Rethinking Currency Factors: The Case for Mean-Variance Optimisation

Minyou Fan[a,*], Fearghal Kearney[a], Youwei Li[b], Jiadong Liu[a]

[a] *Queen's Business School, Queen's University Belfast, UK*
[b] *Hull University Business School, University of Hull, UK*

Abstract

We show that a simple mean-variance (MV) optimisation can substantially enhance the performance of established currency factor strategies such as Carry, Value, and Momentum. We find that the improved performance is due to the stronger cross-sectional predictability of the optimised strategies. International diversification plays a key role in optimisation due to 1) the low correlation between developed and emerging currencies, and 2) the low level of comovement across emerging currencies. We also find that the outperformance of our proposed MV optimised factor portfolios is positively related to the standard deviation of currency abnormal returns over time. Our asset pricing tests suggest that the MV optimised factors subsume the corresponding plain currency factors.

Keywords: Currency factor; Foreign exchange; Portfolio management; Mean-variance optimisation; International diversification

JEL: F31, G11, G12, G15

*Corresponding author. Queen's Business School, Queen's University Belfast, Riddel Hall, 185 Stranmillis Road, BT9 5EE, UK. email: minyou.fan@qub.ac.uk.

1. Introduction

Empirical studies in equity markets have shown that more sophisticated portfolio optimisation approaches often fail to outperform naive diversification (DeMiguel et al., 2009, Hsu et al., 2018). However, this dynamic is not observed in currency markets, where mean-variance optimisation has been proven to deliver consistently significant out-of-sample profits (Baz et al., 2001, Della Corte et al., 2009, Ackermann et al., 2017, Daniel et al., 2017). Chernov et al. (2023) further argued that an unconditional mean-variance efficient portfolio can effectively price currency risks. This argument provides insights that currency asset pricing can shift away from the most prominent strand of literature on studying currency characteristics such as Carry (Lustig et al., 2011), Momentum (Menkhoff et al., 2012*b*) and Value (Asness et al., 2013) factors.

In this paper, we provide a link between currency factors and portfolio optimisation by proposing a simple mean-variance optimiser to existing currency factors. Specifically, we optimise three of the most extensively documented currency factor portfolios, namely carry, momentum and value. We term these portfolios as mean-variance (MV) optimised factors. This means the MV optimised factors are constructed based on information conditional on expected currency returns, the covariance matrix and currency characteristics. These MV optimised factors significantly beat both their corresponding plain currency factors and the mean-variance optimised portfolio unconditional on currency characteristics.

Through a performance attribution analysis, we investigate whether the outperformance of the MV optimised factors is due to its enhanced time series or cross-sectional predictability. Using the decomposition model of Lo and MacKinlay (1990) and Lewellen (2002) for currency trading strategies, we find that the negative cross-covariance component is much stronger for the optimised factors in comparison with the naive factors. This means that the optimised currency factors can better capture the cross-sectional variation of currency returns.

We find that international diversification plays a pivotal role in explaining the enhanced performance of the MV optimised factors, as their performance is more pronounced in periods when the investable currencies are more diversified between developed and emerging economies. More specifically, to rationalise the high excess returns of the MV optimised factors, we highlight two reasons: 1) the low correlation between developed and emerging currencies, and 2) the low level of comovement across the emerging currencies relative to developed currencies. More recent studies on currency market optimisation, such as Ackermann et al. (2017), Maurer et al. (2022), Maurer et al. (2023) and Chernov et al. (2023), mainly focus on using a small cross-section of currencies.[1] However, we provide new evidence that highlights the importance of diversification by adopting a relatively large cross-section of 48 currencies as in Menkhoff et al. (2012*a*) and Menkhoff et al. (2012*b*). Similar to the dynamic in equity markets, where a stock with a negative beta is deemed a valuable asset in a portfolio, a currency with a lower correlation to the main currencies is also deemed to be valuable.

DeMiguel et al. (2009) find that the optimal portfolios consistently fail to outperform the naive portfolio based on the assumption that CAPM α is equal to zero. However, even though empirically the mean of α is quite close to zero (Jarrow, 2010), the assumption of a precisely zero CAPM is not realistic in the real world. Motivated by this, Platanakis et al. (2021) allow for a zero mean distribution of CAPM α and a standard deviation of 30 basis points. Their results show that the MV portfolios significantly outperform the 1/N rule. Our empirical finding in the currency market is consistent with Platanakis et al. (2021), where the outperformance of MV factor portfolios is positively related to the cross-sectional standard deviation of currency abnormal returns. According to our results of the F-test, the cross-sectional standard deviation of emerging currencies' abnormal return, α, significantly exceeds that of the developed sample at the level 1%, confirming our argument that the emerging currencies feature higher return dispersion. We use a simple linear model to show the positive relationship between the outperformance of MV factor portfolios and the standard deviation of currency abnormal returns.

We further explore the asset pricing implication of the MV optimised factors and examine the relationship between these optimised factors and their naive counterparts in currency pricing models. First, spanning tests reveal that abnormal returns of the MV carry and momentum factors cannot be fully explained by naive factors. Second, based on the framework of Fama and French (2015), we assess the performance of MV currency factors in multi-factor

[1] Ackermann et al. (2017) demonstrate that mean-variance portfolio optimisation beats naive diversification in currency markets as it benefits from a small cross-section.

currency pricing models. Replacing the naive factors with their MV counterparts, the predictability of asset pricing models on currency excess returns is significantly improved across all four model evaluation metrics. Notably, a model comprising all three MV optimised factors and the naive momentum factor outperforms all other specifications, providing better insight for currency pricing research.

Our study contributes to our understanding of the benefit of MV optimisation in improving the performance of existing currency factors. Our research is different from studies that directly adapt optimisation for currencies without factor signals, such as Baz et al. (2001) and Ackermann et al. (2017), or that focus on improving the estimation of currency expected returns estimation based on factor signals, such as Opie and Riddiough (2020), Maurer et al. (2022), and Chernov et al. (2023). Kroencke et al. (2014) proposed a combined signal based on carry, momentum and value, but simply averages the ranks across three signals without using additional information apart from currency characteristics. In contrast, we directly employ optimisation to the constituents of the long and short legs of currency factors to improve performance.

Finally, these MV optimised factors provide economic value by enhancing currency factor portfolios, leading to superior performance compared to traditional currency factors. Despite their optimisation, the MV factors preserve the original portfolio structure, ensuring balance between long and short positions with no net exposure. Unlike Della Corte et al. (2009), Barroso and Santa-Clara (2015), and Maurer et al. (2023), our strategy avoids leverage, ensuring that portfolio weights remain stable over time. Incorporating notional values may further refine MV optimisation, improving performance beyond the current framework. For further analysis of the profitability of optimised currency factors, see Fan et al. (2024).

The remainder of this paper is organised as follows. In Section 2, we describe our dataset and the main methodologies used in constructing currency factors and in mean-variance optimisation. Next, Section 3 compares the performance between the proposed MV optimised factors and those using naive diversification. Section 4 analyses the outperformance attribution of our proposed strategies. In Section 5, we test whether the MV optimised factors can be employed in currency asset pricing. Finally, Section 6 concludes.

2. Data and Factor portfolios

This section begins by outlining our data universe and the methodology for computing currency excess returns and transaction costs. We then proceed to describe the construction of currency factor portfolios and the proposed mean-variance optimisation process.

2.1. Data

We use a comprehensive FX market dataset proposed by Menkhoff et al. (2012*a*) and Menkhoff et al. (2012*b*), which consists of 48 currencies from the following countries or regions: Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Cyprus, Czech Republic, Denmark, Egypt, Euro area, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, India, Indonesia, Ireland, Israel, Italy, Japan, Kuwait, Malaysia, Mexico, Netherlands, New Zealand, Norway, Philippines, Poland, Portugal, Russia, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Ukraine, the United Kingdom, and the United States. Our sample period begins in November 1983 and ends in November 2020, with summary statistics reported in Appendix A. We obtain spot and 1-month forward exchange rates for each currency from Barclays Bank International (BBI) and WM Reuters (WMR) via Datastream.

We first measure the monthly gross return as the purchase of one unit of a foreign currency. More precisely, following previous literature (e.g., Fama (1984)), we calculate the logarithms of our spot rate and compute the monthly gross return of currency k in month t as $r_t^k = s_t^k - s_{t-1}^k$, where s_t^k represents the spot exchange rate at the end of the month, measured in US Dollars for each unit of the foreign currency k. When r_t^k is negative, currency k depreciates against the US Dollar.

Next, we measure the monthly excess returns as:

$$rx_t^k = (i_{t-1}^k - i_{t-1}) - r_t^k, \tag{1}$$

where i_{t-1}^k is the foreign interest rate in region k in month $t-1$; and i_{t-1} is the US interest rate. Therefore, $i_{t-1}^k - i_{t-1}$ refers to the interest differential between economy k and the US. We follow Taylor (1989), Akram et al. (2008), Menkhoff et al. (2012*a*) and Menkhoff et al. (2012*b*) in estimating the interest rate differentials using forward discounts. Thus, we have the short-term interest differential $i^k_{t-1} - i_{t-1} = f^k_{t-1} - s^k_{t-1}$, where f^k_{t-1} is the one-month forward exchange rate between currency k and the US dollar in month $t-1$ (foreign exchange rate per unit of US dollar). Then, we compute the monthly excess returns of currency k in month t as:

$$rx^k_t = (i^k_{t-1} - i_{t-1}) - r^k_t \approx f^k_{t-1} - s^k_t. \tag{2}$$

In practice, transaction costs can have a significant impact on the measurement of investment revenues, with Timmermann and Granger (2004) documenting that excess returns generated by exchange rate spreads can often be eliminated in the presence of transaction costs. Menkhoff et al. (2012*b*) estimates transaction costs using the full quoted bid-ask spread, however, Neely and Weller (2013) proposed that the full quoted bid-ask spreads tend to be larger than the effective spread rates that are actually traded. We follow Neely and Weller (2013) and compute one-way approximate transaction costs as one-third of the one-month forward bid-ask spread. This transaction cost measurement is extensively used in FX studies, e.g., Hsu et al. (2016) and Filippou et al. (2018). Therefore, the transaction-adjusted portfolio return is:

$$R^p_t = ln\{1 + \sum_{k=1}^{N} [exp(rx^k_t) - 1] w^k_t - c^k_t \times \sum_{k=1}^{N} |w^k_{t+1} - w^k_t|\}, \tag{3}$$

where N refers to the number of currencies in the portfolio, c^k_t represents the transaction cost for currency k in month t, w^k_t is the weight of currency k before the rebalancing at the end of month t, and w^k_{t+1} is the weight in the coming month. w^k_t is positive (negative) if the asset k is assigned to the long (short) portfolio. In the next section, we present the performance of the naive and optimised factors.

2.2. Currency factors

A factor portfolio is usually constructed by sorting on an asset characteristic and calculating the return difference between *high* and *low* portfolios (Fama and French, 1992). We implement three widely used currency market anomalies: the carry factor of Lustig et al. (2011) and Menkhoff et al. (2012*a*), the currency momentum factor of Menkhoff et al.

(2012*b*), and the currency value factor of Moskowitz et al. (2012) and Menkhoff et al. (2017).[2]

Carry refers to the interest differential measured by the difference between a spot rate and a one-month forward rate, as defined by Lustig et al. (2011). Consequently, a carry strategy involves taking a long position in currencies with higher interest rates and a short position in currencies with lower interest rates. This is based on the premise that higher interest rate currencies tend to appreciate, while lower interest rate currencies tend to depreciate. As the covered interest parity theorem holds at the frequency used in our study (Akram et al., 2008), we estimate the cost of carry using the forward discount, $f_t - s_t$ in line with the approach in Menkhoff et al. (2012*a*). A carry factor portfolio is then created by buying (selling) the currencies with high (low) costs of carry.

Momentum is a cross-sectional trend market anomaly which has been extensively documented in various financial markets. In a momentum strategy, as defined by Jegadeesh and Titman (1993), abnormal profits are generated by buying the securities that are past winners and selling those that are past losers. More recently, Menkhoff et al. (2012*b*) and Asness et al. (2013) uncover momentum profits in the currency market. In line with Menkhoff et al. (2012*b*), we adopt a momentum strategy with a three-month formation period as a proxy of the currency momentum anomaly. A momentum factor portfolio is subsequently constructed by going long (short) the currencies with high (low) cumulative returns over the formation period.

