



Harris, R. D. F., Shen, J., & Yilmaz, F. (2022). Maximally Predictable Currency Portfolios. *Journal of International Money and Finance*, 128, Article 102702. <https://doi.org/10.1016/j.jimonfin.2022.102702>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1016/j.jimonfin.2022.102702](https://doi.org/10.1016/j.jimonfin.2022.102702)

[Link to publication record on the Bristol Research Portal](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Elsevier at <https://doi.org/10.1016/j.jimonfin.2022.102702> . Please refer to any applicable terms of use of the publisher.

University of Bristol – Bristol Research Portal

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available: <http://www.bristol.ac.uk/red/research-policy/pure/user-guides/brp-terms/>

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Maximally predictable currency portfolios

Richard D.F. Harris^{a,*}, Jian Shen^b, Fatih Yilmaz^c^a University of Bristol, United Kingdom^b University of Exeter, United Kingdom^c Eurizon SLJ Asset Management, United Kingdom

ARTICLE INFO

Article history:

Available online 7 July 2022

Keywords:

Currencies
 Predictability
 Trading strategies
 Maximally predictable portfolio
 Momentum and reversals

ABSTRACT

We investigate the predictability of the G10 currencies with respect to lagged currency returns from the perspective of a U.S. investor, using the maximally predictable portfolio (MPP) approach of [Lo and MacKinlay \(1997\)](#). We show that, out-of-sample, the MPP yields a higher Sharpe ratio, higher cumulative return and lower maximum drawdown than both a naïve equal-weighted portfolio of the currencies and an equal-weighted portfolio of momentum trading strategies, and that a mean–variance investor would be willing to pay a performance fee to switch from the naïve and momentum portfolios to the MPP. The MPP has performed particularly well since the 2008 financial crisis, in contrast with the momentum portfolio, the value of which declined significantly over this period. Our results are robust to the estimation window length, the type and level of portfolio weight constraints and transaction costs.

© 2022 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In this paper, we estimate the maximally predictable portfolio (MPP) of the G10 currencies with respect to lagged currency returns from the perspective of a U.S. investor and evaluate its out-of-sample performance. The MPP is defined as the portfolio with the highest coefficient of determination (i.e., R^2) between the realized portfolio return and its expectation conditional on a set of predictor variables (see, for example, [Lo and MacKinlay, 1997](#); [Cheung, He and Ng, 1997](#)). Ex post, the MPP provides an upper bound on the predictability of the portfolio return and can be useful for identifying the sources of predictability in returns, as well as providing a benchmark against which to evaluate other predictive models of returns that condition on the same predictor variables. Ex ante, the MPP approach can be used to construct portfolios that optimally exploit the information content of the predictor variables—in this case, lagged returns—in order to predict future portfolio returns out-of-sample as part of a dynamic investment strategy.

The predictability of currency returns—and their out-of-sample predictability in particular—is an important issue both for the light that it sheds on market efficiency and, relatedly, for the practical implications that it has for investors, including those who invest in foreign exchange directly and those who invest in other asset classes internationally and are consequently exposed to foreign exchange risk. Many models have been developed to predict currency returns, broadly falling into three categories: those that use information contained in lagged exchange rates (i.e., momentum/reversal strategies); those that exploit the failure of uncovered interest parity and use information contained in forward exchange rates (i.e., carry strategies); and those that are based on economic fundamentals (i.e., value strategies). There is a large literature that

* Corresponding author. School of Accounting and Finance, University of Bristol, United Kingdom.
 E-mail address: Richard.Harris@bristol.ac.uk (R.D.F. Harris).

documents the performance of these strategies and, over the half a century since the advent of floating exchange rates, the findings are mixed. Many early studies report evidence of significant predictability, particularly using momentum and carry, with returns that are sufficiently large to overcome reasonable transaction costs.¹ More recently, however, the evidence of predictability in the foreign exchange market has diminished, with the decline in predictability most notable following the financial crisis of 2008. To illustrate this, Panel A of Fig. 1 plots the cumulative return over the period February 1994 to December 2020 (the out-of-sample period used in our empirical analysis) of (a) the passive equal-weighted long-only portfolio of the G10 currencies, (b) an equal-weighted portfolio of second order moving average (MA(2)) momentum strategies applied to the individual currencies and (c) an MA(2) momentum strategy applied to the equal-weighted long-only portfolio. While the two momentum strategies outperformed the benchmark passive strategy until 2008, they have significantly underperformed it since then, delivering negative cumulative returns.² The failure of such strategies to generate positive returns since the financial crisis is not just of academic interest. Momentum, carry and value underpin the investment strategies of many professional investment companies that specialise in currency investment and a large number of these funds have closed over the last decade. Panel 2 of Fig. 1 plots the Citi Parker Global Currency index (a widely used index that tracks the performance of active foreign exchange funds) from its inception in January 2003 until its termination in August 2019. The index peaked in June 2009 and has been in secular decline since then. By the end of the sample, it was about 15 percent below its maximum.

We consider an alternative approach to forecasting currency returns that, like the momentum trading strategy, exploits information contained in lagged returns. However, unlike the momentum trading strategy, our implementation of the MPP conditions each currency's return not just on its own lagged values, but also on the lagged values of all other currencies, thus exploiting important information in cross-correlations, which is typically ignored in a momentum trading strategy. Moreover, a notable feature of this approach is that it focuses on the predictability of the portfolio as a whole, rather than on the predictability of the individual currencies, and as such it could be expected to perform well in terms of the portfolio level evaluation measures against which it is assessed. We apply the MPP approach to the G10 currencies from the perspective of a U.S. investor (a total of nine currencies expressed in terms of the USD) using monthly returns over the period February 1974 to December 2020. We estimate the MPP using the procedure developed by Lo and MacKinlay (1997), which in turn is based on the canonical correlation analysis of Box and Tiao (1977).^{3,4} Lo and MacKinlay show that in the unconstrained case the weights of the MPP and the resulting maximal R^2 are determined by the solution to a generalized eigenvalue problem, which we set out in the following section. As noted by Lo and MacKinlay, the MPP approach allows one to characterize the sources of predictability in returns. In particular, it helps to distinguish the factors that best *explain* returns from those that best *predict* returns. To illustrate this in the context of the G10 currency portfolio, we use principal components analysis to capture the underlying factor structure of the currency returns. We show that the principal components with the greatest explanatory power for returns are those related to the level of the USD, the 'slope' between high interest rate currencies and low interest rate currencies (i.e., a carry factor), and the level of the JPY. Together, these account for 82 percent of the variation in G10 currency returns. Whilst such factors might at first sight be natural choices as conditioning variables (because they are clearly important in explaining returns contemporaneously), we show that they are not the most predictable factors and hence not the most useful for constructing predictable portfolios of the G10 currencies. For example, the R^2 from a regression of the first principal component (i.e., the USD 'level' factor) returns on its first lag is 0.002 and on the lags of all nine principal components is 0.011. In contrast, the sixth principal component—which explains only 3.2 percent of the variation in returns contemporaneously—has an R^2 of 0.025 on its own lag and 0.052 on the lags of all nine principal components. Consequently, a portfolio that loads heavily on the sixth principal component will, *ceteris paribus*, be more predictable than one that loads heavily on the first principal component. While not the focus of this paper, a similar reasoning would apply when using factors other than lagged returns to model currency returns, such as those based on economic fundamentals.

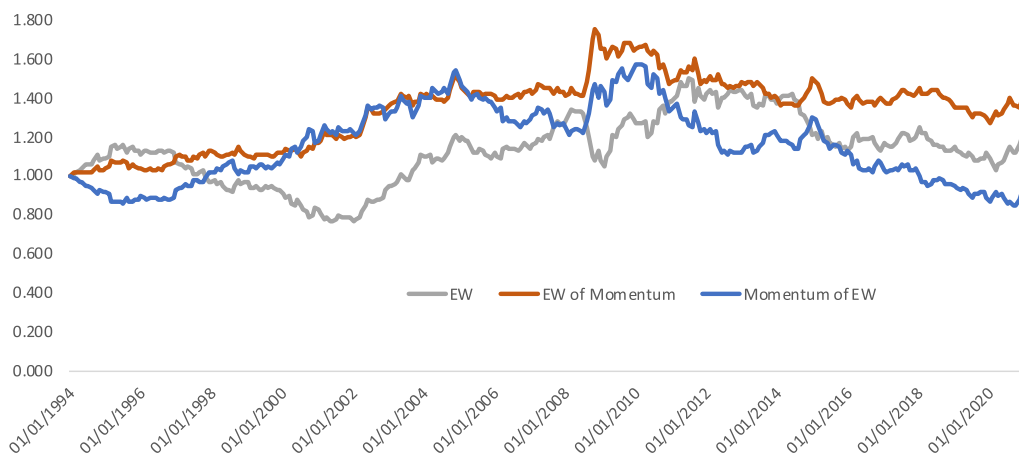
We first estimate the in-sample MPP. As noted above, rather than working in terms of the individual currency return series, we choose instead to use their principal components. The full set of principal components contain exactly the same information as the individual return series and consequently our results are independent of this choice. However, the use of principal components allows us to attach some economic meaning to the most important predictor variables that would be difficult to disentangle from the individual return series. Moreover, by using principal components, which are orthogonal by construction, the total (as opposed to marginal) contribution of each predictor variable can be easily evaluated without

¹ For evidence on momentum trading strategies, see, for example, Sweeney (1986), Taylor and Allen (1992), Levich and Thomas (1993), LeBaron (1999), Okunev and White (2003) and Harris and Yilmaz (2009). For evidence on carry trading strategies, see Burnside, Eichenbaum, and Rebelo (2008) and Della Corte, Sarno and Tsiakas (2009). Models based on fundamentals have fared less well, with little evidence of economically significant predictability over anything but the longest horizons. See, for example, Cheung, Chinn and Pascual (2005) and Della Corte, Sarno and Tsiakas (2009).

