Review of Finance, 2023, 223–246
https://doi.org/10.1093/rof/rfac014
Advance Access Publication Date: 7 March 2022

OXFORD

Market Timing and Predictability in FX Markets

Thomas A. Maurer[1], Thuy-Duong Tô[2], and Ngoc-Khanh Tran[3]

[1]Faculty of Business and Economics, The University of Hong Kong, Hong Kong, [2]UNSW Business School, The University of New South Wales, Australia, and [3]Pamplin College of Business, Virginia Tech, USA

Abstract

We study the economic value of market timing in foreign exchange (FX) markets, that is, using information about the conditional Sharpe ratio to adjust the notional value of a conditionally mean–variance efficient currency portfolio. Our strategy trades more (less) aggressively when the conditional risk-return trade-off is more (less) favorable. This leads to a significant improvement in the out-of-sample unconditional Sharpe ratio, skewness, and maximum drawdown per 1% expected excess return. The strategy's market timing predicts returns, volatility, and skewness in FX markets. Popular currency pricing factors do not explain the strategy's high average excess returns. Our findings suggest that it is costly to impose leverage or risk (i.e., conditional volatility) limits or other inferior market timing policies when constructing currency trading strategies.

Keywords: Market timing, Leverage limits, Risk limits, FX markets, Mean–variance optimization, Exchange rates, Estimation errors

JEL classification: F31, F37, G11, G12, G15, G17

Received March 10, 2020; accepted February 13, 2022 by Editor Amit Goyal.

1. Introduction

Mean–variance optimized portfolios earn high out-of-sample (OOS) returns in foreign exchange (FX) markets (Baz *et al.*, 2001; Della Corte, Sarno, and Tsiakas, 2009; Ackermann, Pohl, and Schmedders, 2016; Daniel, Hodrick, and Lu, 2017). We investigate the time-series performance of these portfolios and show that the conditional Sharpe ratio and crash risks are time-varying and predictable. A mean–variance optimized currency strategy that exploits this information, referred to as MV hereafter, trades more (less) aggressively when prices of risks in FX markets are higher (lower). This market timing is valuable and significantly improves the performance compared with other mean–variance optimized strategies that use inferior market timing policies. MV has a desirable unconditional return distribution with high average excess returns, a high unconditional Sharpe ratio, positive skewness, and low maximum drawdown (MDD) per 1% expected excess return.

The performance of MV has important implications for asset allocations, pricing risks, and parameter estimations in FX markets. Regarding asset allocations, our results suggest that leverage or risk limits (i.e., targeting a constant notional value or volatility) are costly when market conditions are time-varying and market timing is highly profitable. In practice, such limits are often imposed on portfolio managers (in addition to less tight regulatory constraints). Thus, it is important to understand the implicit costs associated with such leverage or risk (i.e., conditional volatility) limits. Regarding the pricing of risks, our paper shows that the high excess returns generated by the market timing of MV cannot be explained by known currency risk factors, indicating the existence of important FX market risk sources that are not yet well understood. Regarding the parameter estimation, our paper shows that a combination of forward discounts (which serve as useful proxies for conditional expected excess returns of currencies) and principal component analysis (PCA, which is useful to obtain a robust estimate of the exchange rate covariance matrix) are key elements for MV's superior OOS performance.

All conditionally mean–variance efficient strategies have proportional portfolio weights, that is, they invest in the same risky asset portfolio. Accordingly, their conditional Sharpe ratios are identical. However, these strategies generally differ with respect to the asset allocation, that is, the time-variation in the notional value or leverage. We denote this time-variation as market timing. We analyze the differences in the market timing of popular mean–variance optimized strategies. In FX markets, Baz *et al.* (2001) minimizes the conditional portfolio variance subject to a constant target expected return. Della Corte, Sarno, and Tsiakas (2009) maximize the conditional expected return subject to a constant target volatility. Ackermann, Pohl, and Schmedders (2016) and Daniel, Hodrick, and Lu (2017) maximize the conditional Sharpe ratio and rescale the portfolio such that the notional value is equal to one. Note that the value of the conditional Sharpe ratio varies over time and these portfolios differ with regard to the time-variation in the notional value (or leverage). This means that they place different weights on the conditional Sharpe ratio. As a result, they have different unconditional Sharpe ratios.

