An anatomy of currency strategies: The role of emerging markets*

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

First draft: September, 2023

This draft: April 21, 2024

Abstract

How do emerging economies contribute to the currency market's conditional risk-return trade-off? We construct an out-of-sample mean-variance efficient portfolio from the G10 and floating-regime emerging-market currencies that prices trading strategies based on all available currencies and characterizes risk premiums at each date in the sample. The risk premium dynamics are consistent with both the substantial decline in average returns to many of the G10 trading strategies over the sample, and the continued high carry returns of the emerging trading strategies. Further, the approach provides a conditional return decomposition into priced and unpriced components. We show that trading strategies, including dollar and carry, are strongly exposed to currency factors that increase return variance, but that do not command a risk premium. This unpriced risk must be hedged out from these strategies if they are to properly characterize risk and return in the currency market. For instance, the carry strategy has a Sharpe ratio of 0.71 that increases to 1.29 after real-time hedging of unpriced risks. We relate these unpriced risks to currency comovements arising from geographically-based factors.

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

---

*We have benefited from the comments of seminar particpants at NYU. We thank Haohang Wu for excellent research assistance. Dahlquist gratefully acknowledges support from the Jan Wallander and Tom Hedelius Foundation.

[†]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 impact of emerging market currencies on the risk-return trade-off implicit in famous 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 cross-sectional carry and momentum strategies for G10 only (labeled G) and all available currencies (labeled GEX). We see that carry based on all currencies has completely decoupled from the G10-based carry during the first decade of the 21st century. The latter has flattened out during the second decade, while the former continued to grow. Momentum is exhibiting less dramatic but similar patterns with the G10-based momentum leveling off in the early 2000s, whereas the full-sample returns continue to grow (though leveling off in the later part of the sample).

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. The second one is based on the fact that risk premiums for certain trading strategies—despite being well-diversified portfolios—seem to approach zero, even though the volatility of the returns from these strategies remains consistent throughout the sample period. Thus, we surmise that the strategies are exposed 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. Another important dimension of this evidence is that it appears especially relevant when considering G10 currencies versus emerging markets.[1]

Our first contribution is to show that the emerging market outperformance arises from the subset of emerging market economies that are under a floating exchange rate regime, or, put differently, 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 construct real-time measures of the conditional expected excess returns to G10 and emerging market currencies, as well as their conditional covariance matrix. Using these measures we build on Chernov, Dahlquist, and Lochstoer (2023) and construct an estimate of the unconditional mean-variance efficient (UMVE) portfolio (Hansen and Richard, 1987).[2] Excess returns on this portfolio represent a real-time, tradeable, single factor. We show empirically that

[1]We refer to non-G10 countries as emerging, although that is not correct, strictly speaking – some of the non-G10 countries with floating exchange rate would be classified as developed. Likewise, 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.

[2]The UMVE portfolio prices all dynamic admissible trading strategies. As a result, it correctly prices the full cross-section of currencies and trading strategies associated with them conditionally as well.

the estimated UMVE factor prices all currency strategies conditionally and unconditionally, in sample and out of sample (OOS). This approach allows us to characterize the risk that drives currency risk premiums at each date throughout the sample, as well as the (conditional) factors that are important for return variance but not risk premiums. That is, this framework allows us the address the two conjectures set out above.

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 about 0.33 to 0.91, the 12-month cross-sectional momentum strategy from about 0.24 to 0.99, and the cross-sectional value goes from about 0.65 to 1.29. These are large increases, which strongly 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.

We also 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 strategies with high SR diminishing the power of tests. Although the model is rejected despite these concerns, we exploit the importance of unpriced risks and implement additional tests of the model on the basis of returns on new strategies. The new strategies are simply the traditional strategies with unpriced risks hedged out using the estimated UMVE. There is no longer a mechanical connection between such returns and candidate factors. Both unconditional and conditional models are rejected. Because dollar and carry are close to the first two principal components of the variation in currency returns, these results indicate that risk-based explanations of currency risk premiums cannot be successful.

Overall, our results lead to the natural question of what these 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 to explain the unpriced risks. Such comovements can be driven by common shocks to the economies of close countries (see, e.g., Lustig and Richmond, 2020). We further verify that a factor model with the dollar and the carry factors cannot explain the average return to the carry strategy after unpriced risks are hedged out using the geographical factors (as opposed to the optimal UMVE-based hedges).

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 are so large as measured by bid-offer spreads 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, or limited set of strategies, or limited set of currencies, we consider a range of possible transaction costs as a fraction of indicative 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 also undertake 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, and 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 when transaction costs are ignored, the costs-based analysis is a robustness check. The main concern is that the UMVE is impacted by transaction costs more than the strategies themselves. We find that it is progressively easier to explain strategies based on the currencies from the largest set as transaction costs increase from 0% to 100% of the quoted bid-offer spread.

Literature. Bansal and Dahlquist (2000) is an early paper 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 in the curency strategy returns, and the level and impact of transaction costs.

Andrews, Colacito, Croce, and Gavazzoni (2023) observe a decline in carry returns for G10 currencies during the post-2008 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 goal although the considered constraints are not as general. 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.

Filippou, Maurer, Pezzo, and Taylor (2024) analyze transaction costs to investor in the foreign exchange markets. They consider a subset of our strategies net of transaction costs using data similar to our GE set, and compare to a cost-optimized strategies with an emphasis on price impact. These authors argue that the constructed cost-optimized conditional mean-variance efficient (CMVE) portfolio experiences a small impact of large quoted costs. The authors do not consider testing cross-sectional pricing. Absent price impact and incorporating full transaction costs because of monthly portfolio turnover, their optimization problem coincides with our SR maximization with returns net of transaction costs. Our finding is that both versions of the UMVE portfolio, cost-optimized and 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. The world today 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 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.

In practice, many currencies are tightly managed by their respective governments. Ilzetzki, Reinhart, and Rogoff (2019) (IRR) introduce currency classification to reflect this. We use their work to separate out a smaller set of currencies, which are loosely classified as free floating, to which we refer as floating emerging currencies. These are currencies with typical IRR scores of 11, 12, or 13, indicating that they have had moving bands (allowing for both appreciation and depreciation over time), managed floats, or free floats. Typically, currencies with these scores remain above 10. Thus, we do not change our classification throughout the sample.

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 includes TWD, which does not have an IRR score.

In the course of analysis we consider three currency datasets:

1. G10 (AUD, CAD, EUR spliced with DEM, JPY, NZD, NOK, SEK, CHF, GBP). We backfill this dataset to 1985 using the data considered in Chernov, Dahlquist, and Lochstoer (2023). We refer to this dataset as G.