² Carry trading strategies have similarly underperformed since the financial crisis. For example, the DB G10 Currency Future Harvest Index, which tracks the performance of a carry trading strategy applied to the G10 currencies, had an annualized return of just 0.08% between its inception in September 2006 and December 2020 (Invesco, 2020).

³ The MPP approach is concerned with optimizing over the set of portfolio weights holding the predictor variables fixed. An alternative but related problem is to find the optimal set of predictor variables for a given portfolio, which is addressed by Foster, Smith and Whaley (1997).

⁴ d'Aspremont (2011) uses the Box and Tiao (1977) procedure to identify linear combinations of asset prices that have maximum mean-reversion and is hence concerned with minimal rather than maximal predictability. Although the author assumes that individual asset prices are stationary, this approach is closely related to the Johansen (1988) approach to the estimation of cointegrating vectors (see Bewley, Orden, Yang and Fisher (1994) for a comparison of the Box and Tiao (1977) and Johansen (1988) approaches).



Panel B: Citi Parker Global Currency Index



Fig. 1. FX Trading Strategy Performance. Panel A shows the cumulative return over the U.S. risk-free rate of (a) the equal-weighted portfolio of the G10 currencies (the grey line), (b) the equal-weighted portfolio of an MA(2) momentum strategy applied to the individual G10 currencies (the orange line) and (c) the MA(2) momentum strategy applied to the equal-weighted portfolio of G10 currencies (the blue line). The sample period is 01/1994 to 12/2020. The cumulative return is set to 1 in 12/1993 for all three strategies. Transaction costs are assumed to be zero. Panel B shows the Citi Parker Global Currency Index for the period 01/2003 to 08/2019. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the need to estimate a separate regression for each variable. Of course, the role of the individual currencies in predicting portfolio returns is readily backed out from that of the principal components.

Knowledge of the in-sample MPP and its corresponding R^2 is useful for characterizing the intertemporal relationships in returns and for determining an upper bound on the achievable predictability of a forecasting model, but it sheds no light on the economic importance of predictability. To address this issue, we go on to examine the out-of-sample performance of the MPP. In particular, we employ a rolling window to estimate the MPP each month and use this to predict the one-month ahead return. We convert the return prediction into a buy/sell signal and then evaluate the performance of this out-of-sample strategy. As with portfolio optimization in the mean-variance framework, unconstrained estimation of the MPP can potentially lead to somewhat extreme portfolio weights and the problem is compounded in the out-of-sample analysis owing to the relatively smaller estimation sample size. Consequently, we impose constraints on the portfolio weights using

three approaches: conventional leverage constraints on the individual currency positions within the MPP, maximum weight constraints in a long only portfolio, and norm constraints on the set of portfolio weights, as in DeMiguel, Garlappi, Nogales and Uppal (2009b), which constrain the size of the positions in aggregate rather than individually. We compare the out-of-sample performance of the MPP with two benchmarks that are also based solely on currency returns. The first is the equal-weighted portfolio of the nine currencies. In the context of equities, DeMiguel, Garlappi and Uppal (2009a) show that the naïve, equal-weighted portfolio is a difficult benchmark to beat, owing primarily to fact that it does not rely on (noisy) estimates of the parameters that are required for optimization. Indeed, among currency funds, equal-weighted portfolios are widely used as benchmarks for actively managed portfolios. The second benchmark is an equal-weighted portfolio of momentum strategies applied to the individual currencies, using a simple but representative moving average filter. We evaluate the MPP and the benchmark strategies over both the full sample and sub-samples either side of the 2008 financial crisis, using a range of statistical and economic criteria.

Our findings can be summarized as follows. The in-sample MPP has an R^2 with respect to the predictor variables that is several orders of magnitude larger than that of a first-order autoregressive model for the individual currencies. It is also substantially higher than that obtained from a regression of each of the individual currencies on all nine lagged principal components, and from a regression of the equal-weighted portfolio on all nine principal components. The MPP has large long/short positions in NZD/AUD and SEK/CHF, which are commonly traded currency pairs in practice. The MPP loads most heavily on the sixth and seventh principal components (which itself has large long positions in NOK, NZD and SEK and large short positions in AUD and CHF), the seventh principal component (which has large long positions in AUD and SEK and a large short position in CAD) and, to a lesser extent, on the first principal component (the USD level factor). In the out-of-sample analysis, the MPP yields a higher Sharpe ratio than the naïve equally weighted portfolio and, at low levels of leverage, higher also than the momentum portfolio. In most cases, it also has a lower maximum drawdown than both the naïve portfolio and the momentum portfolio. We show that an investor with preferences that can be approximated by a mean–variance utility function would be willing to pay a fee of up to 166 basis points per year to switch from the naïve portfolio to the MPP, and up to 115 basis points per year to switch from the momentum portfolio. The MPP performs especially well in the period following the 2008 financial crisis, during which the momentum portfolio experienced negative returns. Our conclusions are robust to the estimation window length, the different types and levels of portfolio weight constraints, and transaction costs.

The outline of the remainder of the paper is as follows. In the following section, we describe the data that we use in the empirical analysis. In Section 3, we outline the MPP methodology, the benchmark portfolios, the estimation approach and the evaluation measures. Section 4 reports the results of the in-sample and out-of-sample analysis. Section 5 explores the robustness of our findings to the estimation sample, alternative leverage constraints and transaction costs. Section 6 provides a summary and offers some concluding remarks.

2. Data

Our sample comprises monthly exchange rate data for the USD against the remaining G10 currencies (AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD and SEK). End-of-month spot exchange rates against the USD were obtained from Refinitiv Datastream for the period January 1974 to December 2020.⁵ Forward exchange rates, which we use to compute returns for the out-of-sample trading strategy, were obtained for the period January 1994 to December 2020. We estimate the MPP using exchange rate returns net of the interest carry component, calculated as $r_{i,t} = S_{i,t}/S_{i,t-1} - 1$, where $S_{i,t}$ is the spot exchange rate at time t expressed as the number of USD per 1 unit of currency i .⁶ Table 1 reports summary statistics of the nine exchange rate return series for the full sample, February 1974 to December 2020. Panel A reports the mean, standard deviation and the first six autocorrelations for each currency, while Panel B reports the first-order cross-correlations between each pair of currencies. The first-order autocorrelation coefficient is positive for six of the nine currencies and negative for three, although in two of these cases the coefficient is very close to zero. In no case is the first-order autocorrelation of a currency larger than all of its first-order cross-correlations with other currencies. Indeed, in two cases (CHF and NZD), all of the currency's cross-correlations are larger in absolute value than its autocorrelation. The average autocorrelation across the nine currencies is 0.033, while the average of their largest cross-correlations is 0.064. Clearly, for every currency the lagged values of the other currencies contain more information than its own lagged value.

Table 2 reports the results of principal components analysis applied to the nine currency return series. Panel A reports the eigenvectors (which are proportional to the currency weights in each of the principal components), the eigenvalues and the percentage and cumulative percentage of the variance explained by each principal component. The first principal component is (approximately) an equal-weighted average of the nine currency return series, thus representing a USD 'level' factor. This accounts for more than 50 percent of the variation in the returns of the nine currencies. The second principal component has large positive weights on AUD and NZD and large negative weights on JPY and CHF and represents the 'carry' factor between

⁵ For the euro, before its introduction in 1999, we use synthetic rates computed by Datastream.

⁶ Using returns including the interest carry component in place of pure exchange rate returns in the estimation of the MPP has a negligible impact on the empirical analysis as the variation in interest rates is several orders of magnitude lower than the variation in exchange rates. However, using pure exchange rate returns makes our approach consistent with the moving average momentum strategy, which also employs data on exchange rates only and not interest rates.

Table 1
Summary Statistics.

	AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK
Panel A: Mean, Standard Deviation and Autocorrelations									
Mean	-0.06%	-0.02%	0.28%	-0.01%	-0.04%	0.23%	-0.02%	-0.06%	-0.05%
St Dev	3.19%	1.97%	3.35%	2.86%	2.89%	3.18%	3.08%	3.41%	3.09%
ρ_1	0.019	-0.042	-0.002	0.037	0.058	0.037	0.022	-0.004	0.075
ρ_2	0.005	0.007	0.049	0.065	0.046	0.071	0.052	0.002	0.035
ρ_3	0.027	-0.047	0.036	0.054	0.023	0.047	-0.007	0.143	0.048
ρ_4	-0.043	0.098	-0.065	-0.020	0.009	0.003	-0.057	-0.080	0.040
ρ_5	-0.028	-0.008	0.008	0.029	-0.009	-0.050	0.002	-0.011	0.013
ρ_6	0.075	-0.002	-0.052	0.010	-0.037	-0.085	0.041	0.072	-0.021
LB(6)	0.210	0.990	0.002	0.795	1.895	0.783	0.264	0.011	3.214
Panel B: First Order Cross-Correlations									
AUD	0.019	-0.011	0.010	0.003	0.004	0.033	0.015	-0.007	0.012
CAD	0.036	-0.042	0.050	0.032	-0.003	0.035	0.040	0.007	0.064
CHF	-0.016	-0.009	-0.002	-0.018	-0.024	-0.015	-0.024	-0.034	-0.041
EUR	0.042	-0.010	0.050	0.037	0.013	-0.002	0.025	-0.002	0.027
GBP	0.052	0.056	0.054	0.053	0.058	-0.001	0.050	0.048	0.069
JPY	0.009	0.010	0.074	0.035	0.010	0.037	-0.020	0.021	-0.012
NOK	0.063	-0.004	0.081	0.043	0.036	0.042	0.022	0.034	0.046
NZD	0.070	0.007	0.071	0.056	0.068	0.041	0.067	-0.004	0.052
SEK	0.062	-0.011	0.097	0.078	0.067	-0.004	0.044	0.024	0.075

The table reports summary statistics for the monthly returns of the nine currencies. Panel A reports the mean and standard deviation of returns, the first six autocorrelations and the Ljung-Box statistic to test the null hypothesis that the first six autocorrelations are jointly equal to zero. Panel B reports the first order cross-correlation between each pair of currencies. Each element gives the correlation between the month t return of the currency in the left column with the month $t-1$ return of the currency in the top row. The sample period is 01/1974 to 12/2020.