Our main focus is the strategy MV that maximizes $\mu_\mathrm{p} - \frac{\vartheta}{2}\sigma_\mathrm{p}^2$, where μ_p and σ_p^2 are the conditional expected excess return and variance of the currency portfolio, and ϑ is a time-invariant parameter. The portfolio weights of MV are $\theta_t^{\mathrm{MV}} = \vartheta\Omega_t^{-1}\mu_t$, where μ_t and Ω_t are the vector of expected excess returns and the covariance matrix at time t. MV delivers not only the optimal conditional Sharpe ratio but also a time-varying notional value that is high (low) when the conditional Sharpe ratio is high (low). This adjustment increases the contribution of the future Sharpe ratio when the risk adjusted FX market return is expected to be high and suppresses the contribution when the risk adjusted return is expected to be low.

Opposite to MV the approach of Baz *et al.* (2001) implies a low (high) notional value when the conditional expected return is high (low) and the conditional Sharpe ratio tends to be high (low). Moreover, the notional value is insensitive to changes in the conditional volatility. Accordingly, the strategy invests less (more) aggressive when conditional prices of risk are high (low) and we expect a lower unconditional Sharpe ratio compared with MV. The strategies of Ackermann, Pohl, and Schmedders (2016) and Daniel, Hodrick, and Lu (2017) have no market timing by design (i.e., constant notional value) and we expect a lower unconditional Sharpe ratio compared with MV. In the approach of Della Corte, Sarno, and Tsiakas (2009), the notional value only varies with the conditional return volatility, but not the conditional expected return. We expect that the unconditional Sharpe ratio is higher than that of the strategy with a constant notional value due to volatility timing, but lower than that of MV as MV also uses information about expected returns for the market timing. We verify these conjectures in the data.

The theoretical properties of MV are well known. However, the empirical OOS performance is not documented in FX markets. Moreover, the empirical literature has not analyzed the importance of market timing with regard to MV versus other mean–variance strategies. In stock markets, it is documented that MV can feature extreme portfolio weights due to parameter uncertainty. In turn, this can lead to a bad OOS performance. Accordingly, portfolio constraints (as a kind of shrinkage) can be beneficial to ensure stability in mean–variance optimized portfolios and improve the OOS performance. In this sense, portfolio constraints have offsetting effects, they are beneficial to address estimation error concerns but costly in terms of market timing. An advantage of FX markets is that estimation errors are not a major concern (Baz *et al.*, 2001; Ackermann, Pohl, and Schmedders, 2016; Daniel, Hodrick, and Lu, 2017). Using FX data, we document that portfolio constraints provide no value in terms of stability, but there are large costs due to the suboptimal market timing.

Ferson and Siegel (2001), Penaranda (2016), and Penaranda and Wu (2020) discuss theoretical properties of mean–variance optimized portfolios and the unconditional return distribution. These are theory papers with interesting qualitative results but the quantitative effect in the data is not studied. Our paper is the first to analyze this difference and quantify it using FX market data.

A common concern is that a mean–variance optimization ignores jump risks and may be subject to large crash or downside risk (Brunnermeier, Nagel, and Pedersen, 2008; Dobrynskaya, 2014, 2015; Lettau, Maggiori, and Weber, 2014; Chernov, Graveline, and Zviadadze, 2018). For instance, the high-minus-low forward discount sorted carry trade strategy of Lustig and Verdelhan (2007) is subject to significant crashes, exhibiting a negative skewness and significant maximum loss. In contrast, MV's return distribution features a positive skewness and significantly lower maximum loss per 1% expected excess return. Thus, MV's high average excess returns do not appear to be a compensation for downside risk. Since a mean–variance optimization with market timing increases average returns and reduces the downside risk, it is unlikely that crash risks are of first-order importance to explain risk premia in FX markets. This finding is consistent with the evidence provided by Bekaert and Panayotov (2018), Daniel, Hodrick, and Lu (2017), and Maurer, To, and Tran (2022).

We further show that the notional value of our strategy predicts future realized returns and the volatility of mean–variance optimized portfolios. This confirms that MV invests more (less) aggressively when the risk-return trade-off is more (less) favorable. We find the predictability in both the first and the second moments of returns, which sets the market timing of our strategy apart from volatility managed portfolios (Fleming, Kirby, and Ostdiek, 2001; Della Corte, Sarno, and Tsiakas, 2009; Moreira and Muir, 2017; Cederburg *et al.*, 2020).