Value is a type of mean reversion currency anomaly outlined in Asness et al. (2013) and Menkhoff et al. (2017). Its premium is sourced from the Real Exchange Rate (RER) as follows:

$$RER_t = s_t \times \frac{c_t}{c_{f,t}}, \tag{4}$$

where $c_{f,t}$ and c_t refer to the foreign and domestic inflation rate in month t, respectively. $\frac{c_t}{c_{f,t}}$ is also known as the Purchasing Power Parity (PPP). Balduzzi and Chiang (2020) further uncovered that the real exchange rate strongly and negatively predicts future currency returns. In line with Asness et al. (2013), we estimate the currency value as the log difference

[2]Given that our optimisation procedure focuses on general cross-sectional factors, we exclude the factors with asymmetric portfolios, such as the dollar factor from Lustig et al. (2011) and the Euro factor from Greenaway-McGrevy et al. (2018).

in the real exchange rates over the past 60 months (5 years) as:

$$V_t = log(\frac{spot_t * c_t}{c_{f,t}}) - log(\frac{spot_{t-60} * c_{t-60}}{c_{f,t-60}}), \tag{5}$$

where $spot_{t-60}$ is the average of spot rates between 4.5 and 5.5 years ago. The premise of the *Value* strategy is that currencies with low RER relative to the US Dollar have higher returns as the strategy is long the low-value currencies and short the high-value currencies. To build the portfolio for each of the above-mentioned currency anomalies, we construct high-minus-low portfolios (or a low-minus-high portfolio for value) by dividing the whole sample into two equal portfolios.

2.3. Mean-variance optimisation

Next, we outline our mean-variance method to optimise currency factor portfolios presented above. Assuming N assets are available at time t, we consider all assets as a universe (U):

$$U = \{a_1, a_2, \cdots, a_N\} \tag{6}$$

where a_N denotes the N-th asset. All assets are sorted in descending order based on currency characteristics for momentum and carry signals and in ascending order for value signals.

When a factor portfolio covers all the available assets, half of the available assets are allocated to the long portfolio, with the other half allocated to the short portfolio. As both the long and short portfolios have $\frac{N}{2}$ assets, the long/short portfolios can be expressed as two subsets of U as follows:

$$\begin{aligned} A &= \{a_1, a_2, \cdots, a_{N/2}\}, \\ B &= \{a_{N/2+1}, a_{N/2+2}, \cdots, a_N\}, \end{aligned} \tag{7}$$

where A contains all assets taking long positions, and B includes all assets in the short portfolio.[3] A given asset can not be allocated into two different factor portfolios simultaneously,

[3]Here, we assume N is an even number. However, if the number of assets in U is odd, we can remove the currency ranked in the middle and take $(N-1)/2$ assets in each of the long/short legs.

so we derive the property that,

$$A \cap B = \emptyset. \tag{8}$$

Two notable properties of sets A and B suggest that the long and short portfolios can be optimised independently. First, as indicated by Equation 8, the long and short portfolios do not overlap. Second, sets A and B exhibit distinct expected return functions, which results in different optimisation targets. In set A, assets with higher expected returns should be assigned greater weights, whereas in set B, assets with lower expected returns should receive higher weights. Therefore, the weights of the long/short portfolio (w_{long}/w_{short}) can be expressed as two vectors,

$$\begin{aligned} w_{long} &= [w_1, w_2, \cdots, w_{N/2}], \\ w_{short} &= [w_{N/2+1}, w_{N/2+2}, \cdots, w_N], \end{aligned} \tag{9}$$

where w_N denotes the weight of a given asset. According to Goyal and Jegadeesh (2017), the asset weights in the long (short) portfolio should add up to one (minus one), so we hold $\sum_{i=1}^{N/2} w_i = 1$ and $\sum_{i=N/2+1}^{k} w_i = -1$. In an original factor portfolio, the weight of each asset in the long (short) portfolio is denoted by $\frac{1}{N/2}$ $(-\frac{1}{N/2})$.

To determine the optimal w_{long} and w_{short}, under mean-variance framework of Markowitz (1952), it is assumed that investors choose the asset weights vector, w_t, that maximises their expected utility:

$$max_{w_t} \quad w_t' \mu_t - \frac{\gamma}{2} w_t' \Sigma_t w_t, \tag{10}$$

where γ is their level of relative risk aversion. The selection of γ becomes inconsequential in our method, as the portfolio is normalised to ensure that the weights sum to 1 regardless. The optimal w_t is not achievable unless we estimate μ_t and Σ_t. Here, the expected return and the covariance matrix are estimated using basic *Exponential Weighted Moving Average (EWMA) estimators*, which is a simple dynamic model in which recent returns have more weight than past returns in the estimation. It estimates the out-of-sample expected return by taking the sum of those weighted returns as follows:

$$\hat{\mu}_t = \sum_{i=1}^{T}(R_i \times \frac{\lambda_i}{\sum_{i=1}^{T}\lambda_i}) \tag{11}$$

Similarly, the out-of-sample covariance matrix can be measured as:

$$\Sigma_t = \lambda\hat{\Sigma}_t + (1-\lambda)\hat{\mu_t}'\hat{\mu_t}, \tag{12}$$

where $\hat{\mu_t}'$ is the transpose of return vector $\hat{\mu_t}$. Following Ardia and Boudt (2015), we set λ to 0.94. It is well-documented that in-sample-based estimators often lead to high estimation errors (Elton and Gruber, 1997, Kolm et al., 2014). By testing our research questions under this worst-case scenario, we demonstrate that if the optimised factor portfolios can outperform the naive ones despite such high estimation errors, they are likely to achieve even greater outperformance when using more improved estimators. In addition, the weight of a single currency is constrained to a range between 0.01 and 0.5 to allow for diversification, and the estimation window is 60 months.[4]

3. Performance Comparison: MV vs Naive factors

In this section, we assess the profitability of the MV currency factors and highlight its superiority to the naive factors and the MV optimised currency portfolios. As our value factor portfolios are available from the 67th month of the full sample period due to the use of a 5.5-year formation period, we measure the performance of the naive and MV factor portfolios starting from the 67th month to ensure that all three factors are measured over the same sample period. This leads to 380 out-of-sample months spanning April 1989 to November 2020.

Table 1 presents an overview of the performance of the naive and MV factors portfolios, incorporating transaction costs as outlined in Section 2.1. For each factor, we present the performance summaries of both the naive and mean-variance (MV) portfolios. The "Diff"

[4] If an asset is selected by the factor portfolio but does not have sufficient data for expected return and covariance estimation, we simply allocate a weight of $\frac{1}{N/2}$ to it and optimise the rest of the assets.

columns refer to the summary statistics of the return differentials between the naive and MV factor portfolios. The performance metrics we use are annualised average return (*Mean*), standard deviation (*Vol*), t-value, Sharpe ratio (*Sharpe*), t-value of the Sharpe ratio, Omega ratio (*Omega*), Dowd ratio (*Dowd*) and Certainty Equivalent Rate (*CER*), as described in Appendix B. Optimal portfolio (Opt) in the last column refers to the performance of the Mean-variance portfolio unconditional on currency characteristics.[5]

For the naive portfolio with equally weighted schemes in both the long and short legs, the momentum factor exhibits the strongest performance among the three factors, with the highest average annual return of 4.84% and a Sharpe ratio of 0.944. Following closely, the Carry factor displays the second strongest profitability with an average return of 3.88% and a Sharpe ratio of 0.674. Notably, both the carry and momentum factors report statistically significant profits at the 1% level. In contrast, the Value factor generates an insignificant and negative annual return at -0.32% using the naive approach. Furthermore, the momentum factor reports the best OOS performance with the highest Omega ratio (2.040), Dowd ratio (0.178), and CER (0.348%).

Remarkably, the out-of-sample performance improves considerably upon incorporating the MV factors. Table 1 reveals that the MV factors consistently yield a mean return that is at least 3% higher than their naive counterparts. The most profitable factor, Carry, stands out with impressive growth in annual return from 3.88% to 10.87% after applying MV optimisation. Furthermore, the optimised Value portfolio even leads to a positive return of 2.99% per annum, versus the loss made following the naive approach.

Although the outperformance of the MV factors is often accompanied by greater return volatility, the MV factors still exhibit higher risk-adjusted performance in comparison with the naive one. The t-values of the Sharpe ratios indicate that the risk-adjusted returns for the Carry and Value factors are significantly improved at the 10% and 5% level, while

[5] In line with our MV factors, we construct the mean-variance optimised portfolio using exponentially weighted moving average (EWMA) estimators for both expected returns and covariance matrices with a rolling 60-month formation period. We impose the same weight constraints such that individual currency allocations are bounded between 0.01 and 0.5.

the improvement for the Momentum factor is statistically insignificant.[6] Notably, the MV Value factor demonstrates the most pronounced improvement in the Sharpe ratio, rising from -0.070 to 0.234, underscoring its superior risk-adjusted performance.

Improvements are observed for the Omega ratios across all three factors, with the most significant change observed in the carry factors (increasing from 1.702 to 2.130). Similarly, the Dowd ratios for the MV factors exceed those of their naive counterparts. It should be highlighted that the optimised momentum factor has a lower Sharpe Ratio than its naive equivalent, at 0.871 versus 0.944. However, the Dowd and Omega ratios and the CER improve using the MV optimised approach.

Table 1: Out-of-sample performance using naive and mean-variance optimised factors

Strategies	Carry			Momentum			Value			Opt
	Naive	MV	Diff	Naive	MV	Diff	Naive	MV	Diff	
Mean	3.88%***	10.87%***	7.00%***	4.84%***	10.52%***	5.69%***	-0.32%	2.99%	3.30%*	6.69%***
t-value	(3.79)	(5.24)	(3.97)	(5.31)	(4.90)	(3.14)	(-0.39)	(1.32)	(1.74)	(4.15)
Vol	0.06	0.12	0.10	0.05	0.12	0.10	0.05	0.13	0.11	0.09
Sharpe	0.67*	0.93**	0.71*	0.94***	0.87***	0.56	-0.07	0.23	0.31*	0.74***
SR.t	(3.79)	(5.23)	(1.83)	(5.30)	(4.90)	(0.45)	(-0.39)	(1.31)	(2.20)	(3.89)
Omega	1.70	2.13	1.79	2.04	2.08	1.59	0.95	1.21	1.29	1.92
Dowd	0.15	0.22	0.15	0.18	0.20	0.12	-0.01	0.05	0.06	0.15
CER	0.25%	0.62%	0.38%	0.35%	0.57%	0.26%	-0.07%	-0.09%	0.04%	0.39%

This table summarises the out-of-sample performance of the naive and mean-variance factor portfolios with transaction costs from April 1989 to November 2020. Opt represents the mean-variance portfolio across all currencies. Diff refers to the series of differences between naive and mean-variance optimised factor returns. *Mean*, t-value, *Vol*, and *Sharpe* denote the annualised average returns, standard deviation, t-statistic of mean, and Sharpe ratio of portfolio returns. SR.t denotes the t-value of the Sharpe ratio based on bootstrapped Sharpe ratio equality test Ledoit and Wolf (2008). *Omega, Dowd,* and *CER* refer to the Omega ratio, Dowd ratio, and Certainty Equivalent Return, respectively. *, **, *** represent that the t-values are statistically significant at 10%, 5% and 1% level.

Figure 1 plots the cumulative performance of the MV optimised and the naive factor portfolios over the investment horizon after accounting for transaction costs. The plot highlights that our MV optimised factor portfolios (red lines) generate significantly superior cumulative performance compared to the original naive factor portfolios, which aligns with our findings in Table 1. For all three factors, the dollar values of investments based on the MV optimised portfolio are broadly similar to the naive portfolio before 2000. The MV Value portfolio greatly exceeds that of the naive Value portfolio, which exhibits a cumulative return of nearly zero over time. As the composition of our sample changes over time,

[6] We measure the t-values of the Sharpe ratios based on the bootstrapped Sharpe ratio equality test of Ledoit and Wolf (2008) to mitigate the impacts from the outliners.

we posit that the performance of the MV factor portfolios may be related to the effect of international diversification. We explore this further in the next section.

We further divide the entire sample period into three sub-periods according to different time regimes. Table 2 presents the analysis of portfolio performance across three the sub-periods, 1) pre-1999 before the launch of the Euro when few emerging currencies were included in the sample, 2) 1999-2009 covering major economic events such as the launch of the Euro and the 2007-2008 global financial crisis, and 3) post-2010 referring to the post crisis period. The MV factors demonstrate consistent outperformance over the naive approach across most combinations of factors and sub-periods. This outperformance is most pronounced during the turbulent period of 1999-2009, where the differences between optimised and naive factors are significant across the three factors. Specifically, the MV strategies outperform the naive strategies by 15.64%, 7.2% and 12.24% per annum for the carry, momentum, and value portfolios, respectively.