Table 2
Principal Components Analysis.

	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9
Panel A: Eigenvectors, Eigenvalues and Proportion of Variance Explained									
AUD	0.298	0.579	0.174	-0.275	0.221	-0.561	0.309	0.103	0.011
CAD	0.145	0.255	-0.085	-0.191	0.342	0.045	-0.770	-0.399	0.018
CHF	0.405	-0.336	0.004	-0.096	-0.555	-0.456	-0.226	-0.129	-0.361
EUR	0.376	-0.182	-0.163	-0.059	-0.140	-0.046	0.009	-0.019	0.879
GBP	0.313	-0.094	-0.204	0.800	0.388	-0.229	0.019	-0.020	-0.088
JPY	0.236	-0.396	0.817	-0.051	0.323	0.115	-0.014	0.001	0.022
NOK	0.394	-0.082	-0.273	-0.239	0.173	0.270	-0.156	0.736	-0.193
NZD	0.349	0.525	0.291	0.344	-0.452	0.437	-0.065	0.021	-0.011
SEK	0.393	-0.083	-0.261	-0.237	0.137	0.378	0.482	-0.520	-0.224
Eigenvalue	0.473	0.126	0.076	0.040	0.038	0.027	0.020	0.017	0.010
Proportion	57.3%	15.3%	9.2%	4.9%	4.6%	3.2%	2.4%	2.0%	1.2%
Cumulative	57.3%	72.6%	81.8%	86.6%	91.2%	94.4%	96.8%	98.8%	100.0%
Panel B: First Order Cross-Correlations									
PC1	0.049	-0.013	0.003	0.016	-0.018	0.080	0.027	0.009	0.020
PC2	0.027	-0.009	0.000	0.024	0.070	-0.007	0.064	-0.050	0.034
PC3	0.011	0.031	0.058	0.011	0.018	-0.043	0.076	-0.004	0.028
PC4	-0.053	-0.031	-0.052	0.022	-0.055	0.019	0.043	-0.010	-0.008
PC5	0.070	0.015	-0.028	-0.012	-0.071	-0.037	0.032	0.095	0.019
PC6	-0.134	0.022	-0.024	0.013	0.037	-0.143	-0.096	-0.033	-0.034
PC7	0.028	0.018	0.052	-0.063	0.007	0.034	-0.014	0.028	-0.035
PC8	0.001	-0.019	0.112	0.016	-0.022	-0.003	0.052	-0.014	-0.033
PC9	-0.012	-0.003	-0.017	-0.064	-0.002	-0.009	-0.052	0.051	0.033

The table reports the results of principal components analysis applied to the monthly currency return series against the USD. Panel A reports the eigenvectors, eigenvalues ($\times 100$), the proportion of variance explained and the cumulative proportion of variance explained. Panel B reports the first order cross-correlation between each pair of principal components. Each element gives the correlation between the month t return of the principal component in the left column with the month $t-1$ return of the principal component in the top row. The sample period is 01/1974 to 12/2020.

high and low interest rate currencies. The third principal component has a high weighting on JPY, thus representing a JPY 'level' factor.⁷ The first three principal components together explain 82 percent of the total variation in returns among the nine currencies. The remaining factors, which have substantially less explanatory power, have less obvious economic interpretations.

Panel B of Table 2 reports the first-order cross-correlations of the nine principal components. Notably, the principal components that have the most explanatory power for returns do not have the highest first-order autocorrelations. For example, the first three principal components (the USD factor, the carry factor and the JPY factor) have first-order autocorrelations of 0.049, -0.009 and 0.058, respectively, while the fifth and sixth principal components have first-order autocorrelations of -0.071 and -0.143, respectively. Consequently, the factors that provide the most explanatory power for returns contemporaneously are different from those that have the most predictive ability one month ahead. A portfolio that loads heavily on the USD, carry and JPY factors will, ceteris paribus, be less predictable than one that loads heavily on the fifth and sixth principal components. As with the raw currency returns, there is significant information in the cross-correlations of the principal component returns: the average autocorrelation of the nine principal components is 0.046 while the average of their largest cross-correlations is 0.084.

The out-of-sample estimation of the MPP uses a rolling window of 240 months, and so the evaluation period is February 1994 to December 2020 (with the period February 1974 to January 1993 used for the estimation of the first rolling window). For the out-of-sample trading strategy, we evaluate the MPP, equal-weighted and momentum portfolios using returns including the interest carry component, calculated as $R_{i,t} = S_{i,t}/F_{i,t-1} - 1$, where $F_{i,t-1}$ is the 1-month forward rate at time $t - 1$, again expressed as USD per units of foreign currency i (see, for example, Barroso and Santa-Clara, 2015). The out-of-sample evaluation period is split into two sub-periods, February 1994 to December 2008 and January 2009 to December 2020.

3. Methodology

We denote the set of nine exchange rate returns against the USD at time t by $\mathbf{r}_t = [r_{1,t} \cdots r_{9,t}]'$. The return of a portfolio of the nine currencies is given by:

$$r_{p,t} = \mathbf{w}'\mathbf{r}_t \quad (1)$$

where \mathbf{w} is the 9x1 portfolio weight vector with $\mathbf{1}'\mathbf{w} = 1$ and $\mathbf{1}$ is an 9x1 vector of 1s. Consider the conditional expectation of the return vector, $\hat{\mathbf{r}}_t = E[\mathbf{r}_t|\Omega_{t-1}]$, where Ω_{t-1} is the information set at $t - 1$. We can write the relation between the return vector and its forecast as:

$$\mathbf{r}_t = \hat{\mathbf{r}}_t + \boldsymbol{\epsilon}_t \quad (2)$$

where $E[\boldsymbol{\epsilon}_t|\Omega_{t-1}] = 0$ and $\text{var}[\boldsymbol{\epsilon}_t|\Omega_{t-1}] = \mathbf{V}$. The corresponding forecast of the portfolio return is given by:

$$\hat{r}_{p,t} = \mathbf{w}'\hat{\mathbf{r}}_t \quad (3)$$

The relation between the portfolio return and its forecast is given by:

$$r_{p,t} = \hat{r}_{p,t} + \varepsilon_{p,t} \quad (4)$$

where $\varepsilon_{p,t} = \mathbf{w}'\boldsymbol{\epsilon}_t$. The predictability of $r_{p,t}$ is measured by the R^2 coefficient between $r_{p,t}$ and $\hat{r}_{p,t}$ (see, for example, Box and Tiao, 1977), which is given by:

$$R^2 = \frac{\text{var}(\hat{r}_{p,t})}{\text{var}(r_{p,t})} = \frac{\mathbf{w}'\hat{\boldsymbol{\Sigma}}\mathbf{w}}{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}} \quad (5)$$

where $\boldsymbol{\Sigma} = \text{var}(\mathbf{r}_t)$ and $\hat{\boldsymbol{\Sigma}} = \text{var}(\hat{\mathbf{r}}_t)$. To implement the MPP approach, we must specify a model for the conditional expectation of \mathbf{r}_t . For the predictor variables, we use the set of lagged principal components of \mathbf{r}_t , which we denote $\mathbf{x}_{t-1} = [pc_{1,t-1} \cdots pc_{9,t-1}]'$. As noted above, this is equivalent to using the lagged returns of the nine currencies but has the advantage of allowing us to attach an economic interpretation to the role of at least some of the predictor variables that would be otherwise difficult to discern. Our model for the conditional expectation of \mathbf{r}_t is therefore given by:

$$\mathbf{r}_t = \mathbf{B}_0 + \mathbf{B}_1\mathbf{x}_{t-1} + \boldsymbol{\epsilon}_t \quad (6)$$

where \mathbf{B}_0 and \mathbf{B}_1 are, respectively, a 9x1 vector and a 9x9 matrix of parameters. The conditional expectation of \mathbf{r}_t is given by:

$$\hat{\mathbf{r}}_t = \hat{\mathbf{B}}_0 + \hat{\mathbf{B}}_1\mathbf{x}_{t-1} \quad (7)$$

The actual portfolio return and forecast portfolio return are given by.