Our asset pricing tests on MV's excess return using popular FX pricing factors produce economically and statistically significant abnormal returns, indicating that MV does not simply invest in a combination of well-known currency risk factors. This result concurs with the finding of Chernov, Dahlquist, and Lochstoer (2021) that unpriced risks constitute a major component of risks in currency markets. Finally, Maurer, To, and Tran (2022) and Chernov, Dahlquist, and Lochstoer (2021) find that mean–variance optimized currency portfolios correctly price the cross-section of FX returns and other popular pricing factors are subsumed. These papers focus on the pricing properties, while our paper analyzes the economic value of market timing based on the time-variation in the conditional Sharpe ratio. We conclude that it is costly to follow suboptimal market timing policies, in particular imposing leverage or risk (i.e., conditional volatility) limits.

Our paper is organized as follows. Section 2 discusses the market timing intuition of our currency portfolio construction and describes FX strategies and the data. Section 3 reports our main empirical results concerning the market timing of MV. Section 4 compares the performance of MV with other popular FX strategies. Section 5 discusses robustness checks. Section 6 concludes. The Appendix and Supplementary (Online) Appendix present further details about strategies and additional tables.

2. Intuition of Market Timing and Currency Strategies

We first discuss the intuition and implications of our FX market timing approach. The intuition provides the motivation for studying various currency strategies.

2.1 Intuitive Aspects of Market Timing

Market timing exploits the time variation in conditional expected returns and risks. It focuses on the capital allocation between the conditionally mean–variance efficient tangency portfolio and the risk-free asset. Hence, market timing dynamically adjusts the leverage or notional value of the conditionally mean–variance efficient strategy with the aim to enhance the unconditional Sharpe ratio. Our specific approach to time the market sets the notional amount equal to the ratio of the conditional expected return and the conditional variance of the strategy. That is, the strategy takes a larger position in the tangency portfolio financed by a larger short position in the risk-free asset (i.e., a larger leverage) when the conditional Sharpe ratio of the tangency portfolio increases.

We now intuitively illustrate the idea of our market timing approach. Consider a currency strategy in a three-date setting $t \in \{\tau, \tau+1, \tau+2\}$. Let rx_t denote the strategy's return in excess of the risk-free rate, $\mu_t \equiv E_t[rx_{t+1}]$ the conditional expected return, and $\sigma_t \equiv \mathrm{Vol}_t[rx_{t+1}]$ the conditional volatility. In the current illustration, the unconditional Sharpe ratio is represented by the two-period Sharpe ratio $\mathrm{SR}_\tau^{(2)}$ of the investment from τ to $\tau+2$. Whereas the conditional Sharpe ratio is represented by one-period Sharpe ratios SR_τ and $\mathrm{SR}_{\tau+1}$. Our market timing is characterized by the time-varying notional values

$$\mathrm{NV}_t = \frac{\mu_t}{\sigma_t^2}, \quad t \in \{\tau, \tau+1\}. \tag{1}$$

Note that the two-period currency excess return realized at $t+2$ is obtained by aggregating the one-period returns scaled by the respective notional amounts $rx_{\tau+2}^{(2)} = \mathrm{NV}_\tau\, rx_{\tau+1} + \mathrm{NV}_{\tau+1}\, rx_{\tau+2}$,[1] implying an intuitive expression for the unconditional Sharpe ratio

$$\begin{aligned} \mathrm{SR}_\tau^{(2)} &= \frac{E_\tau[\mathrm{NV}_\tau\, rx_{\tau+1} + \mathrm{NV}_{\tau+1}\, rx_{\tau+2}]}{\sqrt{\mathrm{Var}_\tau[\mathrm{NV}_\tau\, rx_{\tau+1} + \mathrm{NV}_{\tau+1}\, rx_{\tau+2}]}} \\ &= \frac{E_\tau\left[\frac{\mu_\tau}{\sigma_\tau}\frac{rx_{\tau+1}}{\sigma_\tau} + \frac{\mu_{\tau+1}}{\sigma_{\tau+1}}\frac{rx_{\tau+2}}{\sigma_{\tau+1}}\right]}{\sqrt{\mathrm{Var}_\tau\left[\frac{\mu_\tau}{\sigma_\tau}\frac{rx_{\tau+1}}{\sigma_\tau} + \frac{\mu_{\tau+1}}{\sigma_{\tau+1}}\frac{rx_{\tau+2}}{\sigma_{\tau+1}}\right]}}. \end{aligned} \tag{2}$$

1 The weight NV_t is determined at date t, which scales up the strategy's return rx_{t+1} realized at the next date $t+1$.

