) and the TS-mom 12 strategy uses the currency's performance in the last twelve months as a signal (Moskowitz, Ooi, and Pedersen, 2012). They both go long (short) currencies with a positive (negative) performance (similar to Moskowitz, Ooi, and Pedersen, 2012).

Lastly, the CS-value strategy uses the real exchange rate signal in Equation (A.4), whereby a relatively low (high) real exchange rate today indicates that the foreign currency is cheap (expensive) (Asness, Moskowitz, and Pedersen, 2013). Specifically, weights are again based on the rank weights as described above.

This set of trading strategies comprises our main set of testing results.

A.4 Estimating conditional mean and covariance of currency returns

Our starting point for μ_t is the RWH for spot exchange rates. The RWH implies that expected excess currency returns are given by:

$$\mu_{it} \equiv E_t(R^{ei}_{t+1}) = \gamma \cdot \left(S^i_t/F^i_t - 1\right)$$

with $\gamma = 1$. This is a particular violation of UIP, which posits $\gamma = 0$. We refer to $S^i_t/F^i_t - 1$ as the (normalized) forward discount. Next, we add mean-reversion and trend signals for exchange rate forecasting. Our trend signal is a one-year depreciation rate.

Our mean-reversion signal is motivated by the literature on the role of RER in forecasting and capturing risk premiums. The RER is defined as

$$Q^i_t = S^i_t \cdot P^i_t/P_t, \tag{A.3}$$

where P_t and P^i_t are the US and foreign consumer price index (CPI), respectively. Given that the CPI is published with a lag, and we want to ensure that all variables are observable at time t, we construct the RER in Equation (A.3) using CPIs lagged by three months. The weak form of PPP implies mean-reversion in the RER. Thus, when the RER is far from its long-run mean it should forecast the currency depreciation. As Jorda and Taylor (2012) emphasize, the RER's long-run mean is not a clearly defined object empirically. We divide each RER by its five-year smoothed lag (specifically the average RER from 4.5 to 5.5 years ago) as a way to remove the dependence on the long-run mean while still preserving the long-run nature of mean-reversion signals:

$$\widetilde{Q}^i_t \equiv Q^i_t \cdot \left(\frac{1}{13}\sum_{j=-6}^{6} Q^i_{t-60+j}\right)^{-1}.$$

Lastly, we cross-sectionally demean the signal at each time t to create a cross-sectional ranking of "cheap" and "expensive" currencies. That is, our signal is

$$z^i_{Qt} \equiv \widetilde{Q}^i_t - \frac{1}{N}\sum_{i=1}^{N} \widetilde{Q}^i_t. \tag{A.4}$$

This definition has the virtue of removing any time and currency fixed effects.

In summary, we forecast excess returns OOS via:

$$\mu_{it} = \gamma_t^i \cdot \left(S_t^i/F_t^i - 1\right) + \delta_t^i \cdot z_{Qt}^i + \phi_t^i \cdot \left(S_t^i/S_{t-12}^i - 1\right). \tag{A.5}$$

We set $\gamma_t^i = 1$ to match the RWH baseline. The coefficients δ_t^i and ϕ_t^i are re-estimated every month t using historical exchange rates up until time t.

Because both the mean-reversion and trend signals rely on spot exchange rates, for which we have data going back to 1976, we have nine years of data to estimate the first conditional means and covariance matrix for January 1985 when the currency excess return sample starts. We then each month expand the sample by one month to update these estimates in an OOS fashion. This strategy gets us to the target Equation (A.5) in two steps. First, we forecast percentage changes in spot rates via:

$$S_{t+1}^i/S_t^i - 1 = \bar{\delta}_t \cdot z_{Qt}^i + \bar{\phi}_t \cdot \left(S_t^i/S_{t-12}^i - 1\right) + \varepsilon_{t+1}^i.$$

The coefficients $\bar{\delta}_t$ and $\bar{\phi}_t$ are re-estimated via a panel regression using historical data up to month t. Second, using the definition of the return on a forward position given in Equation (1), the time t expected excess currency return is obtained via:

$$\begin{aligned} \mu_{it} &= (S_t^i/F_t^i) \cdot E_t(S_{t+1}^i/S_t^i) - 1 \\ &= (S_t^i/F_t^i - 1) + (S_t^i/F_t^i)\bar{\delta}_t \cdot z_{Qt}^i + (S_t^i/F_t^i)\bar{\phi}_t \cdot \left(S_t^i/S_{t-12}^i - 1\right) \end{aligned}$$

with $(S_t^i/F_t^i)\bar{\delta}_t$ and $(S_t^i/F_t^i)\bar{\phi}_t$ corresponding to δ_t^i and ϕ_t^i in Equation (A.5), respectively.