The carry factor exhibits the most compelling evidence of MV strategy effectiveness. During the pre-1999 period, the carry portfolio shows modest outperformance (Sharpe ratio differential of 0.082). However, following the Euro's introduction in 1999, the performance differential became significantly more pronounced, with the MV strategy achieving a Sharpe ratio of 1.922 compared to 1.449 for the naive strategy during 1999-2009. The Sharpe ratio of the difference is 1.380 (t-statistics=1.793). This outperformance persists into the post-2010 period, maintaining a significant Sharpe ratio differential of 0.428 (t-statistics=1.820).

The momentum portfolio demonstrates more modest but still notable improvements using the MV strategy, particularly during the 1999-2009 period. Finally, the value portfolio also benefits from optimisation, showing its strongest performance during the 1999-2009 period. However, unlike the carry factor, both the momentum and value portfolios exhibit less consistent outperformance across different sub-periods. To sum up, the performance reveals that MV factors outperformed naive approaches across three key economic periods. The most substantial outperformance occurred during 1999-2009, with the carry factor showing the most robust and sustained results. In the next section, we further explore the source of the outperformance of the MV optimised factors.

Figure 1: Cumulative performance of the currency factor portfolios

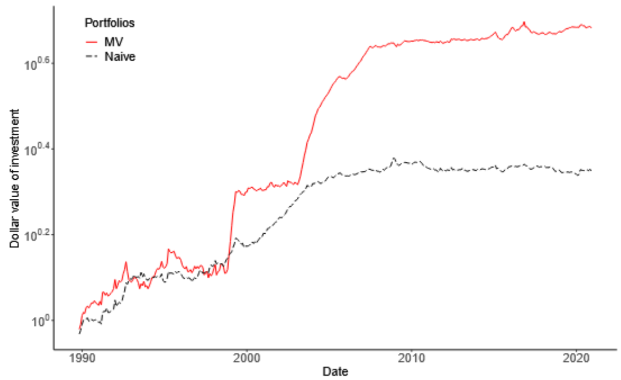

Carry

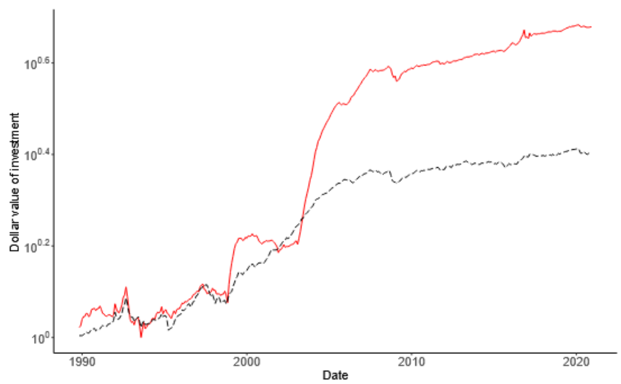

Momentum

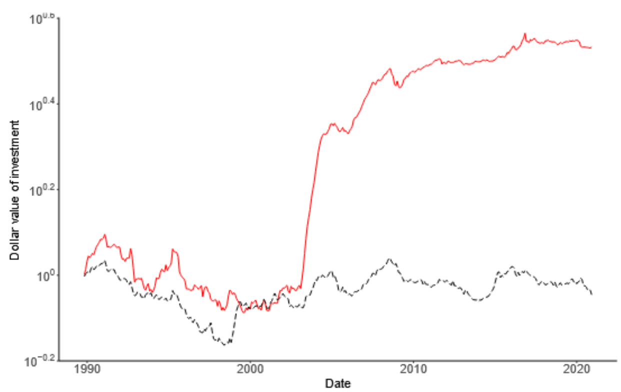

Value

These plots exhibit the cumulative performance of the MV optimised and the naive factor portfolios. The dollar value of investment (y-axis) is logarithmically scaled, given the considerable difference between the two weighting schemes. MV and Naive refer to the mean-variance diversified factor and the original equally weighted factors, respectively.

Table 2: Sub-period out-of-sample performance summary

	Carry			Momentum			Value		
	Naive	MV	Diff	Naive	MV	Diff	Naive	MV	Diff
Panel A: Pre 1999									
Mean	4.42%**	5.12%	0.70%	2.59%	1.58%	-1.02%	-2.53%*	-1.19%	1.34%
t-value	(2.107)	(1.454)	(0.258)	(1.37)	(0.46)	(-0.34)	(-1.732)	(-0.34)	(0.506)
Vol	0.066	0.110	0.084	0.059	0.107	0.093	0.046	0.109	0.082
Sharpe	0.675**	0.466	0.082	0.439	0.147	-0.109	-0.555*	-0.109	0.162***
SR.t	(2.133)	(1.472)	(0.278)	(1.342)	(0.462)	(0.241)	(-1.725)	(-0.337)	(2.228)
Omega	1.704	1.432	1.075	1.407	1.137	0.908	0.668	0.912	1.158
Dowd	0.151	0.103	0.016	0.077	0.031	-0.020	-0.089	-0.026	0.041
CER	0.003	0.002	-0.001	0.001	-0.001	-0.003	-0.003	-0.003	0.000
Panel B: 1999 - 2009									
Mean	8.02%***	23.66%***	15.64%***	1.90%	9.10%*	7.20%*	9.36%***	21.60%***	12.24%***
t-value	(4.805)	(6.374)	(4.579)	(1.300)	(1.923)	(1.770)	(6.377)	(5.028)	(3.410)
Vol	0.055	0.123	0.113	0.049	0.157	0.135	0.049	0.143	0.119
Sharpe	1.449***	1.922***	1.380*	0.392	0.580*	0.534	1.923***	1.516***	1.028
SR.t	(4.469)	(6.686)	(1.793)	(1.288)	(1.939)	(1.145)	(4.810)	(4.774)	(0.137)
Omega	3.062	5.012	2.921	1.335	1.532	1.479	3.864	3.341	2.272
Dowd	0.343	0.677	0.371	0.071	0.123	0.108	0.409	0.377	0.254
CER	0.006	0.017	0.010	0.001	0.002	0.002	0.007	0.014	0.007
Panel C: Post 2010									
Mean	-0.79%	3.13%	3.91%	2.28%*	7.35%**	5.07%*	-0.57%	0.57%	1.14%
t-value	(-0.531)	(0.961)	(1.414)	(1.735)	(2.414)	(1.928)	(-0.469)	(0.174)	(0.400)
Vol	0.049	0.108	0.091	0.043	0.101	0.087	0.040	0.108	0.094
Sharpe	-0.161	0.291	0.428*	0.525*	0.731**	0.583	-0.142	0.053	0.121
SR.t	(-0.527)	(0.947)	(1.820)	(1.707)	(2.257)	(1.105)	(-0.47)	(0.173)	(1.129)
Omega	0.879	1.275	1.428	1.506	1.843	1.602	0.899	1.045	1.116
Dowd	-0.032	0.052	0.074	0.099	0.146	0.106	-0.024	0.009	0.019
CER	-0.001	0.000	0.002	0.002	0.004	0.003	-0.001	-0.002	-0.001

This table summarises the out-of-sample performance of the naive and mean-variance factor portfolios with transaction costs over three sub-periods. Panel A presents the performance summary before January 1999, before the Euro was launched. Panel B covers the pre-crisis and the crisis period between January 1999 and December 2009. Panel C reports the performance after the financial crisis. Diff refers to the series of differences between naive and mean-variance portfolio returns. *Mean*, *t-value*, *Vol*, and *Sharpe* denote the annualised average returns, standard deviation, t-statistic of mean, and Sharpe ratio of portfolio returns. SR.t denotes the t-value of the Sharpe ratio based on bootstrapped Sharpe ratio equality test Ledoit and Wolf (2008). *Omega, Dowd,* and *CER* refer to the Omega ratio, Dowd ratio, and Certainty Equivalent Return, respectively. *, **, *** represent that the t-values are statistically significant at 10%, 5% and 1% level.

4. The source of the outperformance

4.1. A decomposition model

To further understand the source of the outperformance of the MV optimised factors, we adopt the expected return model of Lo and MacKinlay (1990) and Lewellen (2002), where the expected return function of a univariate-sorted portfolio across N currencies as:

$$E[\pi_t^p] = \frac{1}{N}\sum_{i=1}^{N} cov(r_{i,-t}^{signal}, r_{i,t}) - cov(\overline{r_{-t}^{signal}}, \overline{r}_t) + \frac{1}{N}\sum_{i=1}^{N}(\mu^i - \overline{\mu})^2, \tag{13}$$

where $r_{i,t}$ represents the excess return of currency i in month t, $r_{i,-t}^{signal}$ denotes the signal return of currency i over the formation period, $\overline{r_{-t}^{signal}}$ is the cross-sectional average of all currencies' signal returns over the formation period, and μ^i represents the unconditional expected return of currency i.[7] This model demonstrates that the expected returns of a currency factor portfolio can be decomposed into three components:

- 1) $\frac{1}{N}\sum_{i=1}^{N} cov(r_{i,-t}^{signal}, r_{i,t})$ reflects the correlation between currency excess returns and the signal returns over the formation period, i.e. the time series predictability.
- 2) $-cov(\overline{r_{-t}^{signal}}, \overline{r}_t)$ represents the negative cross-covariances where a past high currency signal is associated with low returns on other currencies.
- 3) $\frac{1}{N}\sum_{i=1}^{N}(\mu^i - \overline{\mu})$ refers to the cross-sectional variance of mean returns, i.e. currencies generate persistently high or low excess returns.

This decomposition reflects that the expected returns of a cross-sectional factor portfolio are sourced from both the time series predictability of the past signals, i.e., Component 1), and the cross-sectional predictability, i.e., Component 2) and 3).

We start by investigating the time series component and hypothesise that the MV weights jointly improve the time series predictability of currency characteristics. To test this hypoth-

[7] We simplify Equation 13 by adopting the assumption of Maurer et al. (2022) that the expected signal return is equal to the expected return of currency. The last term should be as $\frac{1}{N}\sum_{i=1}^{N}(\mu^{signal} * \mu^i) - \overline{\mu^{signal}} * \overline{\mu}$, where μ^{signal} denotes the unconditional expected signal return of currency i. We adopt this accurate measurement to conduct the empirical results.

esis, we further apply pooled panel regressions:

$$r_{i,t} = \alpha + \beta_1 r_{i,-t}^{signal} + \beta_2 w_{i,t-1}^{signal} + \beta_3 r_{i,-t}^{signal} * w_{i,t-1}^{signal} + \epsilon_t, \quad (14)$$

where $w_{i,t-1}^{signal}$ denotes the weight of currency i determined at the end of month $t-1$. Other notations are as described in Equation 13.

Table 3: time series predictability of the currency weights

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Carry	0.829***		0.834***						
	(20.65)		(17.42)						
W^{carry}		0.009	0.003						
		(1.67)	(1.24)						
Carry*W^{carry}			-0.209						
			(-1.04)						
Mom				0.348***		0.340***			
				(56.45)		(72.96)			
W^{mom}					0.006	-0.003***			
					(1.16)	(-5.12)			
Mom*W^{mom}						0.004			
						(0.40)			
Value							-0.027***		-0.029***
							(-4.89)		(-4.39)
W^{value}								-0.002	0.002
								(-1.31)	(1.10)
Value*W^{value}									-0.008
									(-0.84)
Constant	-0.001	0.001	-0.001	0.021***	0.001	0.026***	-0.006*	0.001	-0.006*
	(-0.02)	(0.17)	(-0.17)	(6.57)	(0.13)	(6.98)	(-1.87)	(0.20)	(-1.83)
Currency fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of currency	48	48	48	48	48	48	48	48	48
Adj. R^2	0.404	0.367	0.452	0.643	0.366	0.647	0.374	0.365	0.380

This table reports the results of panel regressions testing the joint predictability between currency characteristics and the mean-variance portfolio weights. W^{carry}, W^{mom}, and W^{value} represent the weights of the optimised carry, momentum, and value portfolios. Models (1) to (3) show the predictability of the currency carry. Models (4) to (6) show the predictability based on the currency's past returns. Models (7) to (9) show the predictability based on the currency value. The t-value (in parentheses) uses heteroskedasticity- and autocorrelation-consistent (HAC) standard errors, as described by Newey and West (1987). The number of observations and the adjusted R^2 are also reported. *, **, *** represent that the t-values are statistically significant at 10%, 5% and 1% levels, respectively.