2. G10 combined with 12 floating energing currencies (BRL, CLP, COP, ISK, INR, ILS, MXN, PLN, SGD, ZAR, KRW, and THB). We refer to this dataset as GE.

3. G10 combined with 12 floating emerging currencies and 38 extended emerging currencies. We refer to this dataset as GEX.

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 (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 their National Statistics.

Let the USD be the measurement (numeraire) currency, that is, all exchange rates are expressed 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 $S_{t+1}^i - F_t^i$. One common way to scale this payoff to define excess return is to divide by F_t^i :

$$R_{t+1}^{ei} = (S_{t+1}^i - F_t^i)/F_t^i. \tag{1}$$

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 CIP holds or not. Appendix A.2 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.

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 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 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. Figure 2 shows that there is substantial heterogeneity in the performance of the trading strategies across time and currency sets. For instance, the dollar strategy appears to have near zero return from the financial crisis and on for all currency sets. 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.

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 increase in trading costs over time. Starting with the average cost in the range of 10–20 bps, these currencies end the sample at a much higher level of about 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 all SR 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 very 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 currency 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 $R_{p,t+1} = w_{pt}^{\top} R^e_{t+1}$.

The UMVE portfolio is a dynamic trading strategy in these assets that obtains the MSR, both conditionally and unconditionally. Ferson and Siegel (2001) and Jagannathan (1996) show that the UMVE portfolio weights are:

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

We largely follow Chernov, Dahlquist, and Lochstoer (2023) to estimate μ_t and Σ_t. We modify the procedure for the covariance matrix to ensure that it applies to a larger and unbalanced panel of currencies. See Appendix A.3 for further details. Denote the excess return on this portfolio $R^*_{t+1} = w_t^{*\top} R^e_{t+1}$.

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.

A CMVE portfolio is a dynamic trading strategy 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. All such portfolios would satisfy the conditional linear pricing relation (3). The UMVE portfolio is the only CMVE portfolio that maximizes the unconditional SR and 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 unconditional MSR of the UMVE constructed from the base assets and test assets is not significantly different from that constructed from base assets alone:

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

where Σ_ε is the covariance matrix of the regression residuals.

It is not clear a-priori that our estimates of μ_t and Σ_t, which are constructed OOS are correct. That is, there is no guarantee that the resulting UMVE portfolio would have the ex ante MSR. That is a source of 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 large set of trading strategies p proposed in the literature.

Affleck-Graves and McDonald (1989) and Zhou (1993) point out that the GRS test tends to over reject 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, more conservative, asymptotic 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:

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

which we solve numerically at each time t.

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

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

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. 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 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.[3]

[3] 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 models, and thus

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:

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

where F_{1t} and F_{2t} are iid true factors with unit variance. Further, assume that F_{1t} is "priced" with $E_t\left(F_{1,t+1}\right) = 1$, and F_{2t} is "unpriced" $E_t\left(F_{2,t+1}\right) = 0$.

Suppose a researcher picks a characteristic with values equal to the loading of a return on F_{1t}. Thus, the characteristic-based factor equals

$$R_t^i - R_t^j = F_{1t} + \left(\beta_i - \beta_j\right) F_{2t}.$$

The long-short combination of returns 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

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.

combination without increasing the expected return of the long-short combination. The SR of the long-short combination will then be lower than for a portfolio with no exposure to the unpriced risk.

Now, choose MVE portfolio weights instead. The mean and covariance 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,

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

Using these optimal weights with currency excess returns, we have:

$$2\frac{\beta_j}{\beta_i + \beta_j}R_t^i - 2\frac{\beta_i}{\beta_i + \beta_j}R_t^j = F_{1t}.$$

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. We also consider a standard factor model used in the literature as an alternative to the UMVE 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 such UMVE 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 and with slightly different estimate of the conditional covariance matrix. Appendix A.4 demonstrates that this is indeed the case. Second, we demonstrate in the same Appendix that this UMVE cannot price strategies constructed from GE currencies. Thus, we no longer consider this UMVE in the sequel.

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 adjusted R^2 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 in this case is 0.653 (0.514 for GRS), 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 the maximal ex post SR 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 test results could be stronger 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. This result implies that there little benefit to trading currencies from the extended emerging set.

Conditional model implications

Since the UMVE factor should price assets both conditionally and unconditionally, we can easily implement an unconditional 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 out-of-sample test of the UMVE portfolio, we follow Chernov, Dahlquist, and Lochstoer (2023) and exploit the fact that conditional linear beta pricing model (3) holds. Therefore, if we remove the priced component of a portfolio return, the residual,

$$R_{h,t+1} = R_{p,t+1} - \beta_{pt} R^*_{t+1}, \tag{7}$$

referred to as hedging portfolio should have zero alpha. The conditional beta is equal to

$$\beta_{pt} = \frac{w_{pt}^\top \Sigma_t w_t^*}{w_t^{*\top} \Sigma_t w_t^*}. \tag{8}$$

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

Table 4 reports the test results applied to $R_{h,t+1}$ 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 hedging return $R_{h,t+1}$ 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 hedging portfolio and the portfolio where unpriced risks are hedged out. The latter simply has returns $R_{p,t+1} - R_{h,t+1}$. Note that the construction of the hedging portfolios is done in real time and thus the hedged returns are indeed tradeable.

The hedging portfolios all have SRs close to zero, while the portfolios with unpriced risks hedged out 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 about 0.33 to 0.91, the 12-month cross-sectional momentum strategy from about 0.24 to 0.99, and the cross-sectional value goes from about 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-pricing due to their contamination from these substantial unpriced risks.

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 popular 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.

In terms of the conditional risk-return trade-off, Figure 8 shows 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. Thus, the floating-regime emerging markets currencies indeed materially affect the conditional risk-return trade-off over the sample.

Testing 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). Here we test the unconditional and conditional implications of this 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 everywhere else 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 (7). Under the null of the conditional two-factor dollar-carry model, the excess returns on the hedging portfolio

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

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 (8).

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 two-factor structure pf 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 the conditional model cannot be implemented without computing the conditional covariance matrix of currency returns, which is a key ingredient to the UMVE portfolio. Second, these tests could have weak power because they do not include the new information from our UMVE construction and the importance of hedging out unpriced risks from the strategy under consideration. Thus, we turn to the implications arising from these findings.

The most direct way to account for the large amount of unpriced risks is to hedge them out as per Equation (7). Next, we ask whether these new strategy returns can be priced by dollar and carry. We test the unconditional and conditional versions of the model again. Table 6 reports the results.