⁷ It is well established in the literature that the first two principal components of USD exchange rates represent the USD factor and the interest rate slope factor, respectively (see, for example, Lustig, Roussanov and Verdelhan, 2011). The fact that the third principal component heavily weights the JPY reflects the effective decoupling of the JPY from other major currencies that took place from the 1980s onwards, partly caused by differential monetary policy.

$$r_{p,t} = \mathbf{w}'\mathbf{B}_0 + \mathbf{w}'\mathbf{B}_1\mathbf{x}_{t-1} + \varepsilon_{p,t} \quad (8)$$

$$\hat{r}_{p,t} = \mathbf{w}'\hat{\mathbf{B}}_0 + \mathbf{w}'\hat{\mathbf{B}}_1\mathbf{x}_{t-1} \quad (9)$$

The R^2 between $r_{p,t}$ and $\hat{r}_{p,t}$ is equal to:

$$R^2 = \frac{\mathbf{w}'\hat{\mathbf{B}}_1\Sigma\hat{\mathbf{B}}_1'\mathbf{w}}{\mathbf{w}'\Sigma\mathbf{w}} \quad (10)$$

Equation (10) is the generalized Rayleigh quotient of $\hat{\Sigma} = \hat{\mathbf{B}}_1\Sigma\hat{\mathbf{B}}_1'$ relative to Σ , and its maximal value is given by the largest eigenvalue, λ^{\max} , of $\mathbf{R}^{-1'}\hat{\Sigma}\mathbf{R}^{-1}$, where $\mathbf{R} = \Sigma^{\frac{1}{2}}$ is the Cholesky decomposition of Σ . This is achieved by the vector $\mathbf{R}^{-1'}\mathbf{v}^{\max}$, where \mathbf{v}^{\max} is the eigenvector associated with λ^{\max} . The solution to the unconstrained maximization of (10) can then be normalized so that the portfolio weights sum to unity:

$$\mathbf{w}_{MPP} = \frac{\mathbf{R}^{-1'}\mathbf{v}^{\max}}{\mathbf{1}'\mathbf{R}^{-1'}\mathbf{v}^{\max}} \quad (11)$$

(see, for example, Box and Tiao, 1977; Lo and MacKinlay, 1997; Cheung, He and Ng, 1997).

3.1. Constraints

Unconstrained estimation of the MPP inevitably leads to weights that can be negative as well as positive and, in some cases, particularly in the out-of-sample analysis, the weights can take extreme values. We therefore consider various constraints on the MPP weights. First, we impose symmetric leverage constraints with respect to both long and short positions:

$$-a \leq \mathbf{w} \leq 100\% + a \quad (12)$$

where $a \in \{0\%, 25\%, 50\%, 75\%, 100\%\}$. Second, we consider a long only portfolio with maximum investment constraints:

$$0\% \leq \mathbf{w} \leq b \quad (13)$$

where $b \in \{25\%, 50\%, 75\%, 100\%\}$. The case of $b = 100\%$ in (13) corresponds to $a = 0\%$ in (12). In Section 6, we consider alternative forms of constraints as a robustness check.

3.2. Estimation

For the in-sample analysis, we calculate the weights of the MPP using the eigen decomposition described above. We also estimate the in-sample MPP with a long-only constraint (i.e., $a = 0\%$). For the out-of-sample analysis, we impose the full range of constraints described above. In order to solve for the MPP in the constrained cases, we use numerical optimization. In particular, we use the unconstrained weights from the eigen decomposition as starting values for the numerical procedure and employ a multi-start algorithm to obtain the final optimized weights.

3.3. Benchmark portfolios

We consider two commonly used benchmark portfolios. The first is the equal-weighted portfolio of the nine currencies, which is a commonly used passive benchmark for currencies, representing a 'market' risk factor (see, for example, Lustig et al., 2011; Menkhoff et al., 2012a, 2012b). More generally, the equal-weighted or 'naïve' portfolio has been found to be a difficult benchmark to beat since it obviates the need to estimate either the mean vector or the covariance matrix of returns (see, for example, DeMiguel, Garlappi and Uppal, 2009a). The second benchmark that we consider is based on a simple momentum strategy applied to each of the individual exchange rates. In particular, for each of the nine currencies, if the exchange rate return between month $t - 1$ and month t is greater than zero, a long position is taken in the currency, otherwise a short position is taken. This is equivalent to a second-order moving average (MA(2)) filter strategy that generates a buy signal for month $t + 1$ if the month t exchange rate is greater than the average exchange rate over month t and month $t - 1$. There are of course many other potential benchmarks that we could consider, such as those based on carry, value, the yield curve and other macroeconomic indicators. However, our aim is to provide a comparison with strategies that condition on the same information set, namely lagged exchange rates. The MPP approach could easily be adapted to include other predictors and could be compared against these alternative models, but we leave that for another paper.

3.4. Out-of-sample trading strategy

We examine the out-of-sample effectiveness of the strategy using a rolling window approach. In particular, we first estimate the principal components of the currency return series and the parameters of the conditional mean model given by (6) using the sample $t = 1, \dots, \tau$, and compute the MPP for month $t = \tau$ using (11). We use this to generate the forecast MPP

return for month $t = \tau + 1$. We then roll the estimation window forward by one month and forecast the MPP return for month $t = \tau + 2$, and so on until the end of the sample. Following Lo and MacKinlay (1997), for the main analysis we use an estimation window length of $\tau = 240$ but we also consider alternative window lengths of $\tau = 120$ and $\tau = 60$ in our robustness tests. If the forecast return is positive, the strategy takes a long position in the MPP, otherwise it takes a short position. To evaluate the portfolios, we first define the vector of returns including the interest carry component for the nine currencies, $\mathbf{R}_t = [R_{1,t} \cdots R_{9,t}]$. We calculate returns for the three portfolios for month t as:

$$R_{MPP,t} = I_{MPP,t} \mathbf{w}'_{MPP} \mathbf{R}_t \quad (14)$$

$$R_{EW,t} = \mathbf{w}'_{EW} \mathbf{R}_t \quad (15)$$

$$R_{Mom,t} = \mathbf{w}'_{EW} (\mathbf{I}_{Mom,t} \odot \mathbf{R}_t) \quad (16)$$

where \mathbf{w}_{EW} is the 9×1 vector of equal weights, $I_{MPP,t} = 1$ if $\hat{r}_{MPP,t} \geq 0$ and -1 if $\hat{r}_{MPP,t} < 0$, $\mathbf{I}_{Mom,t} = [I_{1,t} \cdots I_{9,t}]$ where $I_{i,t}$ is equal to 1 if $r_{i,t-1} \geq 0$ and equal to -1 if $r_{i,t-1} < 0$, and \odot is the Hadamard product. The full out-of-sample evaluation period is February 1994 to December 2020. In addition to the full sample, we report results for two sub-samples: February 1994 to December 2008 and January 2009 to December 2020. These are motivated by the fact that, as noted above, the financial crisis that occurred towards the end of 2008 marked a distinct change in the performance of conventional foreign exchange trading strategies.

3.5. Evaluation measures

We evaluate the out-of-sample performance of each of the three portfolios $j = (MPP, EW, Mom)$ using the following statistical and economic measures:

(a) Mean monthly return.

$$\hat{\mu}_j = \frac{1}{T - \tau} \sum_{t=\tau+1}^T R_{j,t} \quad (17)$$

(b) Standard deviation of monthly returns.

$$\hat{\sigma}_j = \left(\frac{1}{T - \tau - 1} \sum_{t=\tau+1}^T (R_{j,t} - \hat{\mu}_j)^2 \right)^{1/2} \quad (18)$$

(c) Sharpe ratio.

$$SR_j = \frac{\hat{\mu}_j}{\hat{\sigma}_j} \quad (19)$$

(d) Maximum drawdown.

$$MD_j = \frac{\max_t \left(\max_s (R_{j,s \leq t}) - R_{j,t} \right)}{\max_s (R_{j,s \leq t})} \quad (20)$$

(e) Terminal value.

$$TV_j = \prod_{t=\tau+1}^T (1 + R_{j,t}) \quad (21)$$

(f) Performance fee.

We measure the maximum fee that an investor would pay to switch from the equal-weight portfolio to the MPP or momentum portfolio, and from the momentum portfolio to the MPP. Following Fleming, Kirby and Ostdiek (2001), we use a quadratic utility function as a second order Taylor series approximation to the investor's true utility function. With this approximation, the investor's utility of wealth at time t , W_t , is given by:

$$U(W_t) = W_t - \frac{a}{2} W_t^2$$

where a is the investor's coefficient of absolute risk aversion. We assume that the investor's coefficient of relative risk aversion $\lambda = aW_t/(1 - aW_t)$ is constant, and so we can estimate the expected the investor's expected utility from investing in portfolio j over the evaluation period as:

$$\bar{U}_j(\cdot) = W_0 \sum_{t=\tau+1}^T \left(R_{j,t} - \frac{\lambda}{2(1 + \lambda)} R_{j,t}^2 \right)$$

The performance fee for portfolio j with respect to a benchmark b is defined as the fee, ϕ_j , that equates the utility from investing in the benchmark portfolio with the utility from investing in portfolio j , but where the return of portfolio j is reduced by the amount of the fee:

$$\sum_{t=\tau+1}^T (R_{j,t} - \phi_j) - \frac{\lambda}{2(1+\lambda)} (R_{j,t} - \phi_j)^2 = \sum_{t=\tau+1}^T R_{b,t} - \frac{\lambda}{2(1+\lambda)} R_{b,t}^2$$

We report the annualised performance fee for the MPP and momentum portfolio with respect to the equal-weight portfolio, and for the MPP with respect to the momentum portfolio. We use a value of $\lambda = 5$ for the risk aversion parameter.⁸

4. Results

In this section we present the results of the empirical analysis. Table 3 reports the results of regressing each of the nine currency return series on the nine lagged principal component return series. The R^2 from these regressions ranges from 0.004 (for the AUD) to 0.024 (for the SEK), and the average value is 0.013. For three of the currencies (AUD, CAD and CHF), none of the predictor variables is statistically significant at conventional significance levels. PC1 is significant at the 10 percent level for SEK, PC5 is significant at the 10 percent level for the JPY, and PC6 is significant at varying levels for the EUR, JPY, NOK, NZD and SEK. At an individual currency level, therefore, the evidence of predictability is minimal, but it is clear that while most of the principal components have little information about subsequent currency returns, PC6, and to a lesser extent PC1 and PC5 contain useful predictive information.