Table 3 presents the regression results of Equation 14 to examine the time series predictability of currency weights for three factors. Interestingly, the optimised currency weights generally do not enhance the time series predictability of the currency characteristics. In the case of the carry factor, column 3 shows that the interaction term, $Carry * W^{carry}$, does not yield a statistically significant coefficient at the 10% level. This suggests that the interaction between the carry factor and its optimised weights does not provide any additional predictive power for future currency returns. The insignificance of this interaction term implies that the optimised carry weights cannot effectively capture time-varying predictability in currency returns, even though the carry factor itself reports highly significant coefficients (0.829, $t = 20.65$).

A similar pattern is observed with the momentum and value factors. In columns 6 and 9, the interaction terms are also insignificant, with a coefficients of 0.004 ($t = 0.40$) and 0.008 ($t = 0.84$). These insignificant results further suggest that the weights associated with the momentum and value factor do not substantially contribute to time series predictability. Although these two currency characteristics are highly significant across specifications, as shown in columns 4 and 7, the optimised weights fail to provide additional predictive insights when interacting with the characteristics. Therefore, we reject the hypothesis that the MV weights enhance the time series predictability of currency characteristics.

In the second step, we investigate which cross-sectional components contribute the most. The decomposition of strategy returns reveals substantial heterogeneity in the impact of optimisation across components. Table 4 presents the performance of three components for the three tested factors under naive and optimised portfolio constructions. For component 1), the differences are statistically insignificant across three factors. The results echo our findings from Table 3 that the optimised weighting scheme does not enhance the time series predictability of currency characteristics.

Of particular note is component 2), which we interpret as the negative cross-covariances. This component exhibits the most statistically significant differences between naïve and optimised approaches. For the carry strategy, the optimisation process leads to a striking improvement in component 2), from 2.15% to 9.08% with a difference at 6.86% (t-statistic = 2.90). The difference of 5.02 percentage points is statistically significant at the 5% level (t-statistic = 2.13). Similarly, the momentum factor experiences a substantial improvement in this component, with the optimised approach yielding a return of 8.69% (t-statistic = 2.92), compared to 4.80% in the naïve implementation. The resulting 3.89% difference is significant at the 10% level (t-statistic = 1.73). The value strategy also shows improvement in component 2), but the difference is insignificant.

Table 4: Performance of decomposed components

Component	Naïve			Optimised			Difference		
	Carry	MOM	Value	Carry	MOM	Value	Carry	MOM	Value
1)	1.41%***	-0.34%	0.27%	0.29%	0.41%	-0.53%	-1.24%	0.64%	-0.89%
	(2.60)	(-0.64)	(0.47)	(0.20)	(0.33)	(-0.43)	(-0.78)	(0.47)	(-0.67)
2)	2.15%*	4.80%***	-0.80%	9.08%***	8.69%***	2.25%	6.86%***	3.89%*	3.06%
	(1.91)	(4.59)	(-0.79)	(3.62)	(3.52)	(0.91)	(2.90)	(1.73)	(1.35)
3)	0.07%***	0.16%***	-0.01%	0.14%***	0.19%***	0.06%**	0.07%***	0.03%	0.05%
	(10.42)	(2.89)	(-0.04)	(6.60)	(8.51)	(2.09)	(3.70)	(0.58)	(0.20)

This table reports the factor log returns per annum created by each component. The last three columns present the differences in log returns between the naive and optimised components. Components 1), 2), and 3) represent the correlation between currency excess returns and the signal returns, the negative cross-covariances, and the unconditional expected returns, respectively. *, **, *** represent that the t-statistics are statistically significant at 10%, 5% and 1% levels, respectively.

Therefore, the outperformance of the optimised factors primarily operates through the channel of negative cross-covariances, i.e. component 2), rather than its time series predictability. The optimal weighting schemes exploit these cross-sectional properties of currency excess returns, allocating relatively low (high) weights to currencies according to the dynamic of the cross-covariances. However, the efficacy of this allocation strategy may diminish when return differentials are less pronounced, i.e. when currency returns are highly correlated. We discuss this in the next sub-section.

4.2. International diversification matters

Our previous finding suggests that utilising covariances matrix information largely improves the predictability of currency factors. Empirically, however, the effectiveness of MV optimisation is highly related to the actual co-movement across different currencies. To further examine this relationship, we divide our sample into two sub-samples, currencies from developed and emerging economies, based on the definition provided by the International Monetary Fund (IMF). 29 out of 48 are considered as developed currencies, including those of Australia, Austria, Belgium, Canada, Cyprus, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong, Iceland, Ireland, Israel, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Singapore, South Korea, Spain, Sweden, Swiss, Taiwan, and the UK, with the remainder designated as emerging currencies.[8]

The inconsistency in the outperformance of the MV optimised factors observed in Figure 1 highlights a regime change in the late 90s/early 2000s. This change corresponds to the launch of the Euro in 1999, when the Euro replaced ten European currencies, including those of Austria, Belgium, Germany, Finland, France, Ireland, Italy, Netherlands, Portugal, and Spain. Furthermore, in the early 2000s, the increase in the number of emerging currencies lowered the proportion of developed currencies in our sample. This creates a better investment environment for the purpose of international diversification. Figure 2 plots the number of available currencies (solid red line) and developed currencies (dashed blue line) during our sample period. Prior to the launch of the Euro, the sample consisted mainly of developed currencies (16 out of 18 currencies), suggesting insufficient international diversification for portfolio construction.

To formally assess the effect of international diversification, we analyse the time-varying average correlation across different groups of currencies. Following Pollet and Wilson (2010), the sample correlation for currencies j and k, $\rho_{jk,t}$ is measured as:

$$\rho_{jk,t} = \frac{\sigma_{jk,t}}{\sigma_{j,t}\sigma_{k,t}}, \tag{15}$$

where $\sigma_{jk,t}$ represents the covariance for currencies j and k, $\sigma_{j,t}/\sigma_{k,t}$ denote the standard deviation of currency j/k. Then, the average correlation is measured as follows:

$$AC_t = \sum_{j=1}^{N} \sum_{k \neq j} w_{j,t} w_{k,t} \rho_{jk,t}, \tag{16}$$

where N currencies are available.

The time-varying average correlation shown in Figure 3 examines the consistency of currencies' comovement for both the developed and the emerging samples. Results suggest

[8]The IMF designates Croatia, Slovakia, and Slovenia as advanced economies in 2023, 2009, and 2007, respectively. Their currencies are excluded from the developed sub-sample before they attain advanced economic status.

Figure 2: Number of available currencies

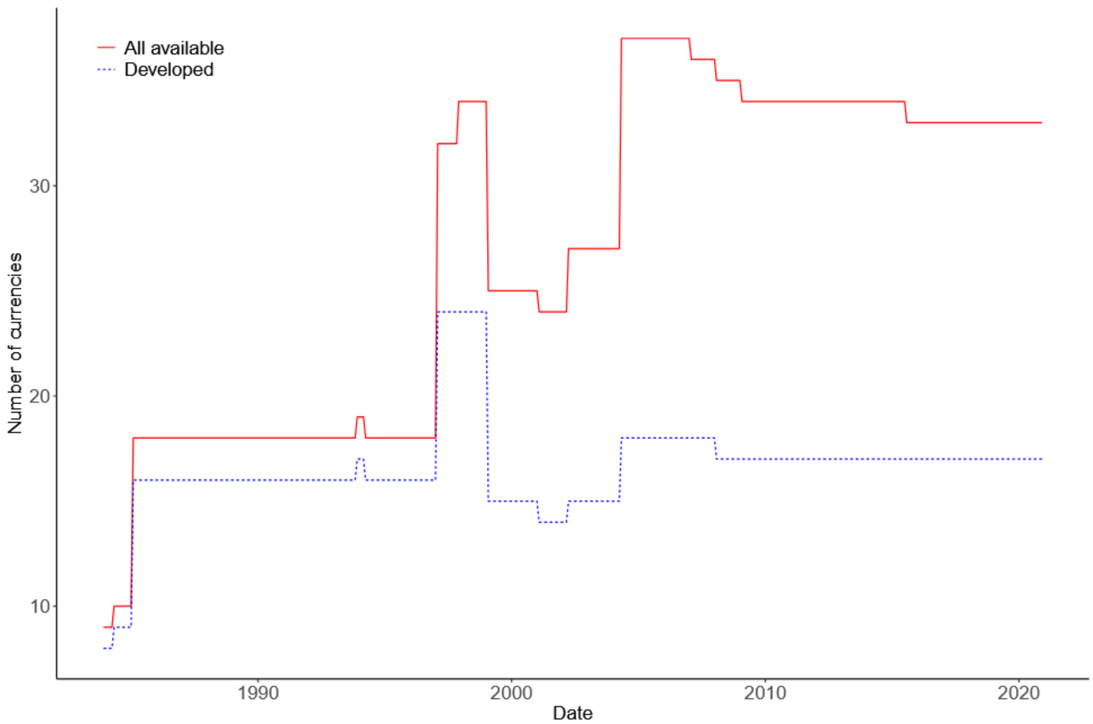

This plot exhibits the number of currencies available to investors. The blue dashed line counts the number of currencies from the developed economies across the sample period, with the solid red line counting the total number of currencies available (both developed and emerging market currencies).

Figure 3: Average pairwise correlation

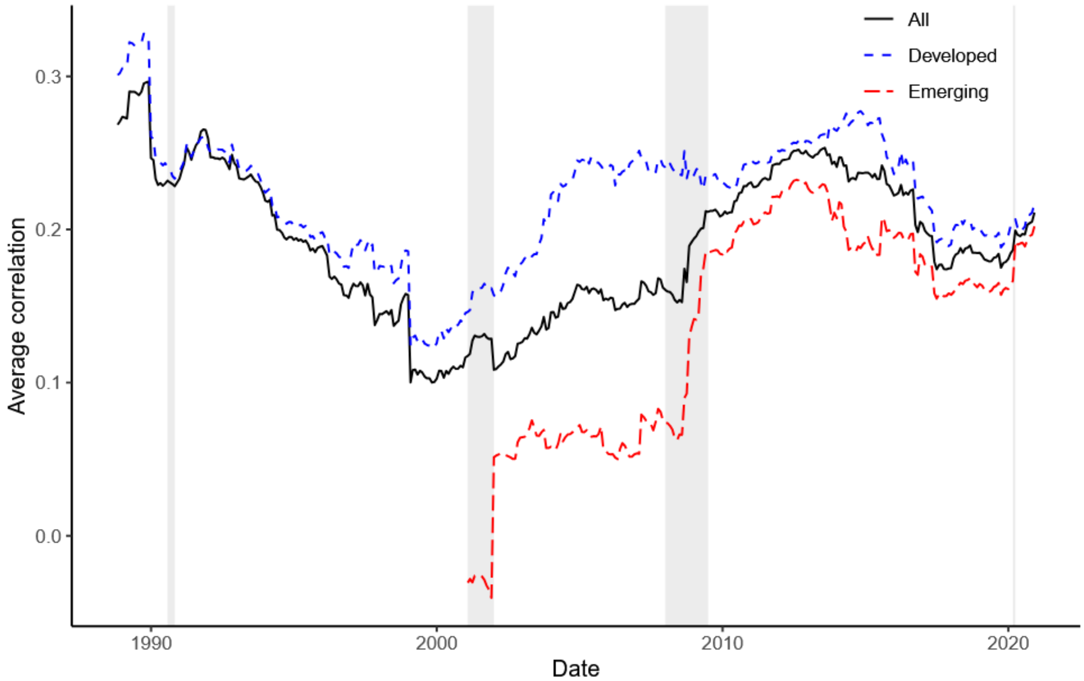

This plot exhibits the average pairwise correlation based on the three samples over time. The sample period is April 1989 to November 2020. The average correlation is calculated as the equally weighted average of pairwise correlations of monthly currency excess returns over the 60 months formation period. For the emerging sample, we omit the period before 2001 as it consists of only two currencies. The shaded ranges represent the recession periods identified by the NBER.

that the degree of comovement in the emerging currencies is much weaker than that of the developed currencies before 2010. These represent periods where MV optimisation works best. After 2010, the average pairwise correlation for the emerging currencies rose, but it is still below that of the developed currencies. Moreover, the emerging currencies exhibit a lower correlation to USD compared to the developed currencies. The average beta for the emerging currencies against the dollar index over the full sample period is 0.704, significantly lower than the value of 0.933 for thedeveloped currencies. Therefore, we argue that the reason the MV optimiser better estimates currency returns is primarily due to the additional information gleaned from the variance-covariance matrix. More specifically, emerging currencies play a pivotal role in the optimisation as they have a lower correlation with USD.