Now there is no mechanical relationship between hedged dollar and carry returns with straight dollar and carry returns. either unconditionally or conditionally. The model struggles with capturing returns on the hedged strategy both unconditionally and conditionally (OOS). All individual alphas are significant and the joint p–value is zero.

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, but it cannot account for the these strategies when unpriced risks are hedged out. A corollary to this is that the two first principal components, which are closely related to the standard dollar and carry factors, are not in fact good factors for determining currency risk premiums.

3.4 Implications of the unpriced risks

In this subsection we delve deeper into the implications of the large unpriced risks that appear in common currency trading strategies. First, we consider, as a robustness check, an alternative method for real-time identification of unpriced risks. Second, we propose a more intuitive description of what these risks actually represents by relating these risks to geographically-based currency factors. Finally, we discuss the implications for models of currency market risk premiums. 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.

An alternative method of real-time identification of unpriced risks

An alternative, but related, way to identify unpriced risks in the carry portfolio is to first assume that conditional expected returns are linear in the interest rate differential of country i versus the dollar. That is, $\mu_{it} = S_t^i/F_t^i - 1$. The conditional covariance matrix, Σ_t, is estimated as before. The standard carry factor does not use Σ_t to determine portfolio weights. Thus, the original way to create the factor is not optimal even in this case of a more simple expected excess return.

Denote the classic factor construction weights as w_{pt}^{carry}. The hedging portfolio then has weights $w_t^* - w_{pt}^{carry}$. In fact, this is what Daniel, Mota, Rottke, and Santos (2020) propose in their analysis of unpriced equity risks.[4] Given this link, we label this alternative way of generating the carry hedge portfolio as the DMRS hedge

[4]They have to approximate w_t^* because the equity market does not allow for estimation of the conditional covariance matrix due to the large number of assets.

portfolio and use it as a robustness test relative to our optimal UMVE-based hedge portfolio. It has the benefit of relying only on the interest differential as the signal for expected returns, so it does not it involve other return signals such as the momentum or value, nor does it require a forecasting regression.

Figure 9 shows the SR of the original carry strategy and the SRs of the UMVE- and DMRS-hedged versions of the carry strategy. The former is 0.71 whereas the latter are 1.29 and 0.99, respectively. That is, there is a substantial increase in the SR. The figure also plots the information ratio (capturing the marginal increase in SR) of the carry hedged for unpriced risks, which is in both cases higher than the SR for the original carry. In sum, the alternative hedging strategy yields similar results with substantial increases in the carry risk-return trade-off.

Unpriced risks and geographic factors

A natural question is what these unpriced risks might represent. In the equity market. 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 factor which is long an equal-weighted portfolio of the remainder of the GE countries.[5] Next, we at each time t project the optimal portfolio weights of either the UMVE- or DMRS-based hedging portfolio

[5] 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).

weights onto the portfolio weights for these two portfolios. We use the projection as our hedge, which then only uses these two geographic portfolios.

The projected weights are shown in Figure 10. The hedging portfolio is generally short Europe and long Rest of the World, with the exception of the European currency crisis in the early 1990s. The weights are relatively stable and of reasonable magnitude.

Figure 9 shows the Sharpe and information ratios of the carry when unpriced risk is hedged out using the UMVE- and DMRS-based weights for the geographic factors as outlined above. In both cases, the SR increases substantially relative to the original carry, although not by quite as much as with the unconstrained UMVE- and DMRS-based hedges. 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.

A new hurdle: Implications for models of currency market risk premiums

Reconstructing popular trading strategies by real-time hedging of unpriced risks represents a new hurdle for testing asset pricing models for the currency market. Table 7 reports standard "alpha" regressions of various versions of the hedged carry on the dollar and the cross-sectional carry factors for the GE currencies. The alphas are large, positive, and strongly statistically significant. That is, the results and conclusions that we have reached using the optimal UMVE-based hedge hold for both DMRS-based simplified version and for further simplification on the basis of the geographical factors. The robustness of this conclusion calls for models that focus on explaining hedged returns.

Further, 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 enhance 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 an out-of-sample 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.

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, 2023, Concealed carry, *Journal of Financial Economics. forthcoming* .

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

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

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.

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.

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.

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.

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.

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

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.

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

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.

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, HEC Montreal.

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

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.

Figure 1
Investments in carry and momentum currency strategies

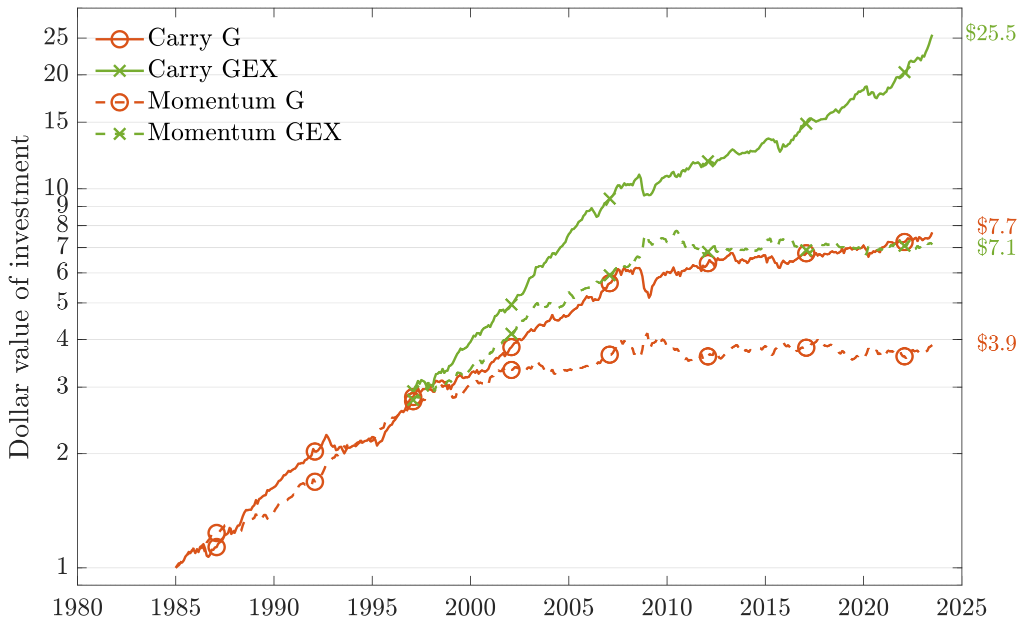

The figure shows dollar values of investments in the cross-sectional carry and momentum strategies (log-scale on the vertical axis). Curry and momentum investments are for both G10 currencies only (labeled G) and all currencies in our sample (labeled GEX). The cumulative gross return of an investment 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 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