Table 4 reports the in-sample estimated unconstrained MPP (Panel A), the constrained long-only MPP (Panel B) and, for reference, the equal-weighted portfolio (Panel C). The unconstrained MPP has positive weights in CAD, EUR, JPY, NOK, NZD and SEK, and negative weights in AUD, CHF and GBP. The approximately equal but offsetting weights in AUD and NZD, CHF and SEK, and GBP and EUR, respectively, are consistent with the idea that these currencies are, in practice, frequently traded as pairs owing to their geographic and economic proximity, and that long-short portfolios in these currencies are often considered to be mean-reverting. The estimated coefficients on the lagged principal components are significantly positive for PC1 and PC7 and significantly negative for PC6. The large positive weights on NZD and SEK and the significance of PC6 reflects the results reported in Table 3 for the individual currencies. In particular, both currencies are positively correlated with PC6. SEK is also individually correlated with PC1, which is significant in the MPP. In the constrained long-only portfolio, only the NZD and SEK are included in the MPP and the only significant predictor is PC6, which is again consistent with the results of Table 3. The R^2 of the unconstrained MPP is 0.069, which is more than four times as large as the average R^2 for the individual currencies reported in Table 3. The constrained long-only MPP (Panel B) comprises positions in only the NZD and SEK and the R^2 is, not surprisingly, significantly lower at 0.026. Indeed, it is only marginally higher than the highest R^2 of the individual currencies, which is for the SEK. For the equal-weighted portfolio (Panel C), PC6 is marginally significant and its R^2 is only 0.011.

Tables 5 and 6 report the results of the out-of-sample analysis for the MPP, the equal-weighted portfolio and the momentum portfolio for the full sample and the two sub-samples, respectively. For the MPP, results are reported for the long only portfolio with maximum weight constraints $b \in \{25\%, 50\%, 75\%, 100\%\}$ and the leveraged portfolio with leverage constraints $a \in \{25\%, 50\%, 75\%, 100\%\}$. For the full sample (Table 5), using both the maximum weight constraint and the maximum leverage constraint, relaxing the constraint reduces the mean return of the MPP, increases its standard deviation and consequently reduces its Sharpe ratio. The maximum drawdown and terminal value also decline as the constraints are relaxed. In all cases, the MPP has a higher Sharpe ratio and terminal value than the equal-weight portfolio and, in all except the highest leverage cases, a lower drawdown. At low levels of leverage (i.e., a maximum weight of 25% or 50% in the long only portfolio, and a maximum leverage of 25%), the MPP also has a higher Sharpe ratio than the momentum portfolio. Even at some higher levels of leverage, the Sharpe ratio is only marginally lower. In terms of cumulative return, the MPP outperforms the momentum portfolio, with a higher terminal value at all but the very highest level of leverage. The last two rows report the performance fee against the equal-weight portfolio and the momentum portfolio respectively. Investors would be willing to pay a fee of up to 166 basis points per year to switch from the equally weighted portfolio to the MPP, and up to 115 basis points per year to switch from the momentum portfolio to the MPP. Indeed, the performance fee is positive for all the MPP portfolios with respect to both the equal-weight portfolio and the momentum portfolio. Thus, in the full sample, the MPP outperforms the equal-weight portfolio and, in almost all cases, the momentum portfolio also.

In the pre-2008 sub-sample (Table 6, Panel A), at lower levels of leverage, the long only MPP with the 25% weight constraint outperforms the EW portfolio, with a higher Sharpe ratio higher terminal value and a positive performance fee, but the remaining MPPs underperform the equal-weight portfolio. The MPP also underperforms the momentum portfolio during this period, which has a higher Sharpe ratio and terminal value than even the best performing MPP. However, in the post-2008 sub-sample (Table 6, Panel B), the MPP very significantly outperforms both the equal-weight portfolio and the momentum portfolio. The Sharpe ratio for the equal-weight portfolio is 0.065, while for the momentum portfolio, it is -0.287 . In contrast, the Sharpe ratio for the MPP ranges from 0.273 to 0.464. Similarly, the terminal value of the MPP ranges from

⁸ In unreported results, we also used values of $\lambda = 1$ and $\lambda = 10$, both of which lead to broadly similar conclusions to the case of $\lambda = 5$.

Table 3
Individual Currency Predictability (In-Sample).

	Const	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	R ²
AUD	−0.001 (0.001)	0.006 (0.02)	−0.011 (0.038)	−0.028 (0.049)	−0.033 (0.068)	0.040 (0.070)	0.030 (0.083)	−0.063 (0.097)	−0.051 (0.105)	−0.061 (0.137)	0.004
CAD	0.000 (0.001)	0.012 (0.012)	−0.013 (0.023)	−0.007 (0.030)	−0.059 (0.041)	−0.020 (0.043)	0.012 (0.051)	−0.136 (0.059)	−0.011 (0.064)	−0.092 (0.084)	0.018
CHF	0.003*** (0.001)	−0.014 (0.021)	−0.011 (0.040)	0.002 (0.052)	−0.021 (0.071)	−0.024 (0.073)	0.095 (0.087)	0.074 (0.101)	−0.028 (0.110)	0.013 (0.143)	0.005
EUR	0.000 (0.001)	0.013 (0.018)	−0.005 (0.034)	0.019 (0.044)	−0.046 (0.06)	−0.034 (0.062)	0.124* (0.074)	−0.071 (0.086)	−0.016 (0.094)	−0.005 (0.122)	0.009
GBP	0.000 (0.001)	0.027 (0.018)	0.018 (0.034)	0.035 (0.044)	0.007 (0.061)	−0.002 (0.063)	0.024 (0.075)	0.004 (0.087)	0.087 (0.095)	−0.064 (0.123)	0.008
JPY	0.002** (0.001)	0.012 (0.019)	−0.021 (0.038)	−0.053 (0.049)	0.019 (0.067)	−0.116* (0.069)	0.147* (0.082)	0.123 (0.095)	0.122 (0.104)	0.062 (0.135)	0.020
NOK	−0.001 (0.001)	0.026 (0.019)	−0.004 (0.037)	−0.034 (0.047)	−0.016 (0.065)	−0.053 (0.067)	0.147* (0.079)	−0.093 (0.093)	0.072 (0.101)	−0.130 (0.131)	0.016
NZD	−0.001 (0.001)	0.033 (0.021)	−0.023 (0.040)	0.011 (0.052)	−0.032 (0.072)	0.075 (0.074)	0.203*** (0.088)	−0.096 (0.102)	−0.096 (0.111)	−0.108 (0.145)	0.021
SEK	−0.001 (0.001)	0.033* (0.019)	−0.010 (0.036)	0.050 (0.047)	0.001 (0.065)	−0.069 (0.067)	0.180*** (0.079)	−0.152 (0.093)	0.075 (0.100)	−0.050 (0.131)	0.024

The table reports the results of regressing individual currency returns on the first lag of each of the nine principal components. For each currency, the table reports the estimated coefficient for each principal component, the standard error in parentheses and the R² coefficient. The sample period is 01/1974 to 12/2020.

Table 4
Maximally Predictable Portfolio (In-Sample).

		AUD	CAD	CHF	EUR	GBP	JPY	NOK	NZD	SEK	
Panel A: Maximally Predictive Portfolio											
Weights		−134.9%	35.2%	−200.9%	39.6%	−32.5%	42.2%	38.2%	112.0%	201.1%	
Regression	Const	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	R ²
Parameter	−0.006***	0.137***	−0.030	0.106	0.014	−0.149	−0.523***	0.538***	0.208	−0.203	0.069
St Error	(0.002)	(0.036)	(0.069)	(0.089)	(0.122)	(0.127)	(0.150)	(0.175)	(0.190)	(0.248)	
Panel B: Maximally Predictive Portfolio (Long-Only)											
Weights		0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	30.9%	69.1%	
Regression	Const	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	R ²
Parameter	−0.001	0.033*	−0.014	0.038	−0.009	−0.024	−0.187***	0.134	0.022	−0.068	0.026
St Error	(0.001)	(0.017)	(0.033)	(0.043)	(0.059)	(0.061)	(0.072)	(0.084)	(0.091)	(0.119)	
Panel C: Equal-Weighted Portfolio											
Weights		11.1%	11.1%	11.1%	11.1%	11.1%	11.1%	11.1%	11.1%	11.1%	
Regression	Const	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	R ²
Parameter	0.000	0.016	−0.009	−0.001	−0.020	−0.023	−0.107*	0.045	0.017	−0.048	0.011
St Error	(0.001)	(0.014)	(0.027)	(0.034)	(0.047)	(0.049)	(0.058)	(0.067)	(0.073)	(0.095)	

The table reports the in-sample regression results for the MPP (Panel A), the long-only MPP (Panel B) and the equal-weighted portfolio (Panel C). In each case, the panel reports the portfolio weights for the nine currencies and the estimated parameters, standard errors and R² coefficient from a regression of the portfolio on the nine lagged principal components. The sample period is 01/1974 to 12/2020.