4.3. Does the distribution of abnormal profits matter?

To further explore the statistical reason for the superiority of the MV factor portfolios, we link the outperformance of the MV factors to the distribution of abnormal returns of currencies, using a framework proposed by DeMiguel et al. (2009) and Platanakis et al. (2021). This approach provides an attribution analysis of the outperformance by examining abnormal returns through established asset pricing models. It allows us to identify the key contributors to excess returns, offering insights into the underlying dynamics of currency market inefficiencies and the potential for strategically exploiting these anomalies.

Specifically, we adopt a currency market version of the factor model proposed by DeMiguel et al. (2009) and Platanakis et al. (2021) and estimate the abnormal returns for each currency as:

$$rx_{k,t} = \alpha + \sum \beta \times r_{factor,t} + \epsilon_t, \tag{17}$$

where $rx_{k,t}$ is the excess return of currency k in month t; α represents the abnormal return; β refers to the factor loading; r_{factor} refers to the excess returns of the control factors, including the dollar (the US dollar index, DXY), carry, momentum, and value factors; and ϵ_t is the noise term distributed as $\epsilon_t \sim N(0, \sigma_\epsilon)$.

Figure 4 plots the distribution of α for both the developed and emerging currencies over the sample period. The means of α in both the developed (dashed red line) and emerging (dashed blue line) samples exceed zero, with the mean in the emerging sample being slightly higher than in the developed sample. Due to the visible fat tails, the standard deviation of α in the emerging sample is higher than in the developed sample. Furthermore, the results of an F-test indicate that the variance of α in both samples is significantly different from zero at the 1% level. Previous literature, e.g., DeMiguel et al. (2009), showed that the optimal portfolios fail to significantly outperform the 1/N rule based on a simulated return series with the zero-alpha assumption. However, as argued by Jarrow (2010), while the zero-alpha assumption may be useful for academic study it is too rigorous empirically. Platanakis et al. (2021) reveal that the performance of optimal portfolios is subject to the standard deviation of α across the available assets. Therefore, they fit α to a normal distribution with a mean of zero and a standard deviation of 30 basis points, finding that the optimal portfolios significantly outperform the naive ones. We empirically test the abnormal returns for each currency and find that most of their abnormal returns are insignificantly different from zero. These results are reported in Appendix C. Motivated by this body of literature and our empirical observations in Figure 4, we now explore if the standard deviation of α affects the performance of the MV factor portfolios.

Figure 4: Distribution of abnormal returns (α) for developed and emerging currencies

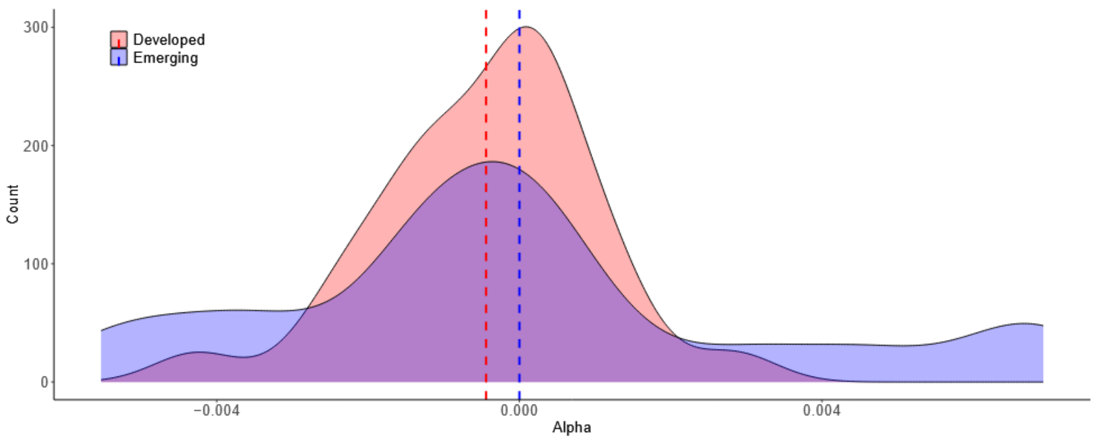

This figure shows the density-fitted areas of distributions of abnormal returns. The dashed lines indicate the mean of each distribution. The red area represents the distribution of abnormal returns for the developed sub-sample, while the blue area denotes the distribution based on the emerging sub-sample.

We hypothesise that the enhanced performance of the MV factor portfolios is derived from an increase in the standard deviation of currencies' abnormal profits. To validate our inference that the performance of the MV factor portfolios is positively related to the standard deviation of α, we first estimate the time-varying standard deviation of α over the entire sample period. In line with our estimation window in Section 2, we implement a 60-month rolling window to estimate the monthly α. Figure 5 plots the time-dynamic standard deviations of α for the full (solid red line), developed (green dotted line), and emerging (blue dashed line) samples. Before 1999, the blue line was even lower than the red and green lines, indicating a lower standard deviation of α for emerging currencies. Then, the standard derivation of α for the emerging sample spiked around 1999, with the blue line moving far above the green and red lines thereafter. The standard deviation of α in the emerging sample persistently exceeds that of the developed sample during 2000-2010.

Figure 5: Time-varying standard deviations of abnormal returns (α) distribution

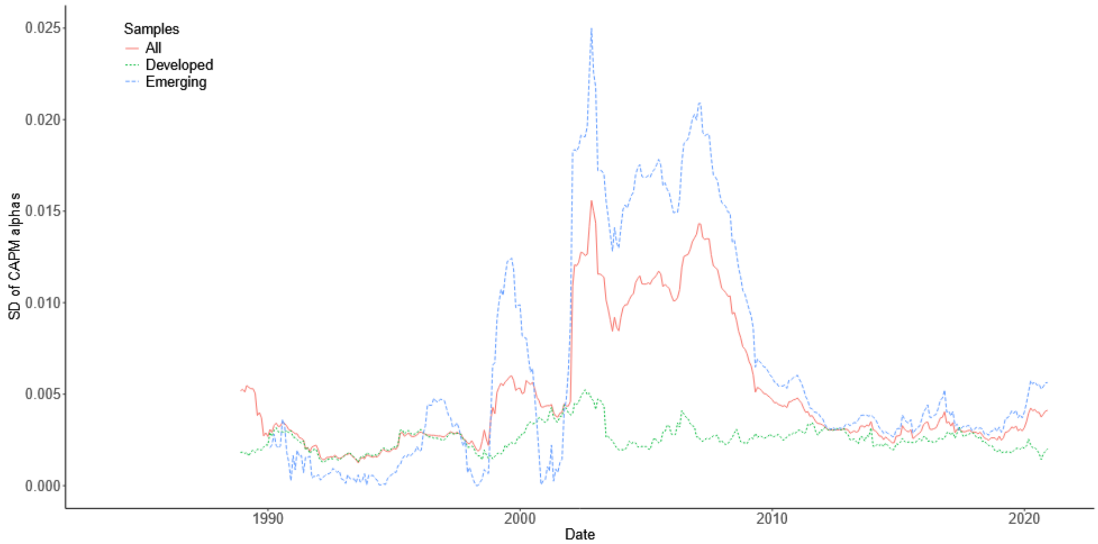

This plot displays the time-dynamic standard deviations of abnormal returns from April 1989 to November 2020 based on the three samples. The red solid line represents the standard deviation of α for the full sample, while the blue and green lines refer to the standard deviations of α for the developed and emerging samples, respectively.

Next, to quantify the relationship between the standard deviation of abnormal returns and the outperformance of the optimised factor portfolios, we regress the outperformance of the optimised factor portfolio on the lagged time-varying standard deviation of abnormal returns, given as:

$$r_{t+1}^{diff} = \alpha + \beta * SD_{t,p} + \epsilon, \tag{18}$$

where r_{t+1}^{diff} represents the differences between the returns of the MV factor portfolio and those of the naive factor in month $t+1$, and $SD_{t,p}$ refers to the standard deviation of currency abnormal returns, α from Equation 17, over the 60-month estimation window based on sample p in month t. We take a month lag between the SD of abnormal returns and r_{t+1}^{diff} as the asset weights in month $t+1$ are determined by the information in month t.

Table 5 reports the beta coefficients and corresponding t-statistics for each optimised factor in terms of Equation 18. Across the full sample period, the correlation coefficients are positive and significantly different from zero at the 1% level for the carry and momentum factors but insignificant for the value factor. This means that the superiority of the MV factor portfolios in month $t+1$ is positively related to the standard deviation of currencies'

Table 5: Relationship between the outperformance of the MV factor portfolios and the standard deviation of abnormal returns

Samples		Carry	Mom	Value
All	β	1.62***	1.56)***	1.08
	t	(2.57)	(2.57	(1.27)
Developed	β	-4.38*	-3.6	-7.59*
	t	(-1.68)	(-1.29)	(-2.09)
Emerging	β	1.09***	1.03***	0.82*
	t	(2.85)	(2.78)	(1.66)

This table reports the regression results of Equation 18. For each factor, we report the correlation coefficients β and the corresponding t-values. The t-value is measured by the heteroskedasticity- and autocorrelation-consistent (HAC) standard errors described by Newey and West (1987). '*', '**', '***' represent that the t-values of coefficients are statistically significant at the 10%, 5% and 1% levels, respectively.

abnormal returns, α, in month t. Further exploring the relationship in the two sub-samples, we find that the beta coefficients based on the developed sample are negative across all three factors, with those of carry and value factors being statistically significant at the 10% level. In contrast, for the emerging currency sample, beta coefficients are all positive and significant for the three factors. These results explain our above-mentioned finding that there is no MV portfolio outperformance when the sample mainly consists of developed currencies.

Statistical evidence from the above results suggests that the outperformance of the MV factor portfolios is positively related to the standard deviation of α. This finding is also consistent with the argument that the emerging currencies provide international diversification benefits, as they display lower return movement and, therefore, are associated with higher alpha standard deviations.

5. Optimised currency factors in asset pricing

Introducing the MV optimised factors raises a new question about whether they can replace the extant currency factors to better explain the cross-section of currency excess returns. First, to investigate whether the returns of the optimised factors subsume those of naive factors, we run spanning tests to assess the relationship following Barillas and Shanken (2017) and Ehsani and Linnainmaa (2022). In Table 6, Models (1) to (3) employ the returns of the optimised factors as the dependent variables, whereas Models (4) to (6) use the returns of the naive factors as the dependent variables. We further adopt the US dollar index, and the economic and inflation momentum strategies of Dahlquist and Hasseltoft

(2020) as macroeconomic controls.

Table 6: Spanning test

	Dependent Variable					
	Optimised factors			Naïve factors		
Factor	Carry	Momentum	Value	Carry	Momentum	Value
Model	(1)	(2)	(3)	(4)	(5)	(6)
α	0.004***	0.003**	0.0003	0.001	0.002**	-0.001
	(2.53)	(2.19)	(0.16)	(1.38)	(2.09)	(-0.89)
Opt				0.306***	0.295***	0.222***
				(8.15)	(10.01)	(10.74)
DXY	0.024	-0.074	-0.118*	-0.003	0.007	0.026
	(0.36)	(-1.08)	(-1.71)	(-0.07)	(0.16)	(0.68)
Carry	0.923***	0.120	0.032		-0.059	0.096**
	(10.06)	(1.30)	(0.25)		(-1.00)	(2.45)
MOM	0.555***	1.296***	0.584***	-0.203***		-0.129**
	(4.17)	(8.55)	(3.23)	(-2.63)		(-2.24)
Value	1.039***	1.168***	1.689***	-0.124	-0.343***	
	(6.81)	(6.23)	(9.55)	(-1.56)	(-4.21)	
EconMOM	-0.013	-0.023	-0.095	0.017	-0.015	0.030
	(-0.17)	(-0.36)	(-1.09)	(0.44)	(-0.62)	(0.89)
InfMOM	0.055	0.068	0.173*	-0.028	0.014	-0.041
	(0.60)	(0.83)	(1.70)	(-0.66)	(0.49)	(-0.97)

This table presents the results of spanning tests for the optimised factors and the naive ones as the dependent variables, respectively. DXY represents the US dollar index. EconMOM and InfMOM denote the economic and inflation momentum factors of Dahlquist and Hasseltoft (2020), respectively. For each model, we report the factor loadings and the corresponding t-values. The t-value is measured by the heteroskedasticity- and autocorrelation-consistent (HAC) standard errors described by Newey and West (1987). '*', '**', '***' represent that the t-values of coefficients are statistically significant at the 10%, 5% and 1% levels, respectively. The sample period ranges from April 1989 to November 2020.