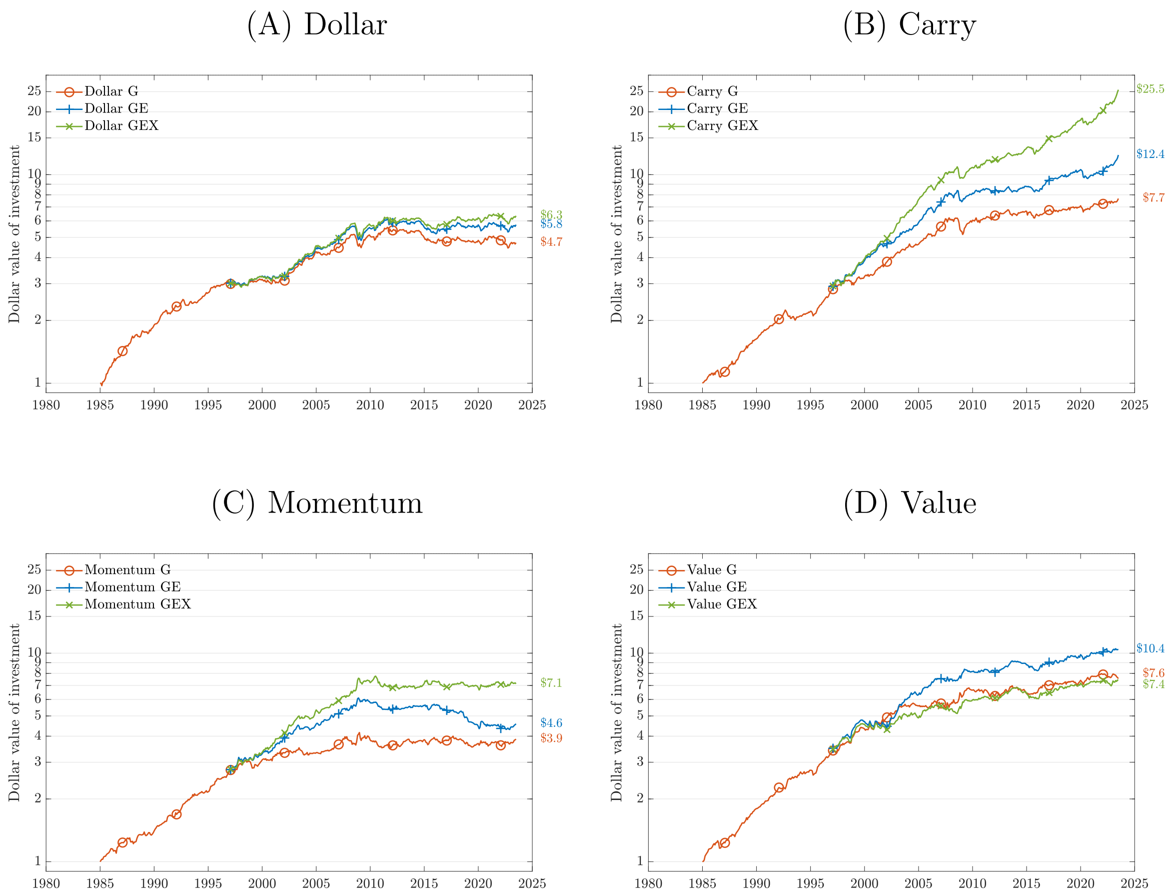

The figure shows dollar values of investments in the dollar, and cross-sectional carry, one-month momentum, and value strategies (log-scale on the vertical axis). The three lines in each panel correspond to strategies constructed from G10 currencies only (G, red line with circles), from G10 and floating-regime 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

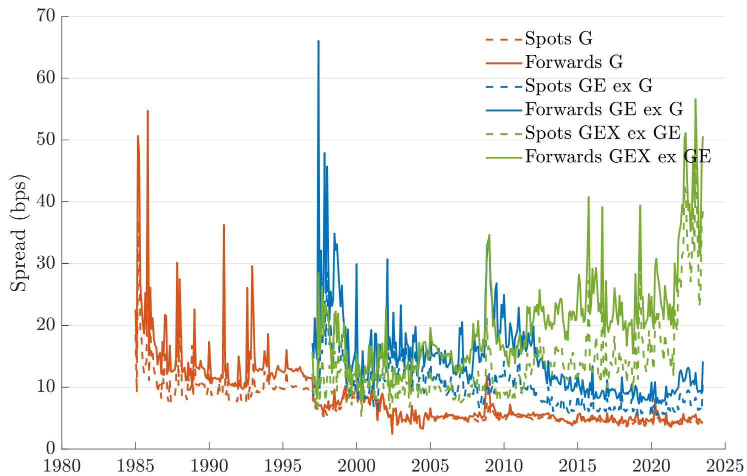

(B) Bid-offer spreads for G10 subgroups

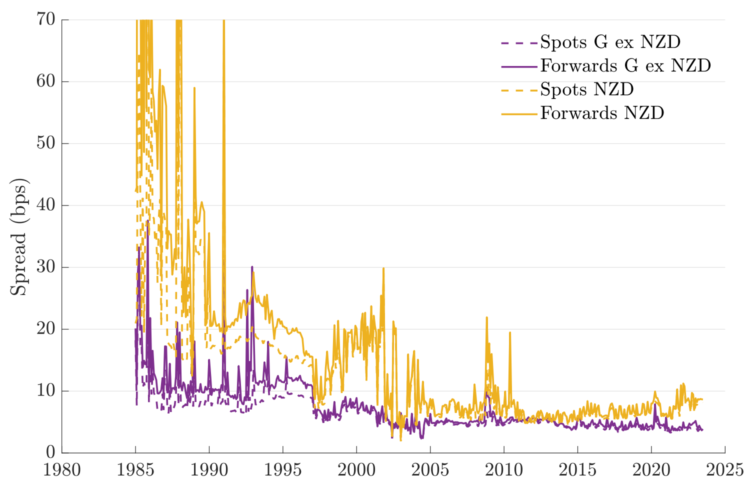

Panel A shows the average bid-offer spreads for spot (dashed lines) and forward (solid lines) markets for the three currency sets: G10 (G), floating emerging (GE ex G), and extended emerging (GEX ex GE). Panel B shows the impact of NZD costs on the G10 set by breaking it up into NZD only and G10 ex NZD. 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