Table 5
Out-of-Sample Performance (Full Sample).

	MPP (Long Only, Max Weight)				MPP (Max Leverage)				Mom	EW
	25%	50%	75%	100%	25%	50%	75%	100%		
$\hat{\mu}$	0.23%	0.18%	0.16%	0.17%	0.21%	0.20%	0.18%	0.18%	0.12%	0.09%
$\hat{\sigma}$	2.25%	2.37%	2.55%	2.63%	2.71%	3.02%	3.43%	3.88%	1.66%	2.25%
SR	0.350	0.269	0.223	0.228	0.272	0.232	0.184	0.160	0.255	0.139
MD	21.3%	25.0%	23.3%	23.4%	23.3%	26.4%	37.5%	47.7%	27.2%	32.8%
TV	1.919	1.656	1.531	1.566	1.767	1.658	1.492	1.403	1.421	1.233
ϕ_{EW}	1.66%	1.11%	0.82%	0.91%	1.37%	1.14%	0.77%	0.57%	0.50%	−
ϕ_{Mom}	1.15%	0.60%	0.32%	0.41%	0.86%	0.64%	0.27%	0.07%	−	−

The table reports the out-of-sample performance of the long-only MPP with maximum weight constraints of 25%, 50%, 75% and 100%, the MPP with maximum leverage constraints of 25%, 50%, 75% and 100%, the equal-weighted portfolio of MA(2) momentum trading strategies, and the equal-weight (EW) portfolio. For each portfolio, the table reports the mean monthly return ($\hat{\mu}$), the standard deviation of monthly returns ($\hat{\sigma}$), the annualized out-of-sample Sharpe ratio (SR), the maximum drawdown over the sample (MD), the terminal value of the cumulative return starting with 1 unit invested at the beginning of the sample (TV), and the annualized performance fee for a mean–variance investor with a coefficient of relative risk aversion equal to 5 with respect to the EW portfolio, (ϕ_{EW}) and the momentum portfolio (ϕ_{Mom}). Results are reported for the out-of-sample period 01/1994 to 12/2020.

Table 6
Out-of-Sample Performance (Sub-Samples).

	MPP (Long Only, Max Weight)				MPP (Max Leverage)				Mom	EW
	25%	50%	75%	100%	25%	50%	75%	100%		
Panel A: Sub-Sample 1 (01/1994 to 12/2008)										
$\hat{\mu}$	0.21%	0.12%	0.11%	0.14%	0.10%	0.09%	-0.04%	-0.02%	0.33%	0.13%
$\hat{\sigma}$	2.05%	2.28%	2.47%	2.56%	2.72%	3.04%	3.46%	3.96%	1.63%	2.12%
SR	0.348	0.186	0.148	0.189	0.129	0.107	-0.036	-0.021	0.704	0.206
MD	21.3%	25.0%	23.3%	23.4%	23.3%	25.8%	31.8%	40.8%	8.0%	32.6%
TV	1.392	1.189	1.144	1.211	1.123	1.090	0.843	0.833	1.769	1.204
ϕ_{EW}	0.98%	-0.08%	-0.32%	0.06%	-0.44%	-0.62%	-2.30%	-2.34%	2.59%	-
ϕ_{Mom}	-1.58%	-2.61%	-2.85%	-2.48%	-2.97%	-3.14%	-4.78%	-4.82%	-	-
Panel B: Sub-Sample 2 (01/2009 to 12/2020)										
$\hat{\mu}$	0.25%	0.26%	0.24%	0.22%	0.35%	0.34%	0.45%	0.43%	-0.14%	0.05%
$\hat{\sigma}$	2.48%	2.49%	2.66%	2.73%	2.71%	2.99%	3.39%	3.78%	1.67%	2.40%
SR	0.354	0.363	0.309	0.273	0.449	0.389	0.464	0.396	-0.287	0.065
MD	17.3%	13.9%	15.9%	16.1%	15.7%	18.4%	26.4%	27.7%	26.4%	32.8%
TV	1.378	1.393	1.338	1.293	1.573	1.521	1.771	1.684	0.803	1.024
ϕ_{EW}	2.51%	2.60%	2.26%	1.97%	3.65%	3.38%	4.71%	4.30%	-2.03%	-
ϕ_{Mom}	4.63%	4.72%	4.38%	4.08%	5.80%	5.51%	6.88%	6.45%	-	-

The table reports the out-of-sample performance of the long-only MPP with maximum weight constraints of 25%, 50%, 75% and 100%, the MPP with maximum leverage constraints of 25%, 50%, 75% and 100%, the equal-weighted portfolio of MA(2) momentum trading strategies, and the equal-weight (EW) portfolio. For each portfolio, the table reports the mean monthly return ($\hat{\mu}$), the standard deviation of monthly returns ($\hat{\sigma}$), the annualized out-of-sample Sharpe ratio (SR), the maximum drawdown over the sample (MD), the terminal value of the cumulative return starting with 1 unit invested at the beginning of the sample (TV), and the annualized performance fee for a mean-variance investor with a coefficient of relative risk aversion equal to 5 with respect to the EW portfolio, (ϕ_{EW}) and the momentum portfolio (ϕ_{Mom}). Results are reported for the two out-of-sample periods, 01/1994 to 12/2008 (Panel A) and 01/2009 to 12/2020 (Panel B).

1.293 to 1.771, compared with 1.024 for the equal-weight portfolio and 0.803 for the momentum portfolio. The performance fee with respect to both the equal-weight and momentum portfolios is positive in all cases, reaching 4.71% per year with respect to the momentum portfolio and 6.88% per year with respect to the equal-weight portfolio. Thus, while the MPP underperforms the momentum portfolio before 2008, it very significantly outperforms it since then. The performance of the MPP in the post-2008 sub-sample is especially notable in light of the general deterioration in the profitability of foreign exchange trading strategies—including both momentum and carry—since the financial crisis.⁹

5. Robustness Tests

We undertake a number of alternative approaches to the implementation of the empirical analysis in order to establish the robustness of our results. We report the results of these tests in this section.

5.1. Alternative estimation window lengths

The main out-of-sample analysis is conducted with an estimation window length of 240 months, as in Lo and MacKinlay (1997). Table 6 reports the results for the full out-of-sample evaluation period (February 1994 to December 2020) using estimation window lengths of 120 months (Panel A) and 60 months (Panel B). The use of shorter estimation samples generally tends to reduce both the mean return and the volatility of the MPP, and reduce the Sharpe ratio, although there are exceptions. For example, with the 120-month estimation window, increasing the level of leverage tends to increase the mean return of the MPP much more than it increases volatility, and so the Sharpe ratio increases. High levels of leverage are associated with very high terminal portfolio values and performance fees with respect to both the equal-weight and momentum portfolios. For the 60-month estimation window, increasing leverage decreases the Sharpe ratio and the performance fee. Generally, our findings in the previous section appear to be robust with respect to the estimation window length. In almost all cases, the MPP has a higher Sharpe ratio and terminal value than the equal-weight portfolio and, for lower levels of leverage, higher also than the momentum portfolio. In unreported results for the two sub-samples, the findings using the shorter estimation samples are also similar to those obtained with a 240-month estimation sample.

⁹ In unreported results, we also examined the performance of the MPP, EW and Momentum portfolios excluding the periods of the Asian financial crisis (Jul 1997-Dec 1998) and the global financial crisis (Apr 2007-Dec 2008). The Asian financial crisis has little impact on our results since the only Asian currency in our sample is the JPY, which was much less affected by the crisis than the other, emerging market Asian currencies. The global financial crisis has a very significant impact on our results since it is well known (and very evident from Fig. 1 above) that momentum trading strategies performed especially well in the second half of 2008, during which time the USD strengthened and so the EW portfolio of G10 currencies performed poorly. The MPP in contrast is much less affected by the global financial crisis. Excluding these two periods from the sample, the Sharpe ratios of the long only MPP with 25% maximum weight, the momentum portfolio and the EW portfolio are 0.277, 0.128 and 0.227, respectively.

5.2. Alternative portfolio weight constraints

In the main analysis we used constraints on the weights of the individual currencies in the MPP, including symmetric leverage constraints and, in the long only portfolio, maximum weight constraints, in each case applied to the individual currency positions. As a robustness check, we apply constraints on the norm of the portfolio weights. First, we impose the 1-norm constraint:

$$\|\mathbf{w}\|_1 \leq k \quad (22)$$

where $\|\mathbf{w}\|_1 = \sum_{i=1}^n |w_i|$ and $k \in \{1, 2, 3, 4, 5\}$. Second, we impose the 2-norm constraint:

$$\|\mathbf{w}\|_2 \leq l \quad (23)$$

where $\|\mathbf{w}\|_2 = (\sum_{i=1}^n w_i^2)^{0.5}$ and $l \in \{0.75, 1.00, 1.25, 1.50, 1.75, 2.00\}$. These constraints are motivated by [DeMiguel, Garlappi, Nogales and Uppal \(2009b\)](#) who develop the use of norm-constrained portfolio optimization in the mean-variance framework in order to reduce the impact of estimation error in the optimization inputs. The values of k and l that we choose are ad hoc but they span the average 1-norm and 2-norm values of the portfolio weights that arise from the leverage constraints defined above.¹⁰ The case of $k = 1$ corresponds to the short selling constraint ($a = 0\%$). The substantive difference between the norm constraint and the leverage constraint is that the latter applies a constraint on each currency in the portfolio, while the former applies a constraint on the total leverage across all currencies, which may be a more realistic restriction in practice. The results using the 1-norm constraint are reported in Panel A of [Table 7](#) and those using the 2-norm constraint are reported in Panel B. Using constraints on portfolio weights in aggregate rather than individually does not significantly affect the performance of the MPP, and the conclusions drawn are similar to those using the maximum weight and maximum leverage constraints. In particular, for all values of the 1-norm parameter k and the 2-norm parameter l , the Sharpe ratio and the terminal value of the MPP are higher than those of both the equal-weight portfolio and the momentum portfolio, and the performance fees are positive in all cases. Indeed, the norm-constrained portfolios seem to outperform the portfolios with conventional leverage constraints. Moreover, the 2-norm constrained MPP portfolios generally perform better than the 1-norm constrained portfolios. In unreported results for the two sub-samples, the findings are again similar to those obtained using the individual leverage constraints in terms of the relative performance of the MPP and the equal-weight and momentum benchmarks.