Results in Table 6 reveal that the optimised carry and momentum strategies generate significant abnormal returns (α of 0.004 and 0.003, with t-statistics of 2.53 and 2.19, respectively) after controlling for their naive counterparts and other risk factors. Conversely, when the naive carry factor is employed as the dependent variable, its α is statistically insignificant at 0.001 (t-statistic = 1.38), indicating that it can be fully explained by the optimised carry factor. The naive momentum factor maintains significant abnormal returns (0.002, t-statistic = 2.09) relative to its optimised counterpart, suggesting that neither factor can fully explain the other. These findings suggest that the optimised carry factor subsumes the naive carry factor, whereas the optimised momentum factor does not. Finally, the value factor demonstrates no significant α in either direction for both optimised (0.0003) and naive

(-0.001) settings. Therefore, the results of the spanning regression for the value factor do not make sense economically.

For a more in-depth examination of the MV optimised factors in multi-factor currency pricing models, we apply the methodology of Fama and French (2015) and construct 27 candidate portfolios to serve as proxies for the currency risk premiums. Table 7 presents the average annualised monthly excess returns for 3×3 double-sorted portfolios formed by sorting currencies based on any two of the three currency factors. The first currency characteristic is displayed in rows, and the second currency characteristic is shown in columns. For example, Panel A reports the Carry-Momentum 3×3 portfolio returns. Then, different rows represent low/medium/high portfolios sorted on carry signals, and different columns represent low/medium/high portfolios sorted on momentum signals.

In Panel A, it is observed that the momentum portfolios exhibit stronger performance in the high-carry group, with high-minus-low profits reaching 8.60%, compared to only 3.83% in the low-carry group. Similarly, Panels B and C of Table 7 report the results of independent sorts based on carry-value, and momentum-value, respectively. The carry and momentum portfolios yield profits across all value groups, highlighting the robustness and persistence of these two effects in the currency market. However, the value effect appears weak. Specifically, low-minus-high value portfolios fail to generate meaningful profits in low carry, low momentum, and high momentum groups. These results are consistent with the low profitability of the value factor observed in Table 1.

We next evaluate the ability of the optimised factors to explain average excess returns on the test portfolios, as reported in Table 8. In Panel A, we employ the dollar index (DXY) as the proxy for the dollar factor proposed by Lustig et al. (2011) to construct four- and five-factor models. For robustness, Panel B incorporates the inflation and economic momentum factors of Dahlquist and Hasseltoft (2020) to develop six- and seven-factor models. The benchmark model includes the most widely used factors—carry, momentum, and value—augmented by additional factors.

Under the null hypothesis of a well-specified asset pricing model, the regression intercepts of excess returns on factor returns should equal zero (Fama and French, 1992). A statistically significant intercept implies the existence of abnormal returns unexplained by the model's

Table 7: Average monthly excess returns for portfolios formed on Carry and Momentum, Carry and Value, Momentum and Value

	Annual average return			T-values		
	Low	Medium	High	Low	Medium	High
Panel A: Carry-Momentum portfolios						
Low	-3.38%	-0.45%	-0.45%	-1.75	-0.37	-0.35
Medium	0.65%	1.98%	2.27%	0.42	1.32	1.51
High	1.23%	1.97%	9.83%	0.53	1.20	4.58
Panel B: Carry-Value portfolios						
Low	-2.64%	-0.76%	-0.73%	-1.35	-0.57	-0.48
Medium	2.35%	1.81%	0.68%	1.37	1.19	0.46
High	5.22%	5.14%	3.82%	2.63	2.52	1.81
Panel C: Momentum-Value portfolios						
Low	-1.71%	0.97%	-1.12%	-0.69	0.59	-0.69
Medium	1.45%	1.68%	0.66%	0.93	1.22	0.43
High	4.43%	3.68%	5.19%	2.19	2.20	2.81

This table reports the annualised profits of the test portfolios based on an approach of Fama and French (2015). We employ three currency characteristics to form the portfolios. In Panel A, currencies are allocated to three groups (Low to High, present in rows) based on carry signals. Currencies are allocated independently to three Momentum groups (Low to High, present in columns), using the past three excess returns breakpoints. The intersections of the two sorts produce nine equaly weighted Carry-Momentum portfolios. The Carry-Value (Panel B) and Momentum-Value (Panel C) portfolios are formed in the same way. The t-value calculation uses heteroskedasticity- and autocorrelation-consistent (HAC) standard errors as described by Newey and West (1987). The sample period ranges from April 1989 to November 2020.

risk factors, suggesting incomplete spanning of the return-generating process. We apply the GRS statistic proposed by Gibbons et al. (1989) on combinations of naive and optimised factors to test this hypothesis.

Our primary focus is on the change in model performance when replacing the naive factors with their optimised counterparts. In Panel A, Model (2) incorporating all three optimised factors yields lower GRS statistics compared to the benchmark, Model (1), indicating an improvement in the model's explanatory power. To identify which optimised factor drives this enhancement, we sequentially replace each naive factor in the benchmark model with its optimised version. Replacing the naive momentum and value factors with their optimised counterparts increases the GRS statistics, suggesting no improvement in these cases. However, substituting the naive carry factor with the optimised carry significantly reduces the GRS statistic from 1.359 to 1.140.

We further replace two naive factors with their optimised counterparts to assess the impact on model performance. Models (6) and (8) demonstrate that substituting the naive carry and momentum factors or the naive momentum and value factors results in increased

GRS statistics, thereby strengthening the rejection of the null hypothesis of Fama and French (2015). In contrast, Model (7), which incorporates the naive momentum factor alongside the optimised carry and value factors, achieves the lowest GRS statistic among these three models, at only 1.128.

Furthermore, Model (9) extends Model (8) by incorporating the U.S. Dollar Index (DXY), the naïve momentum factor, and all three MV-optimised factors. Given that the naive and optimised momentum factors fail to explain each other in Table 6, it is evident that these two factors capture distinct currency characteristics that are not accounted for by each other or the control variables. Thus, it is rational to include both factors in an asset pricing model. Notably, the five-factor model achieves the lowest GRS statistic among all models. Compared to the benchmark Model (1), replacing and adding optimised factors in Model (9) reduces the GRS statistic from 1.359 to 1.077. The GRS statistics indicate that Model (9) provides the most robust explanation for the average returns on currency portfolios across various constructions.

Table 8 also shows two ratios introduced by Fama and French (2015) that estimate the proportion of the cross-section of expected returns left unexplained by competing models. The numerator of each measure reflects the dispersion of intercepts produced by a given model for a set of test portfolios, while the denominator captures the dispersion of test portfolios' expected returns. Let R_i represent the time series average excess returns for portfolio i, $\bar{R}$ denote the cross-sectional average of R_i, and r_i be portfolio i's deviation from the cross-sectional average, defined as $r_i = R_i - \bar{R}$.

The first estimate is expressed as $\frac{A(|\alpha_i|)}{A(|r_i|)}$, which is the average absolute value of the intercepts, $|\alpha_i|$, divided by the average absolute value of r_i. This metric evaluates the proportion of the dispersion in intercepts relative to the dispersion in portfolio deviations from the cross-sectional average return.

The results for $\frac{A(|\alpha_i|)}{A(|r_i|)}$ in Table 8 show that, for various test portfolios, the benchmark model's absolute intercept, $A(|\alpha_i|)$, accounts for 47.5% of $A(|r_i|)$. This finding indicates that the benchmark model fails to explain 47.5% of the dispersion in average excess returns, as measured in units of return. When replacing one or two naive factors with their optimised counterparts, the proportion of unexplained dispersion in average excess returns ranges from

42.0% to 63.1%. More importantly, the five-factor model, Model (9), achieves a remarkable improvement by reducing the ratio by more than 25% compared to the benchmark model, lowering it to 0.354. This result implies that only 35.4% of the average excess returns remain unexplained, representing the best performance across all models.

Table 8: Summary statistics for multi-factor model test

	Model	GRS	p-value	$A\|\alpha\|$	$\frac{A\|\alpha_i\|}{A\|\bar{r_i}\|}$	$\frac{A(\hat{\alpha}_i^2)}{A(\hat{\mu}_i^2)}$	A(R^2)
Panel A: dollar index							
(1)	Carry+MOM+Value	1.359	0.113	0.083	0.475	0.218	0.600
(2)	Opt.Carry+Opt.MOM+ Opt.Value	1.207	0.222	0.084	0.480	0.276	0.541
(3)	Opt.Carry+MOM+Value	1.140	0.291	0.083	0.474	0.220	0.592
(4)	Carry+Opt.MOM+Value	1.422	0.082	0.094	0.539	0.316	0.568
(5)	Carry+MOM+Opt.Value	1.396	0.094	0.081	0.465	0.207	0.579
(6)	Opt.Carry+Opt.MOM+Value	1.455	0.070	0.110	0.631	0.391	0.563
(7)	Opt.Carry+MOM+Opt.Value	1.128	0.304	0.073	0.420	0.175	0.573
(8)	Carry+Opt.MOM+Opt.Value	1.397	0.094	0.078	0.446	0.268	0.548
(9)	Opt.Carry+MOM+Opt.MOM+Opt.Value	1.077	0.364	0.062	0.354	0.145	0.579
Panel B: dollar index with the economic and inflation momemtum							
(1)	Carry+MOM+Value	1.347	0.120	0.083	0.474	0.217	0.602
(2)	Opt.Carry+Opt.MOM+ Opt.Value	1.189	0.240	0.082	0.469	0.270	0.544
(3)	Opt.Carry+MOM+Value	1.133	0.298	0.083	0.474	0.220	0.594
(4)	Carry+Opt.MOM+Value	1.413	0.086	0.094	0.538	0.315	0.570
(5)	Carry+MOM+Opt.Value	1.385	0.099	0.081	0.465	0.206	0.581
(6)	Opt.Carry+Opt.MOM+Value	1.447	0.073	0.110	0.630	0.391	0.565
(7)	Opt.Carry+MOM+Opt.Value	1.118	0.315	0.073	0.420	0.174	0.575
(8)	Carry+Opt.MOM+Opt.Value	1.379	0.102	0.077	0.439	0.263	0.550
(9)	Opt.Carry+MOM+Opt.MOM+Opt.Value	1.066	0.379	0.061	0.351	0.144	0.581

This table evaluates the performance of multi-factor models in explaining the monthly excess returns of the 27 test portfolios detailed in Table 7. Panel A presents results where the factors augment the Dollar Index (DXY), while Panel B extends the analysis by including DXY along with economic and inflation momentum factors from Dahlquist and Hasseltoft (2020). The GRS statistic tests whether the expected values of the 27 intercept estimates are jointly zero, with the corresponding p-values reported. Key measures include $A|\alpha|$, $\frac{A|\alpha|}{A|\bar{r}|}$, $\frac{A(\alpha^2)}{A(\bar{r}^2)}$, and A($R^2$). Specifically, $A|\alpha|$ represents the average absolute value of the intercepts, $\frac{A|\alpha|}{A|\bar{r}|}$ captures the ratio of the average absolute intercept to the average absolute value of $\bar{r_i}$, and $\frac{A(\alpha^2)}{A(\bar{r}^2)}$ measures the ratio of the average squared intercept to the average squared value of $\bar{r_i}$. Here, $\bar{r_i}$ denotes the average return of portfolio i adjusted by subtracting the mean return across all test portfolios. A(R^2) refers to the average R^2 across all regressions.

The second estimate captures the extent of error inflation by measuring increases in both the average absolute intercept, $A(|\alpha_i|)$, and the average absolute deviation, $A(|r_i|)$. The estimated intercept, a_i, is the sum of the true intercept, α_i, and an estimation error, e_i, expressed as $a_i = \alpha_i + e_i$. Similarly, r_i, the deviation of portfolio i's return from the cross-sectional average, is the sum of μ_i, portfolio i's expected deviation from the grand mean, and an estimation error, ϵ_i, such that $r_i = \mu_i + \epsilon_i$. To account for the impact of measurement error, we focus on squared intercepts and squared deviations, which mitigate the inflation effects introduced by the errors and provide a more accurate adjustment. The cross-sectional average of μ_i is zero, so $A(\mu_i^2)$ represents the cross-sectional variance of expected portfolio returns, and the ratio $A(\alpha_i^2)/A(\mu_i^2)$ quantifies the proportion of $A(\mu_i^2)$ left unexplained by a model. Since α_i is a constant, the expected value of the squared estimated intercept is equal to the squared value of the true intercept plus the sampling variance of the estimate, $E(a_i^2) = \alpha_i^2 + E(e_i^2)$. Our estimate, $\hat{\alpha}_i^2$, of the squared true intercept, α_i^2, is calculated as the difference between the squared value of the estimated regression intercept and the square of its standard error. Similarly, our estimate of μ_i^2, $\hat{\mu}_i^2$, is derived as the difference between the squared realised deviation, r_i^2, and the square of its standard error. The ratio of averages, $A(\hat{\alpha}_i^2)/A(\hat{\mu}_i^2)$, then provides an estimate of the proportion of the variance of the expected returns of the test portfolio that remains unexplained by the model.