Several observations rationalize the choice and the implications of the market timing $\mathrm{NV}_t = \frac{\mu_t}{\sigma_t^2}$. First, while the choice of NV_t does not affect the one-period Sharpe ratio, it crucially matters for the two-period Sharpe ratio. Two strategies may offer identical and optimal one-period Sharpe ratio but different two-period Sharpe ratios. Our market timing preserves the one-period (or conditional) mean–variance efficiency and further increases the two-period (or unconditional) Sharpe ratio.

Second, in the two-period Sharpe ratio (2), the next-period risk-adjusted return $\frac{rx_{t+1}}{\sigma_t}$ is scaled by the current price of risk of the strategy $\frac{\mu_t}{\sigma_t}$, $t \in \{\tau, \tau+1\}$. Therefore, a larger position in the currency strategy is taken when its current price of risk is higher. Intuitively, if the risk-return trade-off is expected to be much more attractive in the first period than in the second, then the market timing strategy invests aggressively in the first period and reduces the risk exposure in the second period to achieve a higher two-period Sharpe ratio. That is, the adjustment helps to boost the two-period Sharpe ratio as it enhances the contribution of the future risk-adjusted return $\frac{rx_{\tau+1}}{\sigma_\tau}$ when this risk-adjusted return is expected to be high (i.e., large $\frac{\mu_t}{\sigma_t}$) and suppresses the contribution when it is expected to be low.

This enhancement and suppression effect can be seen in the following analogous but simplified inequality (which approximates Equation (2) when the serial correlation of the return is insignificant and omitted),

$$\frac{x_1 + x_2}{y_1 + y_2} \le \frac{x_1 \frac{x_1}{y_1} + x_2 \frac{x_2}{y_2}}{y_1 \frac{x_1}{y_1} + y_2 \frac{x_2}{y_2}}.$$

Note that this inequality always holds for any four positive real numbers x_1, x_2, x_3, x_4, because the difference between the right-hand and left-hand sides, $(x_1 + x_2)^{-1}(y_1 + y_2)^{-1}\left(x_1\sqrt{\frac{y_2}{y_1}} - x_2\sqrt{\frac{y_1}{y_2}}\right)^2$, is always positive. Consider the right-hand side: (i) when $\frac{x_1}{y_1} > \frac{x_2}{y_2}$, we are enhancing the dominant pairs x_1, y_1 by multiplying them with the larger ratio $\frac{x_1}{y_1}$ and suppressing the dominated pairs x_2, y_2 with the smaller ratio $\frac{x_2}{y_2}$ and (ii) when $\frac{x_1}{y_1} < \frac{x_2}{y_2}$, we are enhancing the dominant pairs x_2, y_2 by multiplying them with the larger ratio $\frac{x_2}{y_2}$ and suppressing the dominated pairs x_1, y_1 with the smaller ratio $\frac{x_1}{y_1}$. As a result, the right-hand side is always larger than the left-hand side.

Finally, the market timing improvement of the two-period Sharpe ratio tends to be larger when the time variation in the strategy's market price is more significant, that is, when the moments of FX market returns are more volatile. Intuitively, this is because the above enhancement and suppression of the return contribution to the two-period Sharpe ratio requires that the return moments be different across the two dates.[2] Quantitatively, the two-period Sharpe ratio is a convex function of the variation across dates of the strategy's return moments, featuring a significant gain when this variation is sufficiently large. An important testable implication is that the market timing leads to a larger improvement in the unconditional Sharpe ratio when changes in FX market conditions are larger. We discuss OOS tests and present supportive empirical evidence for this key implication in Section 3.3.