We use daily data within each month to construct monthly realized variance for each currency depreciation rate. We compute conditional variance by running a panel AR(1) with currency fixed effects on the monthly realized variances up until time t, and forecast the realized variance for month $t+1$ based on this estimation. We proceed in an expanding manner through the sample so our conditional variance estimates are computable in real time. Next, we estimate the conditional correlation matrix using the last five years of daily depreciation rate data, where we normalize the depreciation rates by their conditional volatility. We take this time t correlation matrix and the vector of conditional variances to form the conditional covariance matrix Σ_t. Our estimation implies different decay rates on historical observations when estimating volatilities versus correlations. This follows, e.g., De Santis, Litterman, Vesval, and Winkelmann (2003), who argue that conditional correlations have much longer half-life than conditional volatilities when considering asset return data. The five-year window for correlation estimation that we apply is advocated by

Frazzini and Pedersen (2014), who also point out that correlations appear to move more slowly than volatilities.

We check if the main objects that we use for constructing the UMVE, $E_t(R^{ei}_{t+1})$ and $V_t(R^{ei}_{t+1})$, are plausible. Specifically, we check if they predict R^{ei}_{t+1} and $(R^{ei}_{t+1} - E_t(R^{ei}_{t+1}))^2$, respectively. We consider the case of GE. See Chernov, Dahlquist, and Lochstoer (2023) for G.

Panel A of Table A1 gives the results from regressing R^{ei}_{t+1} on $E_t(R^{ei}_{t+1})$ and $(R^{ei}_{t+1} - E_t(R^{ei}_{t+1}))^2$ on $V_t(R^{ei}_{t+1})$ for individual currencies in a panel setting. Panel B reports the same analysis for strategies. Under the null hypothesis that we have identified the correct conditional means and variances, the expected slope coefficients equal 1. Indeed, slopes in both regressions are found to be significantly different from 0 and insignificantly different from 1.

Lastly, Panel C explores the role of the three signals that we used to construct conditional expectations in terms of their effect on the UMVE portfolio. In the first row we report the unconditional SR for the full model, which equals 1.34. We emphasize that UMVE portfolio returns correspond to a real-time implementable trading strategy. In the next row, we consider the UMVE formed on the basis of using the forward discount signal alone (e.g., Daniel, Hodrick, and Lu, 2017, Maurer, To, and Tran, 2022). This UMVE has a lower SR of 1.06.

In order to assess the economic significance of using all the signals, we implement the Barillas and Shanken (2017) test and check if the UMVE formed using all three signals has an alpha with respect to the alternative UMVE, and vice-versa. The alpha of regressing the full-model UMVE on the forward-discount-based is significant, while the alpha from the reverse regression is not. Thus, the UMVE portfolio from the first row explains the alternative UMVE returns, but not the other way around.

A.5 Testing the UMVE constructed from G10 currencies

Tables A2 and A3 report tests of trading strategies using the G-UMVE.

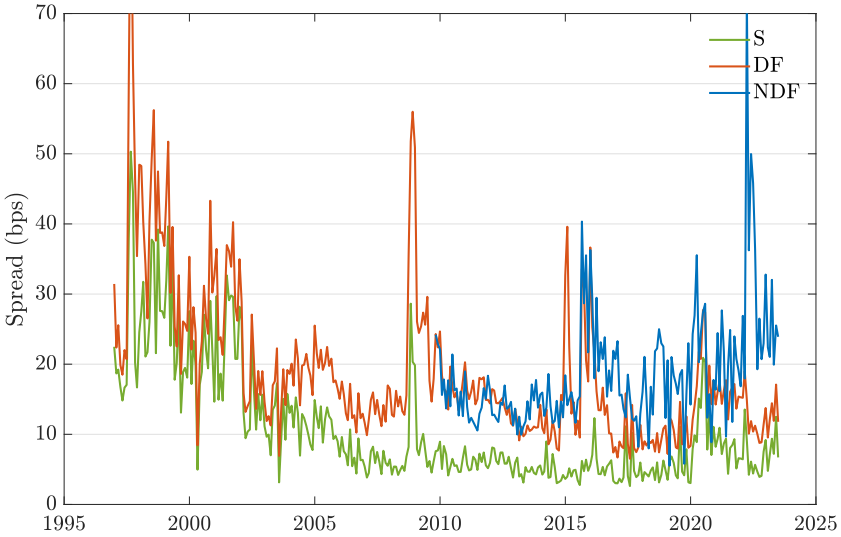
A.6 Tests with short-run consumption

We repeat our analysis reported in Table 7 in the case of short-run, or one-quarter, consumption growth as a candidate factor. Table A4 reports the results, which