5.3. Transaction costs

The costs of buying and selling foreign exchange are among the lowest of any financial asset, and the currency pairs that we use—all involving the USD—are some of the most cheaply traded of all. For the exchange rates in our sample, the average percentage transaction cost is currently about 0.015%, ranging from about 0.010% for EUR/USD to about 0.025% for USD/NOK. Trading costs for the MPP arise from three sources: (a) changes in the buy/sell signal, (b) changes in the portfolio weights as a result of re-estimation of the MPP and (c) rebalancing required to compensate for relative changes in exchange rates. By far the largest of these is changes in the buy/sell signal, since this involves swapping the sign of the position in each currency in the portfolio. The costs associated with changes in the buy/sell signal increase as the maximum weight and maximum leverage constraints are relaxed since this increases the average absolute position size. Changes in portfolio weights as a result of re-estimation of the MPP are generally small, since the composition of the MPP typically changes quite slowly over time, although occasionally there are step changes in the portfolio weights. The rebalancing required to maintain portfolio weights in the face of exchange rate changes is negligible relative to the first two sources of portfolio weight changes. To examine the robustness of our results to the transaction costs arising from such portfolio changes, [Table 8](#) reports the performance of the MPP over the full sample for two levels of transaction costs: a lower level of 0.02%, which would comfortably exceed the true current cost of implementing the MPP strategy (Panel A) and a higher level of 0.04%, which conservatively allows for the fact that transaction costs have reduced over the sample period (Panel B). In each case, these are applied to the total absolute change in the portfolio weights each month, which we compute for portfolio j as:

$$\Delta w_j = \sum_{t=\tau+1}^T \mathbf{1}' |\mathbf{w}_{j,t} - \mathbf{w}_{j,t+1}| \quad (24)$$

where $\mathbf{w}_{j,t}$ is the weight vector of portfolio j at time t , $\mathbf{w}_{j,t+1} = \mathbf{w}_{j,t} \odot (1 + \mathbf{R}_t) / (1 + R_{j,t+1})$ is the weight vector before rebalancing at time t . The table also reports the corresponding results for the equal weight and momentum portfolios. Note that for the equal-weight portfolio, transaction costs are close to zero, since the weights are constant and so the only rebalancing arises from changes in relative exchange rates. For the long-only MPP, the mean return is reduced by approximately the amount of the transaction cost, irrespective of the maximum weight constraint. This is because the MPP trading strategy generates a change in the buy/sell signal in approximately half of the months, and each time the sign of the position in the MPP changes, the 2-way change in weights is equal to 200%. In the remaining months, the changes in the weights—

¹⁰ In particular, the leverage constraints $a \in \{0\%, 25\%, 50\%, 75\%, 100\%\}$ yield average 1-norm values in the out-of-sample analysis of 1.00 (by construction), 2.77, 3.62, 4.15 and 4.86, respectively, and average 2-norm values of 0.72, 1.17, 1.49, 1.71 and 2.00, respectively.

Table 7
Robustness Tests: Estimation Window Length.

	MPP (Long Only, Max Weight)				MPP (Max Leverage)			
	25%	50%	75%	100%	25%	50%	75%	100%
Panel A: 120-Month Estimation Window								
$\hat{\mu}$	0.13%	0.11%	0.07%	0.05%	0.15%	0.33%	0.38%	0.42%
$\hat{\sigma}$	2.20%	2.18%	2.24%	2.27%	2.52%	2.98%	3.36%	3.73%
SR	0.205	0.181	0.112	0.081	0.204	0.387	0.396	0.389
MD	21.7%	19.3%	24.2%	25.1%	22.4%	22.7%	30.9%	32.2%
TV	1.409	1.338	1.167	1.093	1.459	2.540	2.887	3.083
ϕ_{EW}	0.49%	0.30%	-0.21%	-0.45%	0.64%	2.76%	3.27%	3.55%
ϕ_{Mom}	-0.01%	-0.20%	-0.71%	-0.95%	0.14%	2.25%	2.76%	3.04%
Panel B: 60-Month Estimation Window								
$\hat{\mu}$	0.11%	0.19%	0.15%	0.14%	0.18%	0.10%	0.11%	0.05%
$\hat{\sigma}$	2.28%	2.36%	2.40%	2.45%	2.63%	3.02%	3.43%	3.84%
SR	0.173	0.276	0.212	0.193	0.236	0.114	0.110	0.044
MD	21.2%	29.4%	29.8%	29.4%	21.5%	31.1%	36.2%	33.0%
TV	1.328	1.678	1.464	1.410	1.594	1.191	1.175	0.920
ϕ_{EW}	0.28%	1.16%	0.65%	0.51%	0.98%	-0.09%	-0.11%	-0.98%
ϕ_{Mom}	-0.22%	0.65%	0.14%	0.01%	0.47%	-0.59%	-0.61%	-1.47%

The table reports the out-of-sample performance of the long-only MPP with maximum weight constraints of 25%, 50%, 75% and 100%, the MPP with maximum leverage constraints of 25%, 50%, 75% and 100%, using an estimation window length of 120 months (Panel A) and 60 months (Panel B). For each portfolio, the table reports the mean monthly return ($\hat{\mu}$), the standard deviation of monthly returns ($\hat{\sigma}$), the annualized out-of-sample Sharpe ratio (SR), the maximum drawdown over the sample (MD), the terminal value of the cumulative return starting with 1 unit invested at the beginning of the sample (TV), and the annualized performance fee for a mean-variance investor with a coefficient of relative risk aversion equal to 5 with respect to the EW portfolio, (ϕ_{EW}) and the momentum portfolio (ϕ_{Mom}). Results are reported for the out-of-sample period 01/1994 to 12/2020.

Table 8
Robustness Tests: Norm Constraints.

	1	2	3	4	5
	Panel A: Maximum 1-Norm				
$\hat{\mu}$	0.17%	0.22%	0.32%	0.31%	0.30%
$\hat{\sigma}$	2.63%	2.79%	3.26%	3.71%	4.28%
SR	0.228	0.272	0.343	0.286	0.245
MD	23.4%	28.6%	29.3%	24.2%	43.8%
TV	1.566	1.787	2.393	2.161	1.991
ϕ_{EW}	0.91%	1.41%	2.54%	2.18%	1.91%
ϕ_{Mom}	0.41%	0.91%	2.03%	1.68%	1.40%
	0.75	1.00	1.25	1.50	1.75
Panel B: Maximum 2-Norm					
$\hat{\mu}$	0.27%	0.25%	0.30%	0.31%	0.35%
$\hat{\sigma}$	2.44%	2.63%	2.98%	3.29%	3.98%
SR	0.378	0.335	0.349	0.323	0.303
MD	15.8%	20.9%	20.4%	21.1%	29.1%
TV	2.150	2.034	2.287	2.269	2.425
ϕ_{EW}	2.09%	1.89%	2.35%	2.34%	2.62%
ϕ_{Mom}	1.58%	1.38%	1.84%	1.83%	2.10%

The table reports the out-of-sample performance of the norm constrained portfolio with maximum 1-norm values of 1, 2, 3, 4 and 5 (Panel A) and maximum 2-norm values of 0.75, 1.00, 1.25, 1.50, 1.75 and 2.00 (Panel B). For each portfolio, the table reports the mean monthly return ($\hat{\mu}$), the standard deviation of monthly returns ($\hat{\sigma}$), the annualized out-of-sample Sharpe ratio (SR), the maximum drawdown over the sample (MD), the terminal value of the cumulative return starting with 1 unit invested at the beginning of the sample (TV), and the annualized performance fee for a mean-variance investor with a coefficient of relative risk aversion equal to 5 with respect to the EW portfolio, (ϕ_{EW}) and the momentum portfolio (ϕ_{Mom}). Results are reported for the out-of-sample period 01/1994 to 12/2020.

and also the transaction costs—are negligible, and hence the average total cost is about 100% multiplied by the nominal transaction cost. With leverage, the 2-way weight changes become much greater and so the reduction in the mean return is correspondingly larger. For example, with 0.02% transaction costs, the mean return of the 100% leverage portfolio is reduced from 0.20% to 0.10% (an average effective cost of 0.10%). With 0.04% transaction costs, the mean return is reduced to zero (an average effective cost of 0.20%). Since the transaction costs have little effect on the volatility, the Sharpe ratio of the MPP also falls further the higher the level of leverage. For the momentum portfolio, transaction costs have a noticeable effect on performance, but less so than for the MPP owing to the fact that position sizes are typically smaller. Nevertheless, with more moderate levels of leverage, the MPP still outperforms both the equal-weight and momentum portfolios, with a higher Sharpe ratio and higher terminal value. Indeed, even at the higher level of transaction costs, the MPP has higher terminal value than both the equal-weight and momentum portfolios in the long only case (both with and without maximum weight constraints) and in the leveraged case with 25% maximum leverage. The long only MPP has a positive performance fee with respect to both the equal-weight and momentum portfolios in all cases.