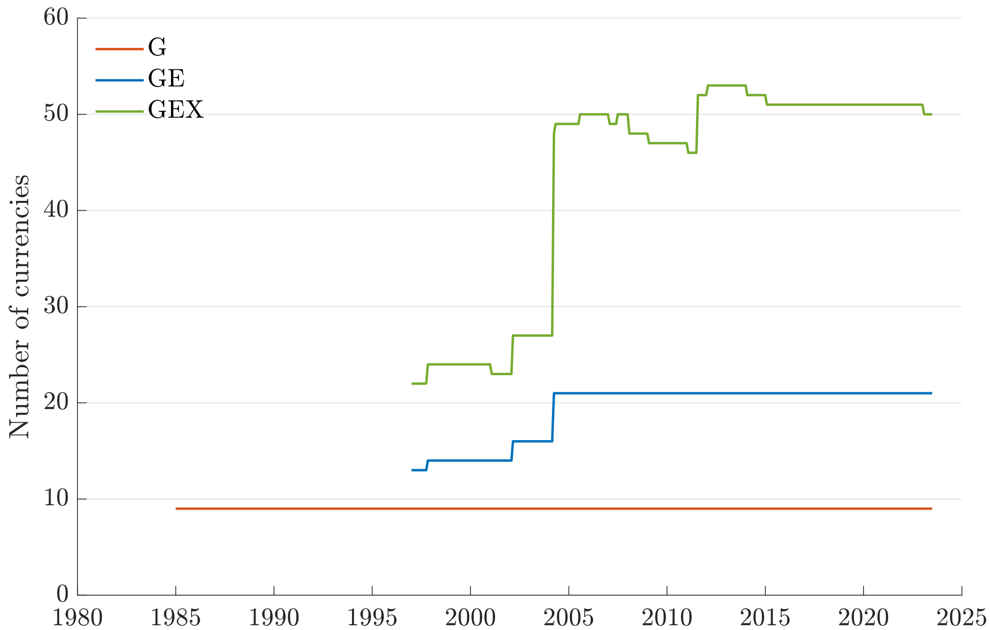

The figure shows the number of currencies within each currency set over time: G10 (G), G10 plus floating-regime emerging markets (GE), and G10 plus all emerging market currencies in our sample (GEX).

Figure 5
Sharpe ratio of strategies with different transaction costs

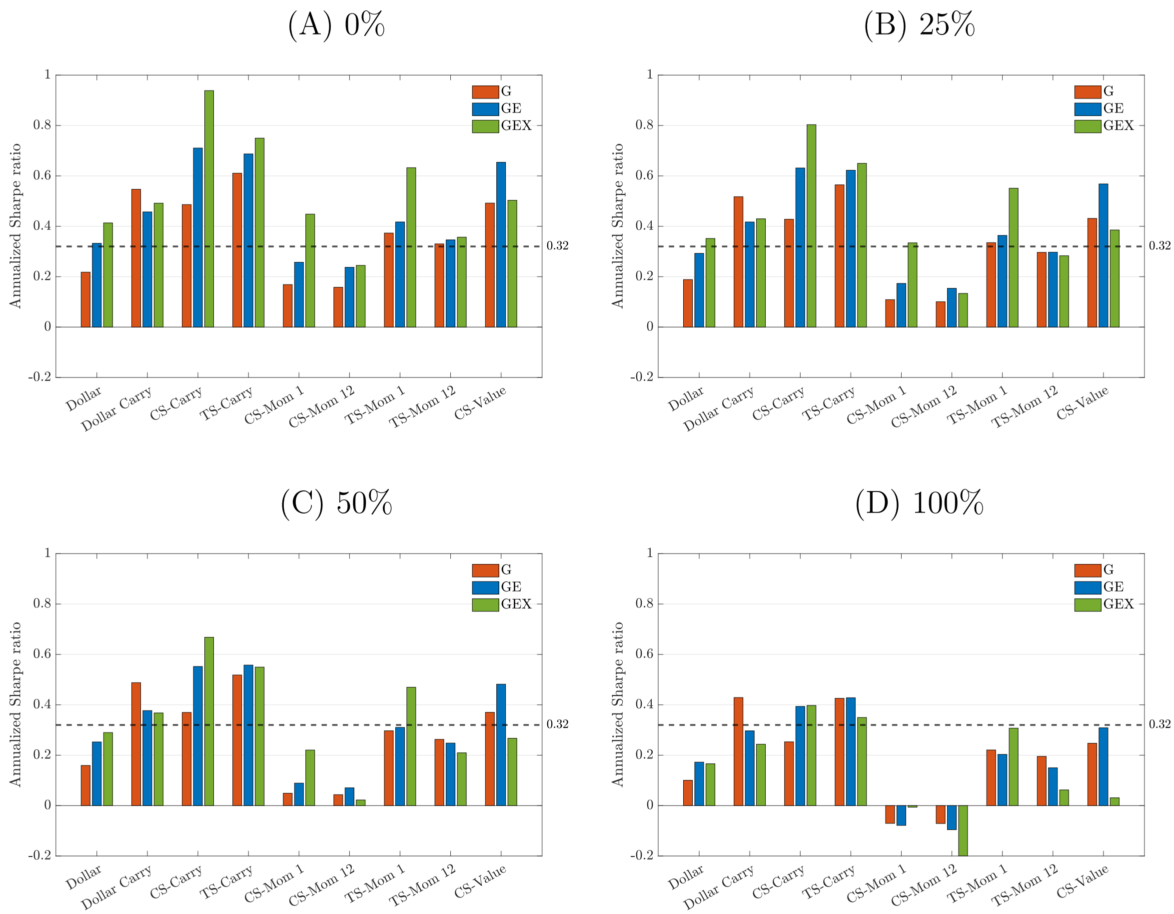

The figure shows annualized sample Sharpe ratios for the nine trading strategies computed for the different levels of transaction costs. The sample is monthly from 1985 to 2023.