are more mixed than in the case of the long-run consumption. Many betas are insignificant even when using hedged strategy returns. Most time-series adjusted R^2 are negative. That said, we reject the null of identical betas across all three sets of returns (original, priced and unpriced components).

As is the case with many other studies, the Covid period could have affected the inference due to wild swings in consumption growth. (This is not an issue for long-run consumption as the 12-quarter period smooths out the Covid impact.) In unreported results we run the same regressions excluding the post 2020 sample. Many results are similar to those reported for the full sample: a lot of insignificant betas and negative adjusted R^2. The big difference is that now the null of identical betas is not rejected across all three sets of returns (original, priced and unpriced components). Thus, there is no clear support for short-run consumption risk being the driver of currency risk premiums.

Figure A1
Average bid-offer spreads for currencies with NDF contracts



The figure shows the average bid-offer spreads for spot (S), deliverable forward (DF), and non-deliverable forward (NDF) markets for eight currencies with NDF data (CNY, IDR, INR, KRW, KZT, MYR, PHP, and TWD). The figure is truncated at 70 bps to facilitate comparison with Figure 3. The sample is monthly from 1997 to 2023.

Table A1: Predictive ability of conditional expectations and variance

Panel A Currency returns

R^{ei}_{t+1}		$(R^{ei}_{t+1} - E_t(R^{ei}_{t+1}))^2$	
$E_t(R^{ei}_{t+1})$	0.666	$V_t(R^{ei}_{t+1})$	0.904
s.e.	0.176	s.e.	0.102
R^2_{adj}	0.008	R^2_{adj}	0.246

Panel B Strategy returns

R^{ei}_{t+1}		$(R^{ei}_{t+1} - E_t(R^{ei}_{t+1}))^2$	
$E_t(R^{ei}_{t+1})$	0.713	$V_t(R^{ei}_{t+1})$	1.143
s.e.	0.178	s.e.	0.115
R^2_{adj}	0.019	R^2_{adj}	0.235

Panel C UMVE portfolios

Signals used	SR	alpha	t–stat	R^2_{adj}
All signals	1.3434	0.50	0.27	0.579
Fwd disc	1.0578	7.16	3.01	0.579

This analysis is implemented on the basis of the GE set. Panel A shows panel regressions of currency excess returns R^e_{t+1} on the conditional expected excess returns as estimated from our advocated model, $E_t(R^e_{t+1})$. We also regress $(R^e_{t+1} - E_t(R^e_{t+1}))^2$ (using the advocated estimate of the conditional expectation) on the conditional variance as estimated by our model, $V_t(R^e_{t+1})$. Panel B shows the same for the strategy returns. Standard errors are clustered by month. Panel C compares the performance of the UMVE from the full model (all signals used) vs the model where conditional mean is estimated using the forward discount only (fwd disc). The panel reports the Sharpe ratios (SR) of the two UMVE portfolios, along with the alpha, t–statistic and R^2 from the Barillas and Shanken (2017) regression of a given UMVE on its competitor. The numbers in Panel C are annualized (except for t–statistics and R^2_{adj}s), and the UMVE portfolios are normalized to have the same volatility as that of the dollar strategy.

Table A2: Testing strategy returns in G only using G-UMVE

Strategy	SR	$E(R^e)$	t–stat	α	t–stat	β	t–stat	R^2_{adj}
Dollar	0.218	1.75	1.35	1.02	0.75	0.09	1.39	0.006
Dollar Carry	0.546	4.34	3.39	2.15	1.63	0.27	4.84	0.071
CS-Carry	0.486	4.19	3.02	0.51	0.36	0.45	6.15	0.172
TS-Carry	0.610	3.12	3.79	1.02	1.23	0.26	9.91	0.160
CS-Mom 1	0.169	1.36	1.05	0.63	0.40	0.09	1.05	0.006
CS-Mom 12	0.158	1.29	0.98	–1.09	–0.78	0.29	4.20	0.080
TS-Mom 1	0.373	2.31	2.32	1.06	1.00	0.15	2.62	0.037
TS-Mom 12	0.330	2.26	2.05	0.26	0.23	0.24	4.36	0.080
CS-Value	0.492	3.66	3.05	0.97	0.72	0.33	5.22	0.123
UMVE SR	1.023	GRS	p–value	0.714				