2.2 Currency Strategies

Let rx_{t+1} denote the vector of excess returns against the USD of N foreign currencies for the holding period from month t to $t+1$, and θ_t^S the vector of portfolio weights of a

2 When return moments are similar at $\tau+1$ and $\tau+2$, the scaling parameters NV_τ and $\mathrm{NV}_{\tau+1}$ do not matter for the two-period Sharpe ratio, offering no market-timing improvement to $\mathrm{SR}_\tau^{(2)}$.

currency strategy S at the end of month t. The sum of weights $\sum_i \theta^S_{i,t}$ does not need to be equal to 1 as rx_{t+1} are excess returns. The excess return of strategy S is $rx'_{t+1}\theta^S_t$ and the total USD risk exposure is measured by its notional value, $\sum_i |\theta^S_{i,t}|$. Next, we describe the currency strategies in our analysis based on their efficiency and market timing characteristics.

We first introduce our principal strategy, which is conditionally mean–variance efficient and uses market timing based on the conditional Sharpe ratio. Second, we describe other conditionally mean–variance efficient strategies. Finally, we discuss other common strategies in the literature.

2.2.a. Conditional mean–variance efficiency with market timing

Our principal currency strategy, MV, is defined by the portfolio weights

$$\theta^{\mathrm{MV}}_t = \vartheta \Omega_t^{-1} \mu_t,$$

where Ω_t and μ_t are the conditional covariance matrix and vector of expected excess returns of the N currencies against the USD and ϑ is a time invariant scaling factor and can be interpreted as the inverse of the relative risk aversion. θ^{MV}_t maximizes $\mu_{\mathrm{p}} - \frac{\vartheta}{2}\sigma^2_{\mathrm{p}}$, where $\mu_{\mathrm{p}} = \mu'_t\theta^{\mathrm{MV}}_t$ and $\sigma^2_{\mathrm{p}} = \theta\prime^{\mathrm{MV}}_t \Sigma_t \theta^{\mathrm{MV}}_t$ are the conditional expected excess return and volatility of the portfolio. MV features both conditional mean–variance efficiency and market timing. As elaborated intuitively in Equation (1), MV's market timing is characterized by its time-varying notional value, which depends on the size of the conditional expected excess returns and the conditional covariance matrix. Note that ϑ is constant, and thus, does not affect the market timing or the Sharpe ratio of the strategy. In our empirical analysis, we assume the investor believes the average conditional Sharpe ratio of MV is equal to 1 and we choose ϑ such that MV targets an unconditional volatility of 10%.

A practical and important challenge of a mean–variance optimization is the estimation of the conditional covariance matrix and expected excess returns. Estimation errors may lead to a deceptively high in-sample but low OOS Sharpe ratio. For instance, DeMiguel, Garlappi, and Uppal (2009) show that mean–variance optimized portfolios in the US stock market earn low OOS returns and are consistently outperformed by equally weighted portfolios. They attribute this finding to estimation errors. We address the estimation error problem for both expected returns and covariances. The realized excess return of currency i denominated in USD is composed of currency i's forward discount and realized exchange rate growth, $rx_{i,t+1} = f_{i,t} + \Delta s_{i,t+1}$. The empirical finding that short-term exchange rate movements $x_{i,t+1}$ are close to a random walk, $E_t[\Delta s_{i,t+1}] \approx 0$ (Meese and Rogoff, 1983) then implies that forward discounts are good predictors of conditional currency excess returns, $E_t[rx_{i,t+1}] \approx f_{i,t}$. Motivated by this observation, we employ the current forward discount $f_{i,t}$ to proxy for the conditional expected excess return $\mu_{i,t}$ of currency i. Baz *et al.* (2001), Della Corte, Sarno, and Tsiakas (2009), Ackermann, Pohl, and Schmedders (2016), and Daniel, Hodrick, and Lu (2017) employ the same proxy to estimate expected excess returns in currency markets. Furthermore, Lustig, Roussanov, and Verdelhan (2011) document that FX markets have a strong factor structure and the first few PCs capture most of the covariation of FX returns. Accordingly, we employ PCA and retain PCs that explain at least 1% of the common variation in returns and construct a robust version of the covariance matrix Ω_t^{-1}.

In the first step, we use an exponentially weighted moving average (EWMA) popularized by J.P. Morgan's RiskMetrics to estimate daily covariances in an expanding window,

$$\widehat{\Omega}_{d,\tau,ij} = \delta\widehat{\Omega}_{d,\tau-1,ij} + (1-\delta)\Delta s_{d,i,\tau-1}\Delta s_{d,j,\tau-1},$$

where $\Delta s_{d,i,\tau}$ is the daily exchange rate growth of currency i against the USD on day τ, the initial value at $\tau = 1$ is $\widehat{\Omega}_{d,\tau,ij} = \Delta s_{d,i,\tau}\Delta s_{d,j,\tau}$, and EWMA weight $\delta = 0.94$, that is, a half-life of 11 trading days (Fleming, Kirby, and Ostdiek, 2001). Let τ_t be the last day of month t and T_t the number of trading days in month t. We define the monthly estimate $\widehat{\Omega}_{t,ij} = T_t\widehat{\Omega}_{d,\tau_t,ij}$.