Figure 6
Sharpe ratio of hedging portfolios and hedged strategies

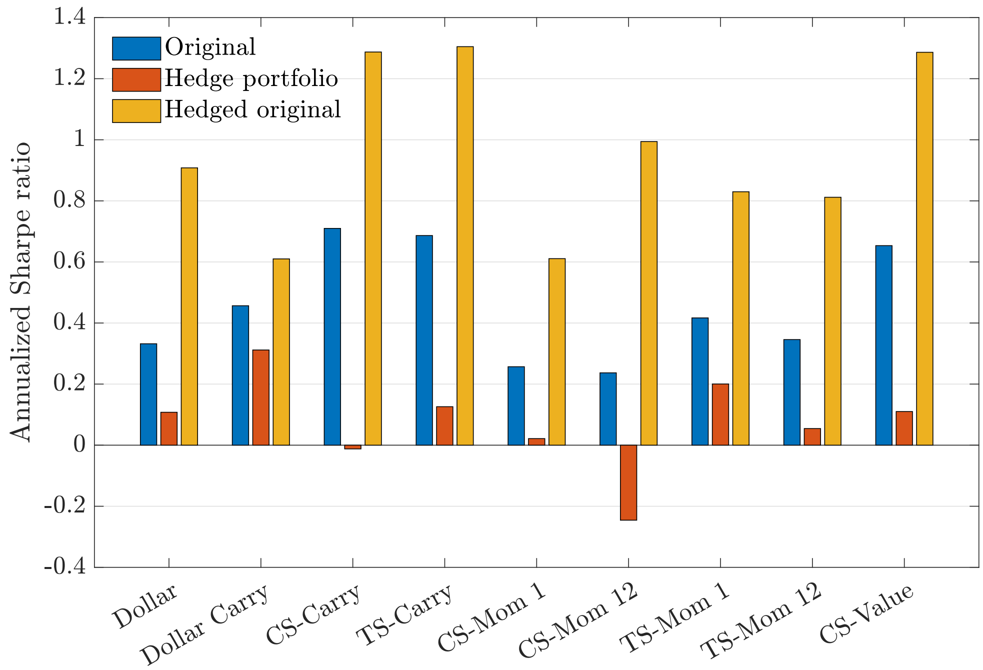

The figure shows annualized sample Sharpe ratios for three portfolios associated with each strategy. The "original" refers to the baseline version of the strategy, "hedge portfolio" refers to the hedging portfolio that hedges out unpriced risks in real time, and "hedged original" refers to the portfolio consisting of the original strategy return minus the hedging portfolio. The sample is monthly from 1985 to 2023.

Figure 7
Mean realized and expected returns in subsamples

(A) 1985–2004

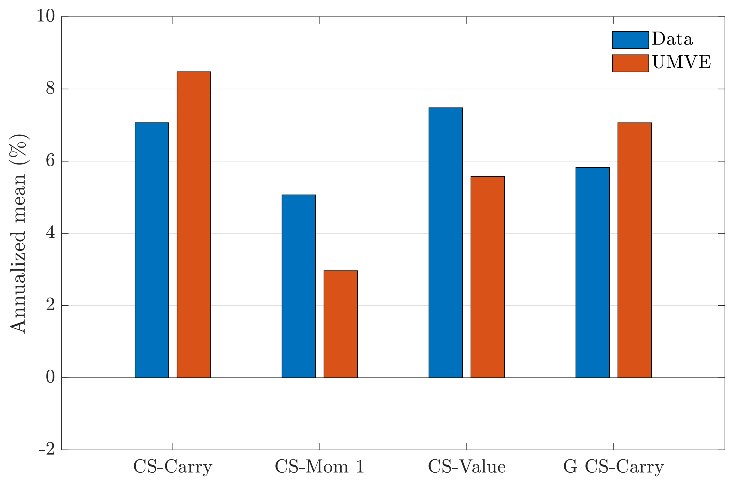

(B) 2005–2023

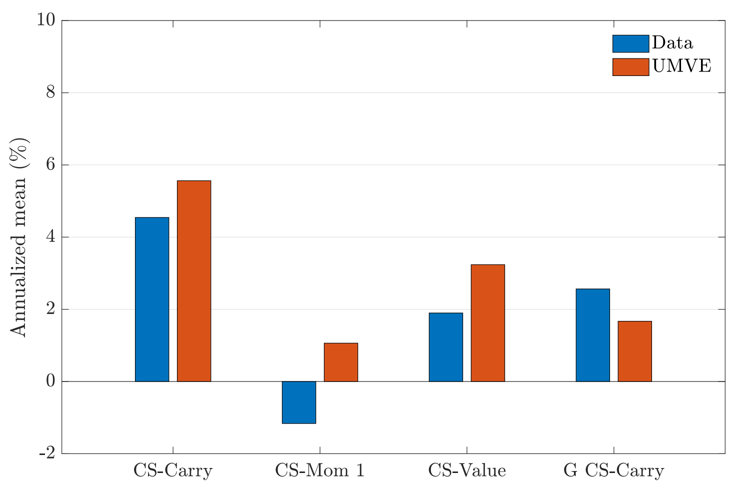

The figures shows sample average annualized returns for four strategies as given on the horizontal axis, as well as the average annualized conditional risk premium as given by our real-time UMVE construction. The upper panel shows these quantities for the first half of the sample, while the lower panel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

Figure 8
Conditional price of risk for G and GE UMVEs

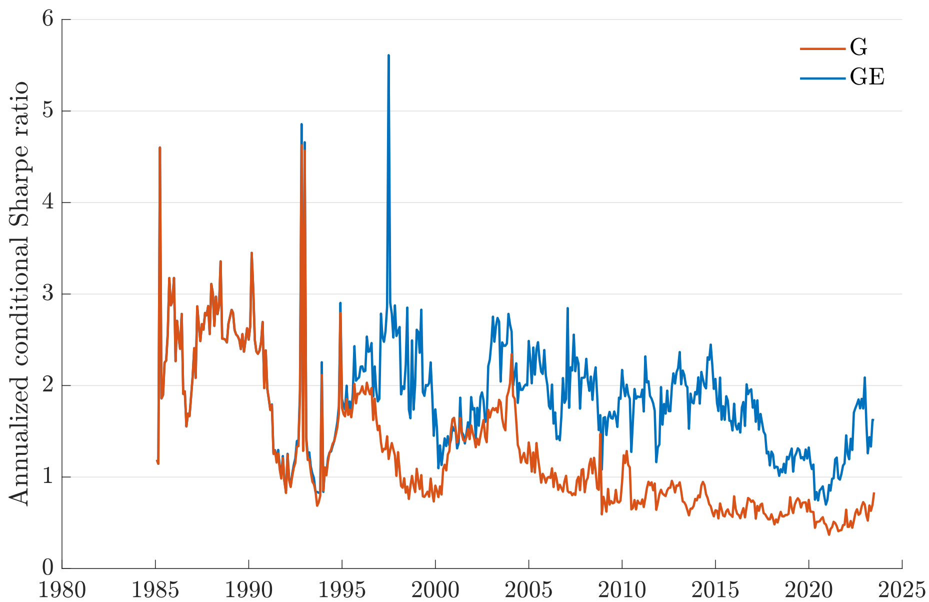

The figure shows the conditional annualized Sharpe ratio of the UMVE portfolio constructed from G10 (G) and G10 plus floating-regime emerging markets (GE) currencies. The sample is monthly from 1985 to 2023.