As highlighted in Fama and French (2015), compared to $A(|\alpha_i|)$, the ratio $A(\hat{\alpha}_i^2)/A(\hat{\mu}_i^2)$ is expressed in units of squared return and accounts for corrections due to sampling error. The estimates indicate that the benchmark model leaves 21.8% of the cross-sectional variance of expected returns unexplained. Replacing one naive factor with its optimised counterpart reduces this proportion slightly to 20.7%, while replacing two naive factors decreases it further to 17.5%. Notably, in our five-factor model, this ratio is reduced to just 0.145, indicating that only 14.5% of the cross-sectional variance of expected returns remains unexplained.

In conclusion, the results for $\frac{A(|\alpha_i|)}{A(|r_i|)}$ underscore the effectiveness of incorporating optimised factors into the model, as they significantly enhance its ability to explain the dispersion of average excess returns, thereby improving the model's overall explanatory power. Similarly, the findings for $A(\hat{\alpha}_i^2)/A(\hat{\mu}_i^2)$ further highlight the substantial gains in explanatory power achieved through the inclusion of optimised factors. These results demonstrate the superiority of the five-factor model in capturing cross-sectional variations in expected returns, reinforcing its robustness and effectiveness in asset pricing applications.

Panel B of 8 presents the robustness check after incorporating economic and inflation momentum factors of Dahlquist and Hasseltoft (2020) as additional macroeconomic level factors. The newly added factors change the five-factor model in Panel A to a seven-factor model in Panel B. The results substantiate the superior performance of the seven-factor model, Model (9), relative to alternative specifications. Specifically, when augmented with two additional macroeconomic control variables, Model (9) exhibits the most favourable pricing performance across multiple statistical metrics. The model yields the lowest GRS statistic of 1.066 (p-value $= 0.379$), indicating that we cannot reject the null hypothesis of zero alphas at conventional significance levels. Furthermore, Model (9) demonstrates the most substantial pricing efficiency as evidenced by both the lowest $\frac{A(|\alpha_i|)}{A(|r_i|)}$ ratio of 0.351 and the lowest $A(\hat{\alpha}_i^2)/A(\hat{\mu}_i^2)$ value of 0.144. These results provide robust evidence that using optimised factors instead of naive factors enhances the asset pricing model's ability to capture the cross-sectional variation in currency portfolio returns.

6. Conclusion

This study proposes the application of MV optimisation on extant currency factors. By adopting the basic mean-variance portfolio construction, we find compelling evidence to support the superiority of the MV optimised factors over the simplistic, naive diversification of currency factors. We find that this superiority is sourced from the extra covariance matrix information added to the currency characteristics and expected mean returns, leading to stronger cross-sectional predictability. We also find that international diversification plays an important role in forming currency portfolios.

Furthermore, our asset pricing test results suggest that these MV optimised currency factors capture the cross-sectional variation of currency returns better than using naive diversification of currency factors. The most efficient currency pricing models consist of all three MV optimised factors and the naive momentum factor.

The economic value of the MV optimised factors is notable, indicating that the profitability of currency factors, when incorporating portfolio optimisation, is substantially improved. Consequently, our findings provide valuable insights for investors by presenting a compelling methodology for portfolio allocation decisions. Our proposed MV-optimised factors operate without leverage and are designed with constant portfolio weights over time. Therefore, future studies may seek to investigate the potential benefits of incorporating notional values or volatility timing to further refine performance.

References

Ackermann, F., Pohl, W. and Schmedders, K. (2017), 'Optimal and naive diversification in currency markets', *Management Science* **63**(10), 3347–3360.

Akram, Q. F., Rime, D. and Sarno, L. (2008), 'Arbitrage in the foreign exchange market: Turning on the microscope', *Journal of International Economics* **76**(2), 237–253.

Ardia, D. and Boudt, K. (2015), 'Implied expected returns and the choice of a mean–variance efficient portfolio proxy', *The Journal of Portfolio Management* **41**(4), 68–81.

Asness, C. S., Moskowitz, T. J. and Pedersen, L. H. (2013), 'Value and momentum everywhere', *The Journal of Finance* **68**(3), 929–985.

Balduzzi, P. and Chiang, I.-H. E. (2020), 'Real exchange rates and currency risk premiums', *The Review of Asset Pricing Studies* **10**(1), 94–121.

Barillas, F. and Shanken, J. (2017), 'Which alpha?', *The Review of Financial Studies* **30**(4), 1316–1338.

Barroso, P. and Santa-Clara, P. (2015), 'Beyond the carry trade: Optimal currency portfolios', *Journal of Financial and Quantitative Analysis* **50**(5), 1037–1056.

Baz, J., Breedon, F., Naik, V. and Peress, J. (2001), 'Optimal portfolios of foreign currencies', *Journal of Portfolio Management* **28**(1), 102.

Chernov, M., Dahlquist, M. and Lochstoer, L. (2023), 'Pricing currency risks', *The Journal of Finance* **78**(2), 693–730.

Dahlquist, M. and Hasseltoft, H. (2020), 'Economic momentum and currency returns', *Journal of Financial Economics* **136**(1), 152–167.

Daniel, K., Hodrick, R. J., Lu, Z. et al. (2017), 'The carry trade: Risks and drawdowns', *Critical Finance Review* **6**(2), 211–262.

Della Corte, P., Sarno, L. and Tsiakas, I. (2009), 'An economic evaluation of empirical exchange rate models', *The Review of Financial Studies* **22**(9), 3491–3530.

DeMiguel, V., Garlappi, L. and Uppal, R. (2009), 'Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?', *The Review of Financial Studies* **22**(5), 1915–1953.

Ehsani, S. and Linnainmaa, J. T. (2022), 'Factor momentum and the momentum factor', *The Journal of Finance* **77**(3), 1877–1919.

Elton, E. J. and Gruber, M. J. (1997), 'Modern portfolio theory, 1950 to date', *Journal of Banking & Finance* **21**(11-12), 1743–1759.

Fama, E. F. (1984), 'Forward and spot exchange rates', *Journal of Monetary Economics* **14**(3), 319–338.

Fama, E. F. and French, K. R. (1992), 'The cross-section of expected stock returns', *The Journal of Finance* **47**(2), 427–465.

Fama, E. F. and French, K. R. (2015), 'A five-factor asset pricing model', *Journal of Financial Economics* **116**(1), 1–22.

Fan, M., Kearney, F., Li, Y. and Liu, J. (2024), 'Optimising currency factors', *Available at SSRN 4965896* .

Filippou, I., Gozluklu, A. E. and Taylor, M. P. (2018), 'Global political risk and currency momentum', *Journal of Financial and Quantitative Analysis* **53**(5), 2227–2259.

Gibbons, M. R., Ross, S. A. and Shanken, J. (1989), 'A test of the efficiency of a given portfolio', *Econometrica* pp. 1121–1152.

Goyal, A. and Jegadeesh, N. (2017), 'Cross-sectional and time-series tests of return predictability: What is the difference?', *The Review of Financial Studies* **31**(5), 1784–1824.

Greenaway-McGrevy, R., Mark, N. C., Sul, D. and Wu, J.-L. (2018), 'Identifying exchange rate common factors', *International Economic Review* **59**(4), 2193–2218.

Hsu, P.-H., Han, Q., Wu, W. and Cao, Z. (2018), 'Asset allocation strategies, data snooping, and the 1/n rule', *Journal of Banking & Finance* **97**, 257–269.

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

Jarrow, R. A. (2010), 'Active portfolio management and positive alphas: fact or fantasy?', *The Journal of Portfolio Management* **36**(4), 17–22.

Jegadeesh, N. and Titman, S. (1993), 'Returns to buying winners and selling losers: Implications for stock market efficiency', *The Journal of Finance* **48**(1), 65–91.

Keating, C. and Shadwick, W. F. (2002), 'A universal performance measure', *Journal of Performance Measurement* **6**(3), 59–84.

Kolm, P. N., Tütüncü, R. and Fabozzi, F. J. (2014), '60 years of portfolio optimization: Practical challenges and current trends', *European Journal of Operational Research* **234**(2), 356–371.

Kroencke, T. A., Schindler, F. and Schrimpf, A. (2014), 'International diversification benefits with foreign exchange investment styles', *Review of Finance* **18**(5), 1847–1883.

Ledoit, O. and Wolf, M. (2008), 'Robust performance hypothesis testing with the sharpe ratio', *Journal of Empirical Finance* **15**(5), 850–859.

Lewellen, J. (2002), 'Momentum and autocorrelation in stock returns', *The Review of Financial Studies* **15**(2), 533–564.

Lo, A. W. and MacKinlay, A. C. (1990), 'When are contrarian profits due to stock market overreaction?', *The Review of Financial Studies* **3**(2), 175–205.

Lustig, H., Roussanov, N. and Verdelhan, A. (2011), 'Common risk factors in currency markets', *The Review of Financial Studies* **24**(11), 3731–3777.

Markowitz, H. (1952), 'Portfolio selection', *The Journal of Finance* **7**(1), 77–91.

Maurer, T. A., Tô, T.-D. and Tran, N.-K. (2023), 'Market timing and predictability in fx markets', *Review of Finance* **27**(1), 223–246.

Maurer, T., To, T.-D. and Tran, N.-K. (2022), 'Pricing implications of covariances and spreads in currency markets', *The Review of Asset Pricing Studies* **12**(1), 336–388.

Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A. (2012*a*), 'Carry trades and global foreign exchange volatility', *The Journal of Finance* **67**(2), 681–718.

Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A. (2012*b*), 'Currency momentum strategies', *Journal of Financial Economics* **106**(3), 660–684.

Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A. (2017), 'Currency value', *The Review of Financial Studies* **30**(2), 416–441.

Moskowitz, T. J., Ooi, Y. H. and Pedersen, L. H. (2012), 'Time series momentum', *Journal of Financial Economics* **104**(2), 228–250.

Neely, C. J. and Weller, P. A. (2013), 'Lessons from the evolution of foreign exchange trading strategies', *Journal of Banking & Finance* **37**(10), 3783–3798.

Newey, W. K. and West, K. D. (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**(3), 703–708.

Opie, W. and Riddiough, S. J. (2020), 'Global currency hedging with common risk factors', *Journal of Financial Economics* **136**(3), 780–805.

Platanakis, E., Sutcliffe, C. and Ye, X. (2021), 'Horses for courses: Mean-variance for asset allocation and 1/n for stock selection', *European Journal of Operational Research* **288**(1), 302–317.

Pollet, J. M. and Wilson, M. (2010), 'Average correlation and stock market returns', *Journal of Financial Economics* **96**(3), 364–380.

Prigent, J.-L. (2007), *Portfolio optimization and performance analysis*, CRC Press.

Sharpe, W. F. (1966), 'Mutual fund performance', *The Journal of Business* **39**(1), 119–138.

Taylor, M. P. (1989), 'Covered interest arbitrage and market turbulence', *The Economic Journal* **99**(396), 376–391.

Timmermann, A. and Granger, C. W. (2004), 'Efficient market hypothesis and forecasting', *International Journal of Forecasting* **20**(1), 15–27.

Appendix A. Summary statistics

This section details our data sample and summary statistics. Table A.1 presents some descriptive statistics of our FX market data. Most currencies yield positive annualised returns over the sample period, with the *Indonesian Rupiah* reporting the highest annualised return at 13.86% per annum. Only 11 out of 48 currencies produce negative returns, with *Finnish Markka* displaying the lowest averaged return, at -7.06% per annum. The annualised standard deviation ranges from 0.003 (*Saudi Arabian* Riyal) to 0.272 (*Indonesian Rupiah*). Finally, the *Slovak Koruna* illustrates the highest Sharpe ratio, at 1.150 per annum; while the *Finnish Markka* shows the lowest value, at -1.078 per annum.