The table shows the annualized Sharpe ratio, average excess return, and t–statistic of the average excess returns to each trading, along with its "alpha," "beta," and R^2 with respect to the UMVE portfolio, which is constructed using the G set of currencies. The t–statistics are heteroskedasticity-adjusted. The p–value is computed using the GRS test. The strategy returns are constructed using the G set. The sample is monthly from 1985 to 2023.

Table A3: Testing strategy returns in GE only using G-UMVE

Strategy	SR	$E(R^e)$	t–stat	α	t–stat	β	t–stat	R^2_{adj}
Dollar	0.332	2.54	2.06	1.95	1.50	0.08	1.15	0.004
Dollar Carry	0.457	3.47	2.84	1.46	1.14	0.26	4.46	0.065
CS-Carry	0.710	5.80	4.41	2.42	1.79	0.43	6.28	0.162
TS-Carry	0.686	3.24	4.26	1.44	1.85	0.23	7.48	0.138
CS-Mom 1	0.257	1.95	1.60	1.24	0.82	0.09	1.09	0.006
CS-Mom 12	0.237	1.80	1.47	–0.32	–0.25	0.27	3.99	0.073
TS-Mom 1	0.417	2.38	2.59	1.29	1.27	0.14	2.42	0.033
TS-Mom 12	0.346	2.12	2.15	0.51	0.49	0.21	4.06	0.064
CS-Value	0.653	4.69	4.06	2.62	1.97	0.27	4.15	0.078
UMVE SR	1.023	GRS	p–value	0.083	(0.041)			

The table shows the annualized Sharpe ratio, average excess return, and t–statistic of the average excess returns to each trading, along with its "alpha," "beta," and R^2 with respect to the UMVE portfolio, which is constructed using the G set of currencies. The t–statistics are heteroskedasticity-adjusted. The p–value is computed using bootstrap (and the GRS test in parenthesis). The strategy returns are constructed using the GE set. The sample is monthly from 1985 to 2023.

Table A4: Consumption betas for the short-run consumption growth factor

	Original strategy			Priced component			Unpriced component		
Strategy	beta	t–stat	R^2_{adj}	beta	t–stat	R^2_{adj}	beta	t–stat	R^2_{adj}
Dollar	–0.000	–0.002	–0.007	0.042	1.176	–0.005	–0.042	–0.179	–0.006
Dollar Carry	0.058	0.189	–0.006	0.044	0.973	–0.005	0.014	0.049	–0.007
CS-Carry	0.160	0.768	–0.005	0.287	2.250	0.009	–0.127	–0.976	–0.005
TS-Carry	–0.010	–0.068	–0.007	0.112	1.886	0.005	–0.123	–0.965	–0.002
CS-Mom 1	–0.106	–0.410	–0.006	–0.033	–0.572	–0.006	–0.073	–0.302	–0.006
CS-Mom 12	0.064	0.271	–0.006	0.083	1.092	–0.004	–0.019	–0.086	–0.007
TS-Mom 1	0.007	0.064	–0.007	0.064	2.026	–0.001	–0.057	–0.541	–0.006
TS-Mom 12	0.140	0.543	–0.004	0.100	1.528	0.002	0.040	0.164	–0.006
CS-Value	0.112	0.669	–0.006	0.155	2.534	0.003	–0.042	–0.244	–0.006
p–value (equal betas)		0.004			0.000			0.002	

In this table we regress the original nine strategy returns, as well as their priced and unpriced components, on contemporaneous quarterly consumption growth. The priced and unpriced components are constructed in real-time in an out-of-sample fashion – see the main text for further details. The data frequency is quarterly, from 1985 to 2023. The consumption betas are computed as $\beta_{i,c} = Cov(R^{ei}_{t,t+1}, \Delta c_{t,t+1})/Var(\Delta c_{t,t+1})$. In the bottom row of the table, we report p–values for testing if all betas are equal within a return set (original, priced, unpriced). Standard errors are computed via GMM adjusting for autocorrelation.