Table 9
Robustness Tests: Transaction Costs.

	MPP (Long Only, Max Weight)				MPP (Max Leverage)				Mom	EW
	25%	50%	75%	100%	25%	50%	75%	100%		
Panel A: Transaction Costs 0.02%										
$\hat{\mu}$	0.21%	0.16%	0.14%	0.15%	0.16%	0.13%	0.10%	0.08%	0.10%	0.09%
$\hat{\sigma}$	2.25%	2.38%	2.55%	2.63%	2.71%	3.01%	3.43%	3.88%	1.66%	2.25%
SR	0.318	0.239	0.194	0.201	0.200	0.147	0.099	0.072	0.213	0.138
MD	22.7%	26.4%	24.6%	24.8%	26.1%	31.8%	43.1%	53.3%	29.3%	32.8%
TV	1.796	1.552	1.431	1.464	1.473	1.307	1.136	1.021	1.331	1.231
ϕ_{EW}	1.41%	0.87%	0.57%	0.66%	0.69%	0.26%	-0.24%	-0.60%	0.26%	-
ϕ_{Mom}	1.15%	0.61%	0.31%	0.40%	0.42%	-0.01%	-0.50%	-0.86%	-	-
Panel B: Transaction Costs 0.04%										
$\hat{\mu}$	0.19%	0.14%	0.12%	0.13%	0.10%	0.05%	0.01%	-0.02%	0.08%	0.09%
$\hat{\sigma}$	2.25%	2.38%	2.55%	2.63%	2.72%	3.01%	3.43%	3.88%	1.67%	2.25%
SR	0.287	0.210	0.166	0.173	0.128	0.062	0.013	-0.015	0.170	0.137
MD	24.0%	27.7%	26.0%	26.1%	30.0%	36.8%	49.5%	59.2%	31.4%	32.9%
TV	1.681	1.455	1.337	1.368	1.228	1.029	0.864	0.743	1.246	1.229
ϕ_{EW}	1.17%	0.63%	0.33%	0.42%	0.02%	-0.62%	-1.24%	-1.76%	0.03%	-
ϕ_{Mom}	1.15%	0.61%	0.30%	0.39%	-0.01%	-0.65%	-1.27%	-1.79%	-	-

The table reports the out-of-sample performance net of transaction costs of the long-only MPP with maximum weight constraints of 25%, 50%, 75% and 100%, the MPP with maximum leverage constraints of 25%, 50%, 75% and 100%, the equal-weighted portfolio and the equal-weight (EW) portfolio of MA(2) momentum trading strategies. For each portfolio, the table reports the mean monthly return ($\hat{\mu}$), the standard deviation of monthly returns ($\hat{\sigma}$), the annualized out-of-sample Sharpe ratio (SR), the maximum drawdown over the sample (MD), the terminal value of the cumulative return starting with 1 unit invested at the beginning of the sample (TV), and the annualized performance fee for a mean-variance investor with a coefficient of relative risk aversion equal to 5 with respect to the EW portfolio, (ϕ_{EW}) and the momentum portfolio (ϕ_{Mom}). Results are reported for the out-of-sample period 01/1994 to 12/2020.

6. Conclusion

In this paper, we investigate the maximally predictable portfolio (MPP) of the G10 currencies with respect to lagged currency returns from the perspective of a U.S. investor. The MPP, which is defined as the portfolio with the highest R^2 between the portfolio return and its predicted value, provides an upper bound on predictability and, when viewed from an in-sample perspective, provides a benchmark against which to evaluate other forecasting models that condition on the same predictor variables. From an out-of-sample perspective, however, it provides a systematic framework for optimally exploiting the information content of the predictor variables—in this case, lagged returns—in order to predict future portfolio returns as part of a dynamic investment strategy. By using lagged currency returns as conditioning variables, our implementation of the MPP approach can be thought of as a generalization of the momentum/reversals strategy. The momentum/reversals strategy, however, relies on autocorrelations of the returns of individual currencies (or portfolios of currencies) in isolation, ignoring potentially useful information in the lagged returns of other currencies. In contrast, by conditioning the portfolio return on the returns of all individual currencies, the MPP exploits significant cross-correlations in returns that are useful for the prediction of individual currency returns, and hence also for the prediction of portfolio returns. We show that the MPP has superior out-of-sample performance relative to both the passive equal-weighted portfolio of currencies and an equal-weighted portfolio of momentum trading strategies. Our results are robust to the choice of estimation window length, portfolio weight constraints and transaction costs.

There are a number of ways that this research could be extended. First, the trading strategy used in this paper is based on a buy/sell signal that depends only on the sign of the forecast return of the MPP, and not on its magnitude. A natural approach would be to use a threshold for the forecast return to enhance the strength of the trading signal and to reduce the number of trades and hence minimize transaction costs. Second, the MPP implemented in this paper uses all nine principal components (or, equivalently, all nine exchange rates) as predictor variables. It is well known that parsimonious models tend to work better for forecasting and so the use of a smaller number of conditioning variables is likely to help trading strategy performance. Finally, we implement the MPP using the R^2 coefficient between the actual and predicted portfolio return as the measure of predictability. An alternative approach, and one that would be more consistent with the way that the MPP forecasts are used in the trading strategy, would be to measure predictability by the accuracy of the forecast sign of the MPP return rather than its level.

CRedit authorship contribution statement

Richard D.F. Harris: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing. **Jian Shen:** Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing. **Fatih Yilmaz:** Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We are grateful for very helpful comments from George Bulkley, Nick Taylor and participants of the University of Bristol Financial Markets Research Group.

References

- Barroso, P., Santa-Clara, P., 2015. Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis* 50 (5), 1037–1056.
- Bewley, R., Orden, D., Yang, M., Fisher, L.A., 1994. Comparison of Box-Tiao and Johansen canonical estimators of cointegrating vectors in VEC (1) models. *Journal of Econometrics* 64 (1–2), 3–27.
- Box, G.E.P., Tiao, G.C., 1977. A canonical analysis of multiple time series. *Biometrika* 64 (2), 355–365.
- Burnside, C., Eichenbaum, M., Rebelo, S., 2008. Carry trade: The gains of diversification. *Journal of the European Economic Association* 6 (2–3), 581–588.
- Cheung, Y.-W., He, J., Ng, L.K., 1997. Common predictable components in regional stock markets. *Journal of Business and Economic Statistics* 15, 35–42.
- Cheung, Y.-W., Chinn, M.D., Pascual, A.G., 2005. Empirical exchange rate models of the nineties: Are any fit to survive?. *Journal of International Money and Finance* 24 (7), 1150–1175.
- d'Aspremont, A., 2011. Identifying small mean-reverting portfolios. *Quantitative Finance* 11 (3), 351–364.
- Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. *Review of Financial Studies* 22 (9), 3491–3530.
- DeMiguel, V., Garlappi, L., Uppal, R., 2009a. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?. *The Review of Financial Studies* 22 (5), 1915–1953.
- DeMiguel, V., Garlappi, L., Nogales, F.J., Uppal, R., 2009b. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science* 55 (5), 798–812.
- Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *The Journal of Finance* 56 (1), 329–352.
- Foster, F.D., Smith, T.O.M., Whaley, R.E., 1997. Assessing goodness-of-fit of asset pricing models: The distribution of the maximal R^2 . *Journal of Finance* 52 (2), 591–607.
- Harris, R.D.F., Yilmaz, F., 2009. A momentum trading strategy based on the low frequency component of the exchange rate. *Journal of Banking and Finance* 33 (9), 1575–1585.
- Invesco (2020) DB G10 Currency Harvest Fund. <https://www.invesco.com/us/financial-products/etfs/product-detail?audienceType=Investor&ticker=DBV> (accessed 9 April 2021).
- Johansen, S., 1988. Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control* 12 (2–3), 231–254.
- LeBaron, B., 1999. Technical trading rule profitability and foreign exchange intervention. *Journal of International Economics* 49 (1), 125–143.
- Levich, R.M., Thomas, L.R., 1993. The significance of technical trading rule profits in the foreign exchange market: A bootstrap approach. *Journal of International Money and Finance* 12 (5), 451–474.
- Lo, A.W., MacKinlay, A.C., 1997. Maximizing predictability in the stock and bond markets. *Macroeconomic Dynamics* 1 (1), 102–134.
- Lustig, H., Roussanov, N., Verdelhan, A., 2011. Common risk factors in currency markets. *Review of Financial Studies* 24 (11), 3731–3777.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012a. Carry trades and global foreign exchange volatility. *Journal of Finance* 67, 681–718.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. *Journal of Financial Economics* 106 (3), 660–684.
- Okunev, J., White, D., 2003. Do momentum-based strategies still work in foreign currency markets?. *Journal of Financial and Quantitative Analysis* 38 (2), 425–447.
- Sweeney, R.J., 1986. Beating the foreign exchange market. *Journal of Finance* 41 (1), 163–182.
- Taylor, M.P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance* 11 (3), 304–314.