Table A.1: Summary statistics

Regions/Currency	Statistics					Sample period	
	Ann.Mean	*Ann.SD*	*Ann.SR*	*min*	*max*	*Start Date*	*End Date*
Australian Dollar	2.41%	0.117	0.205	-15.11%	10.88%	Jan-85	Nov-20
Austrian Schilling	-6.08%	0.101	-0.605	-6.69%	5.66%	Jan-97	Dec-98
Belgian Franc	3.82%	0.125	0.306	-10.34%	7.36%	Nov-83	Dec-98
Brazilian Real	4.51%	0.160	0.282	-15.64%	12.12%	Apr-04	Nov-20
Bulgarian Lev	-0.36%	0.095	-0.038	-9.14%	10.29%	Apr-04	Nov-20
Canadian Dollar	0.64%	0.074	0.086	-11.66%	9.25%	Jan-85	Nov-20
Croatian Kuna	0.56%	0.099	0.057	-9.76%	7.82%	Apr-04	Nov-20
Cyprus Pound	4.65%	0.074	0.630	-4.60%	4.42%	Apr-04	Dec-07
Czech Koruna	1.46%	0.125	0.117	-13.11%	10.10%	Jan-97	Nov-20
Danish Krone	2.08%	0.105	0.198	-10.31%	10.12%	Jan-85	Nov-20
Egyptian Pound	11.65%	0.152	0.766	-49.88%	16.59%	Apr-04	Nov-20
Euro	-0.58%	0.097	-0.060	-9.45%	9.89%	Jan-99	Nov-20
Finnish Markka	-7.06%	0.102	-0.691	-6.72%	5.63%	Jan-97	Dec-98
French Franc	4.04%	0.116	0.347	-9.99%	7.93%	Nov-83	Dec-98
German Mark	2.14%	0.120	0.179	-10.45%	7.19%	Nov-83	Dec-98
Greek Drachma	-4.76%	0.111	-0.427	-10.88%	7.19%	Jan-97	Dec-00
Hong Kong Dollar	-0.25%	0.006	-0.402	-1.34%	0.79%	Nov-83	Nov-20
Hungarian Forint	2.47%	0.138	0.178	-18.88%	11.10%	Nov-97	Nov-20
Icelandic Krona	1.53%	0.149	0.102	-26.71%	20.41%	Apr-04	Nov-20
Indian Rupee	1.57%	0.073	0.217	-8.43%	8.16%	Nov-97	Nov-20
Indonesian Rupiah	13.86%	0.272	0.509	-59.81%	34.85%	Jan-97	Nov-20
Irish Punt	1.96%	0.081	0.243	-5.98%	5.33%	Nov-93	Dec-98
Israeli Shekel	1.92%	0.080	0.240	-7.61%	7.30%	Apr-04	Nov-20
Italian Lira	4.13%	0.117	0.354	-14.11%	8.59%	Apr-84	Dec-98
Japanese Yen	-0.16%	0.109	-0.014	-10.29%	16.14%	Nov-83	Nov-20
Kuwaiti Dinar	0.50%	0.022	0.225	-6.23%	2.15%	Jan-97	Nov-20
Malaysian Ringgit	2.45%	0.158	0.155	-31.31%	12.57%	Jan-85	Nov-20
Mexican Peso	3.01%	0.119	0.253	-20.38%	12.24%	Jan-97	Nov-20
Netherlands Guilder	2.28%	0.120	0.190	-10.48%	7.55%	Nov-83	Dec-98
New Zealand Dollar	4.86%	0.125	0.391	-14.74%	13.55%	Jan-85	Nov-20
Norwegian Krone	1.89%	0.111	0.170	-11.71%	8.64%	Jan-85	Nov-20
Philippine Peso	1.17%	0.076	0.154	-11.99%	9.28%	Jan-97	Nov-20
Polish Zloty	2.69%	0.139	0.194	-16.41%	10.45%	Mar-02	Nov-20
Portuguese Escudo	-5.26%	0.097	-0.543	-6.51%	5.51%	Jan-97	Dec-98
Russian Ruble	0.36%	0.149	0.024	-16.45%	11.89%	Apr-04	Nov-20
Saudi Riyal	0.12%	0.003	0.407	-1.02%	0.52%	Jan-97	Nov-20
Singapore Dollar	0.31%	0.055	0.056	-8.58%	6.08%	Jan-85	Nov-20
Slovak Koruna	13.37%	0.116	1.150	-9.69%	10.55%	Mar-02	Dec-08
Slovenian Tolar	2.45%	0.081	0.302	-4.54%	4.47%	Apr-04	Dec-06
South Africa Rand	0.03%	0.157	0.002	-19.69%	13.92%	Nov-83	Nov-20
South Korean Won	1.70%	0.109	0.156	-13.72%	12.93%	Mar-02	Nov-20
Spanish Peseta	-4.89%	0.101	-0.483	-7.04%	5.70%	Jan-97	Dec-98
Swedish Krona	1.20%	0.112	0.107	-14.49%	11.06%	Jan-85	Nov-20
Swiss Franc	0.46%	0.114	0.040	-14.11%	12.65%	Nov-83	Nov-20
Taiwan Dollar	-1.20%	0.053	-0.226	-8.76%	5.72%	Jan-97	Nov-20
Thai Baht	1.24%	0.099	0.125	-19.96%	16.33%	Jan-97	Nov-20
Ukrainian Hryvnia	-3.43%	0.198	-0.174	-48.39%	12.77%	Apr-04	Jun-15
UK Pound	1.20%	0.103	0.117	-13.02%	13.74%	Nov-83	Nov-20

This table reports the summary statistics of monthly excess returns. *Ann.Mean* denotes annualised excess return. *Ann.SD* is annualised standard deviation of excess returns. *Ann.SR* refers to the Sharpe ratio. *max* and *min* represent the highest and lowest single-month excess return, respectively. *Start Date* and *End Date* are the start and end months of the sample period, respectively.

Appendix B. Performance metrics

We employ various performance metrics to evaluate the out-of-sample performance of both the naive and optimised factors. Initially, we calculate the mean ($\bar{\mu_P}$) and volatility (σ_P) of the excess returns of the out-of-sample (OOS) factors to serve as fundamental profitability and risk metrics.

Subsequently, we apply the widely recognised Sharpe ratio by Sharpe (1966) as a risk-adjusted performance measure:

$$SR_P = \frac{\bar{\mu_P}}{\sigma_P}. \tag{B.1}$$

Furthermore, aligning with Platanakis et al. (2021), we incorporate alternative out-of-sample performance metrics. First, the Certainty Equivalent Return (CER) for mean-variance investors is approximated and computed as:

$$CER = \bar{\mu_P} - \frac{\lambda \sigma^2}{2}, \tag{B.2}$$

where λ represents the relative risk aversion parameter, set at 0.25 as per Platanakis et al. (2021).

Next, the Omega ratio by Keating and Shadwick (2002), with a target return of zero, is computed as the average gain to the average loss ratio:

$$Omega = \frac{\sum_{t=1}^{\tau} max(0, R_{p,t})}{\sum_{t=1}^{T} max(0, -R_{p,t})}, \tag{B.3}$$

where τ signifies the sample size of the OOS observations, and $R_{p,t}$ denotes the OOS portfolio return at time t. The Omega ratio has the advantage of not relying on assumptions about the distribution of portfolio returns.

The Dowd ratio, introduced by Prigent (2007), is the out-of-sample mean excess portfolio return divided by the portfolio value-at-risk (VaR), which is measured as:

$$Dowd = \frac{\bar{\mu_P}}{VaR_{0.95}}, \tag{B.4}$$

with the VaR calculated at the 95% confidence level over the OOS period.

Appendix C. Regression results based on the factor model

Table C.1 summarises the coefficient correlations of asset pricing model for each currency over the entire sample period. Panel A contains the regression results of developed currencies, as defined in Section 4 and Panel B lists the results of emerging currencies.

We find that the average alpha in the developed sample, -0.044, is lower than in the emerging sample, 0.0001. Two alphas in Panel A are significantly different from zero at the 10% level, and three currencies are significantly different from zero at the 10% level in Panel B. The standard deviation of α in the emerging sample exceeds that of the developed sample, which supports our findings in Section 4. Furthermore, the regression results further show that σ_ϵ in the emerging sample, 0.025, is higher than in the developed sample, 0.015. This aligns with the proposition described by Platanakis et al. (2021) that increased idiosyncratic volatility can benefit optimal portfolios.

Table C.1: Currency abnormal profits

	α	$t(\alpha)$	β_{DXY}	$t(\beta_{DXY})$	σ_ϵ
Panel A: Developed					
Australian Dollar	-0.195	-1.321	-0.650	-5.895	0.028
Austrian Schilling	-0.269	-1.472	-1.237	-11.708	0.009
Belgian Franc	0.123	1.110	-1.152	-35.499	0.012
Canadian Dollar	-0.148	-1.818	-0.378	-5.720	0.018
Cyprus Pound	0.284	3.107	-0.722	-16.436	0.005
Czech Koruna	0.039	0.354	-1.348	-28.371	0.019
Danish Krone	0.043	0.843	-1.155	-52.758	0.009
Euro	-0.064	-1.508	-1.183	-57.674	0.007
Finnish Markka	-0.427	-1.954	-1.328	-13.524	0.009
French Franc	0.144	1.900	-1.139	-36.581	0.010
German Mark	0.021	0.293	-1.174	-41.296	0.009
Greek Drachma	0.014	0.053	-1.210	-7.302	0.018
Hong Kong Pound	-0.022	-1.754	-0.008	-1.959	0.002
Irish Punt	0.037	0.128	-0.796	-12.330	0.037
Israeli Shekel	-0.067	-0.598	-1.081	-11.450	0.013
Italian Lira	0.103	0.966	-0.562	-10.212	0.018
Icelandic Krona	0.021	0.140	-1.066	-23.913	0.017
Japanese Yen	0.095	0.676	-0.663	-9.598	0.026
Netherlands Guilder	0.027	0.369	-1.177	-41.803	0.010
New Zealand Dollar	-0.008	-0.054	-0.770	-6.722	0.028
Norwegian Krone	-0.122	-1.505	-1.088	-24.827	0.017
Portuguese Escudo	-0.227	-1.229	-1.202	-14.310	0.009
Singapore Dollar	-0.082	-1.410	-0.422	-13.185	0.012
South Korean Won	-0.116	-0.765	-0.836	-6.327	0.023
Spanish Peseta	-0.126	-0.942	-1.213	-13.872	0.009
Swedish Krona	-0.126	-1.493	-1.120	-31.652	0.016
Swiss Franc	0.005	0.067	-1.127	-26.916	0.016
Taiwan Dollar	-0.208	-2.335	-0.329	-9.865	0.013
UK Pound	-0.026	-0.271	-0.913	-17.643	0.019
Average	-0.044	-0.359	-0.933	-20.322	0.015
SD	0.143	1.267	0.342	14.835	0.008
Panel B: Emerging					
Brazilian Real	-0.136	-0.681	-0.946	-9.078	0.029
Bulgarian Lev	-0.005	-0.125	-1.172	-44.754	0.007
Croatian Kuna	0.044	0.810	-1.187	-42.257	0.009
Egyptian Pound	0.693	1.896	-0.136	-0.941	0.041
Hungarian Forint	-0.129	-0.963	-1.473	-21.149	0.021
Indian Rupee	-0.099	-0.857	-0.401	-7.712	0.018
Indonesian Rupiah	0.653	0.927	-0.967	-5.278	0.071
Kuwaiti Dinar	0.020	0.540	-0.174	-7.688	0.005
Malaysian Ringgit	0.120	0.246	-0.416	-3.858	0.043
Mexican Peso	-0.311	-1.930	-0.533	-5.610	0.026
Philippine Peso	-0.061	-0.435	-0.253	-6.190	0.019
Polish Zloty	-0.155	-0.907	-1.391	-18.459	0.021
Russian Ruble	-0.370	-1.521	-0.918	-8.661	0.029
Saudi Riyal	0.008	1.461	-0.003	-2.238	0.001
Slovak Koruna	0.462	2.434	-1.052	-8.840	0.014
Slovenian Tolar	0.301	5.097	-0.712	-13.858	0.004
South Africa Rand	-0.553	-2.884	-0.877	-8.226	0.036
Thai Baht	-0.007	-0.038	-0.428	-5.745	0.025
Ukrainian Hryvnia	-0.473	-1.131	-0.333	-2.404	0.050
Average	0.0001	0.102	-0.704	-11.734	0.025
SD	0.330	1.749	0.433	11.990	0.017

This table reports the regression results for each currency based on Equation 17. α and β represent the intercept and slope coefficients of the dollar index, and $t(\alpha)$ and $t(\beta)$ report the t-values measured by Heteroskedasticity- and autocorrelation-consistent (HAC) standard errors of Newey and West (1987). σ_ϵ refers to the standard deviation of the noise term ϵ_t in Equation 17, which is also known as idiosyncratic volatility. '*', '**', '***' represent that the t-values of α are statistically significant at the 10%, 5% and 1% levels, repsectively.