In the second (PCA) step, we diagonalize the covariance matrix $\widehat{\Omega}_t = W_t\Lambda_t W_t'$, with $\Lambda_t = \mathrm{Diag}(\lambda_{1,t},\ldots,\lambda_{N,t})$. To generate a robust (inverse) of the covariance matrix $\tilde{\Omega}_t^{-1} = \tilde{W}_t\tilde{\Lambda}_t^{-1}\tilde{W}_t'$, we remove all K eigenvalues $\lambda_{k,t}$ for which $\frac{\lambda_{k,t}}{\sum_{h=1}^{N}\lambda_{h,t}} < \overline{\lambda} = 1\%$, that is, the PCs which explain less than $\overline{\lambda} = 1\%$ of the common variation in exchange rates. $N \times (N-K)$ matrix $\tilde{W}_t$ is obtained by removing the K columns of the $N \times N$ rotation matrix W_t associated with the K removed eigenvalues $\lambda_{k,t}$, and $(N-K)\times(N-K)$ matrix $\tilde{\Lambda}_t$ by removing K eigenvalues from the $N \times N$ matrix Λ_t.

As a result, MV is a mean–variance optimized strategy that is linear in the $(N-K)$ dominant PCs of FX markets. Because the removed PCs explain only a small fraction of the exchange rate covariation, this procedure reduces the covariance matrix estimation errors and alleviates in-sample over-fitting and near-arbitrage opportunities arising from factors with unreasonably large in-sample Sharpe ratio (Ross, 1976; Kozak, Nagel, and Santosh, 2018).

2.2.b. Other conditionally mean–variance efficient strategies

For comparison, we consider several alternative currency strategies that also offer the optimal conditional Sharpe ratio but differ from MV with respect to the market timing, that is, the time-variation in the notional value. All conditionally mean–variance efficient strategies, including MV, have proportional portfolio weights, that is, invest in the same risky asset portfolio. They only differ with respect to the notional value. For the construction of the following strategies, we define the conditional expected excess return and variance of MV, $\mu_t^{MV} = \mu_t'\theta_t^{MV}$, and $\sigma_t^{MV} = \sqrt{\theta_t^{MV\prime}\Omega_t\theta_t^{MV}}$.

MV_{CN} (or notional target) is a conditionally mean–variance efficient strategy whose portfolio weights are scaled from MV's weights to keep the notional value constant, $\theta_t^{MV_{CN}} = \frac{1}{\sum_i|\theta_{i,t}^{MV}|}\theta_t^{MV}$. We choose a notional value of 4 in order to generate an annualized volatility of roughly 10%. The constant notional value requirement of MV_{CN} precludes market timing. This strategy is examined by Ackermann, Pohl, and Schmedders (2016) and Daniel, Hodrick, and Lu (2017). Note that both of these papers use forward discounts to proxy for expected excess returns but they do not employ our PCA approach to address estimation errors in the covariance matrix of currency returns.

MV_{CV} (or volatility target) with $\theta_t^{MV_{CV}} = \frac{\overline{\sigma}}{\sigma_t^{MV}}\theta_t^{MV}$ is a conditionally mean–variance efficient strategy with portfolio weights rescaled to keep the conditional volatility constant and equal to $\overline{\sigma}$. We choose $\overline{\sigma} = 10\%$. MV_{CV} is implemented by Della Corte, Sarno, and Tsiakas (2009) in FX markets.

MV_{CY} (or yield target) with $\theta_t^{MV_{CY}} = \frac{\overline{\mu}}{\mu_t^{MV}}\theta_t^{MV}$ is a conditionally mean–variance efficient strategy with portfolio weights rescaled to keep the conditional expected excess return constant and equal to $\overline{\mu}$. We choose $\overline{\mu} = 10\%$. Baz *et al.* (2001) analyze MV_{CY} in FX markets.