Figure 9
Sharpe and information ratios for hedging strategies

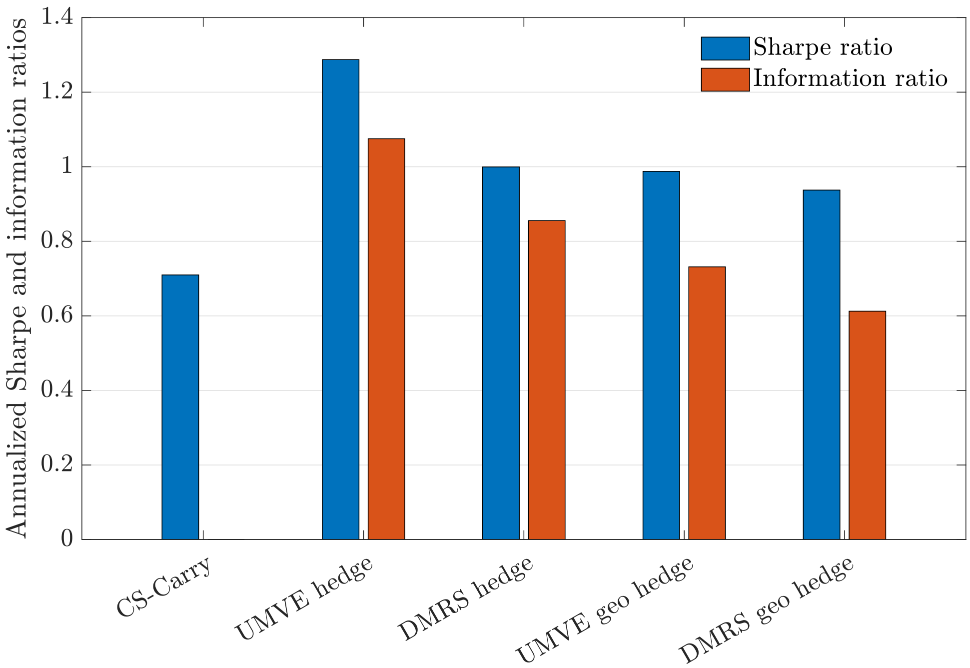

The figures gives sample annualized sample Sharpe ratios and information ratios for the cross-sectional carry and various versions of the strategy where different proxies for unpriced risks are hedged out. The sample is monthly from 1985 to 2023.

Figure 10
Geographic hedging exposures

(A) UMVE-based hedging betas

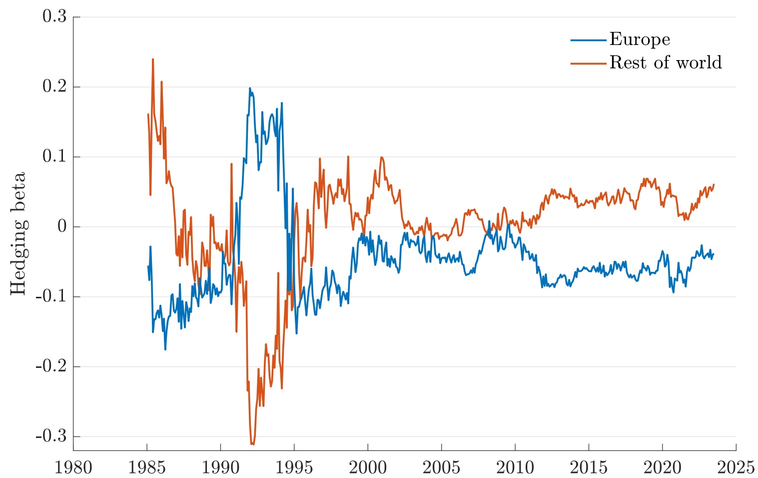

(B) DMRS-based hedging betas

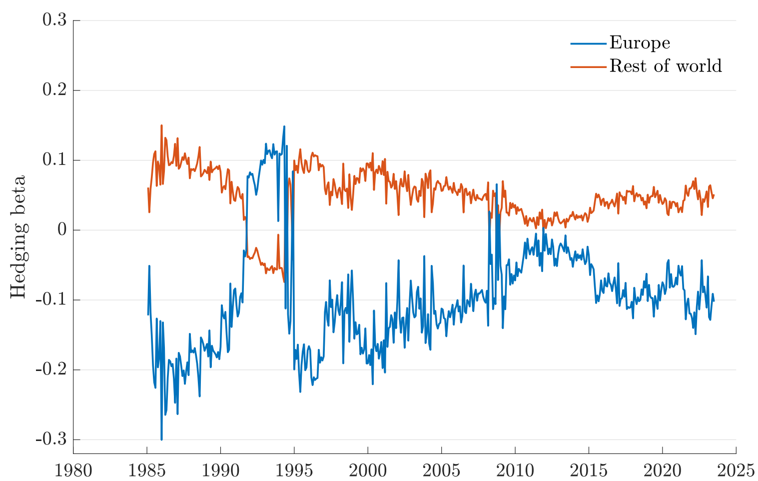

The figure shows the conditional beta on two geographic factors for the hedging portfolio for hedging out unpriced risks from the cross-sectional carry factor. The geographic factors are European currencies and Rest of the World currencies. The currency set this refers to is the G10 plus floating-regime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

Table 1: Summary statistics

		TC = 0			TC =0.25			TC =0.50			TC =1.00		
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	5.13	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	0.99	3.35	0.87	0.39	2.50	0.39	-0.21	1.65	-0.57	-1.41	-0.05
	std	8.04	7.39	7.49	8.04	7.39	7.50	8.03	7.39	7.51	8.03	7.39	7.53
	SR	0.17	0.13	0.45	0.11	0.05	0.33	0.05	-0.03	0.22	-0.07	-0.19	-0.01
	skew	0.14	0.33	0.20	0.12	0.31	0.20	0.11	0.30	0.20	0.09	0.28	0.19
CS-Mom 12	mean	1.29	0.13	1.88	0.82	-0.46	1.03	0.35	-1.05	0.17	-0.58	-2.24	-1.54
	std	8.14	7.43	7.69	8.14	7.44	7.69	8.14	7.45	7.51	8.14	7.47	7.71
	SR	0.16	0.02	0.24	0.10	-0.06	0.13	0.04	-0.14	0.22	-0.07	-0.30	-0.20
	skew	-0.48	-0.26	-0.54	-0.49	-0.28	-0.54	-0.50	-0.30	0.20	-0.52	-0.34	-0.55
TS-Mom 1	mean	2.31	1.97	3.12	2.07	1.68	2.72	1.84	1.39	2.32	1.36	0.82	1.52
	std	6.19	5.60	4.94	6.19	5.60	4.95	6.19	5.60	4.95	6.18	5.59	4.95
	SR	0.37	0.35	0.63	0.33	0.30	0.55	0.30	0.25	0.47	0.22	0.15	0.31
	skew	0.53	0.62	0.45	0.52	0.60	0.44	0.52	0.59	0.43	0.50	0.56	0.40
TS-Mom 12	mean	2.26	1.19	1.90	2.03	0.90	1.51	1.80	0.62	1.11	1.34	0.04	0.33
	std	6.86	6.14	5.32	6.85	6.14	5.32	6.85	6.14	5.32	6.85	6.14	5.33
	SR	0.33	0.19	0.36	0.30	0.15	0.28	0.26	0.10	0.21	0.20	0.01	0.06
	skew	-0.20	-0.37	-0.50	-0.20	-0.38	-0.51	-0.21	-0.39	-0.51	-0.22	-0.42	-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	-01.6	-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 statisticians 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, which is constructed using the GE set of currencies. The t–statistics are heteroskedasticity-adjusted. Panel A shows results for strategies constructed from GE currencies. Panel B considers GEX currencies. The sample is monthly from 1985 to 2003.

Table 3: Unconditional GRS-style tests with transaction costs

	Transaction costs			
	0	0.25	0.5	1
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. The 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 hypothesis 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 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 asymptotic GRS test. The sample is monthly from 1985 to 2023.

Table 4: Conditional tests of GE UMVE

Strategy	SR	$E\left(R^e\right)$	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 unconditional tests of the conditional implications of the GE UMVE model. The p–value for joint test is computed by bootstrap (p–value reported in parentheses is computed using the asymptotic GRS statistic). The sample is monthly from 1985 to 2023.

Table 5: Tests of dollar/carry 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)		0.017	(0.017)	

This table shows unconditional and conditional tests of the of the dollar/carry model using strategy returns. The p–value for joint test is computed by bootstrap (p–value reported in parentheses is computed using the asymptotic GRS statistic). The first (second) set corresponds to the unconditional (conditional) model. The sample is monthly from 1985 to 2023.

Table 6: Tests of dollar/carry model on hedged strategies

Strategy	SR	$E\left(R^e\right)$	α	t–stat	OOS α	t–stat
Dollar	0.91	1.75	1.10	3.71	0.81	4.16
Dollar Carry	0.61	1.19	0.77	2.47	0.65	3.31
CS-Carry	1.29	5.89	3.97	6.29	3.00	6.36
TS-Carry	1.31	2.72	1.82	6.36	1.40	6.51
CS-Mom 1	0.61	1.81	1.43	2.71	1.01	3.54
CS-Mom 12	0.99	3.46	2.44	4.85	2.07	5.87
TS-Mom 1	0.83	1.31	0.97	3.89	0.56	3.46
TS-Mom 12	0.81	1.78	1.12	3.56	0.90	4.49
CS-Value	1.29	3.97	2.84	6.56	1.99	6.13
p–value			0.000	(0.000)	0.000	(0.000)

This table shows unconditional and conditional tests of the of the dollar/carry model using strategy returns with unpriced risks hedged out (using the UMVE). The p–value for joint test is computed by bootstrap (p–value reported in parentheses is computed using the asymptotic GRS statistic). The first (second) set corresponds to the unconditional (conditional) model. The sample is monthly from 1985 to 2023.

Table 7: Hedged CS-carry versus standard dollar/carry factor model

CS-carry hedged by:	$E(R^e)$	α	t–stat	R^2_{adj}
UMVE	5.89	3.97	6.29	0.34
DMRS	8.62	7.02	4.90	0.06
UMVE Geo	6.81	2.45	4.17	0.74
DMRS Geo	6.43	2.22	3.50	0.69

This table shows the result from regressing different versions of hedged CS-carry on a standard factor model consisting of the dollar factor and the CS-carry factor. All factors and test assets are constructed from the G10 plus the floating-regime emerging markets currencies (GE). The labels in the leftmost column refer to the dependent variable in each regression and the method by which unpriced risks are hedged in real time for the hedged CS-carry strategy. UMVE refers to the optimal hedging given the UMVE portfolio construction. DMRS refers to the optimal hedging assuming the only expected return signal is the interest rate differential. UMVE Geo refers to the hedging portfolio that obtains from a conditional projection of the optimal hedging weights onto two geographical factors – a Europe factor and a Rest of the World factor. Finally, DMRS Geo refers to a conditional projection of the DMRS hedging weights onto the same geographic factors. The sample is monthly from 1985 to 2023, and all hedging is done in an out-of-sample fashion to ensure that the test assets are tradeable portfolios.

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 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:

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

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.3 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 $S_t^i/F_t^i - 1$ 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

$$Q_t^i = S_t^i \cdot P_t^i / P_t, \tag{A.3}$$

where P_t and P_t^i 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:

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

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

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

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:

$$S^i_{t+1}/S^i_t - 1 = \bar{\delta}_t \cdot z^i_{Qt} + \bar{\phi}_t \cdot \left(S^i_t/S^i_{t-12} - 1\right) + \varepsilon^i_{t+1}.$$

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:

$$\begin{aligned}\mu_{it} &= (S^i_t/F^i_t) \cdot E_t(S^i_{t+1}/S^i_t) - 1 \\ &= (S^i_t/F^i_t - 1) + (S^i_t/F^i_t)\bar{\delta}_t \cdot z^i_{Qt} + (S^i_t/F^i_t)\bar{\phi}_t \cdot \left(S^i_t/S^i_{t-12} - 1\right)\end{aligned}$$

with $(S^i_t/F^i_t)\bar{\delta}_t$ and $(S^i_t/F^i_t)\bar{\phi}_t$ corresponding to δ^i_t and ϕ^i_t 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.

A.4 Testing the UMVE constructed from G10 currencies

Tables A1 and A2 report tests of trading strategies using the G-UMVE.

Table A1: 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 asymptotic GRS test. The strategy returns are constructed using the G set. The sample is monthly from 1985 to 2023.

Table A2: 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 asymptotic GRS test in parenthesis). The strategy returns are constructed using the GE set. The sample is monthly from 1985 to 2023.