The volatility and yield targets generate market timing patterns in MV_{CV} and MV_{CY} that are different from the market timing of MV. In the next section, we show that the market timing of MV generates a more desirable unconditional return distribution.

MV_{FS} is a conditionally mean–variance efficient strategy with portfolio weights $\theta_t^{MV_{FS}} = \frac{\zeta \mu_t^{MV}}{(\mu_t^{MV})^2 + (\sigma_t^{MV})^2} \theta_t^{MV}$, where ζ is a time-invariant constant parameter.[3] Ferson and Siegel (2001) and Penaranda (2016) introduce and discuss theoretical properties of MV_{FS}. Moreover, Chernov, Dahlquist, and Lochstoer (2021) test the empirical properties of MV_{FS} as a pricing factor in FX markets. MV_{FS} features a different market timing than MV. That is, MV_{FS} reduces its holdings uniformly across currencies when the conditional Sharpe ratio increases. A comparison of MV_{FS} and MV in our empirical section reveals that MV's market timing approach leads to a more attractive OOS performance.

2.2.c. Other Common Strategies

For completeness, we also consider other currency strategies that are commonly used in the literature. We briefly characterize some of the key strategies and relegate details to Appendix. These strategies include the Dollar strategy DOL (Lustig, Roussanov, and Verdelhan, 2011), the Dollar carry strategy DDOL (Lustig, Roussanov, and Verdelhan, 2014), the momentum strategy MOM (Burnside, Eichenbaum, and Rebelo, 2011; Menkhoff *et al.*, 2012b), the value strategy VAL (Bilson, 1984; Menkhoff *et al.*, 2017), and the equally weighted high-minus-low forward discount carry strategy HML (Lustig and Verdelhan, 2007).

SW (or spread weighted carry strategy) is defined by portfolio weights that are set equal to the forward discounts, $\theta_t^{SW} = f_t$. The SW strategy is dynamic and features a market timing based on the absolute size of the forward discounts. SW enhances the strategy HML by taking into account the size and time variation in the forward discounts.

HML_{VM} (or the volatility managed portfolio of HML) and SW_{VM} (or the volatility managed portfolio of SW) are strategies with portfolio weights $\theta_t^{HML_{VM}} = \frac{1}{(\sigma_t^{HML})^2} \theta_t^{HML}$ and $\theta_t^{SW_{VM}} = \frac{1}{(\sigma_t^{SW})^2} \theta_t^{SW}$. The conditional volatilities σ_t^{HML} and σ_t^{SW} of the HML and SW are computed using daily returns of these factors over the past month (Moreira and Muir, 2017).

2.3 Data

We collect daily spot and 1-month forward exchange rates from October 31, 1983 to January 30, 2016 from Reuters via Datastream. We use quotes of the last day of the month to compute monthly returns.

In the main analysis, we focus on the data of fifteen developed countries reported and analyzed in Lustig, Roussanov, and Verdelhan (2011). Trading frictions are typically lower for currencies of developed countries (e.g., they have a large active trading volume, there are less capital controls, liquidity is higher, and transaction costs are lower) than for currencies of emerging countries. Thus, our strategy may be easier and cheaper to implement using only developed currencies. We check the robustness of our results using the dataset of

3 In theory, if MV is conditionally mean–variance efficient and for $\vartheta = 1$, then the conditional Sharpe ratio of MV is $SR_t^{MV} = \sigma_t^{MV} = \sqrt{\mu_t^{MV}}$, and thus, $\theta_t^{MV_{FS}} = \frac{\zeta}{1 + (SR_t^{MV})^2} \theta_t^{MV}$.

twenty-nine developed and emerging currencies. The currencies included are the standard ones considered in the literature, see for example, Lustig, Roussanov, and Verdelhan (2011). All our results remain qualitatively the same. The results for this twenty-nine-currency set are available upon request.

3. Economic Value of Market Timing

Section 3.1 documents that the OOS performance of MV dominates other mean–variance optimized strategies. Section 3.2 examines how the market timing of MV, which is characterized by its time-varying notional value, predicts future returns and risks. Section 3.3 validates the key implication that market timing is more beneficial when market conditions and return moments are more volatile. Finally, Section 3.4 compares market timing based on expected returns versus volatility.

3.1 Performance of Mean–Variance Efficient Strategies

Our empirical analysis starts with a performance evaluation of MV in comparison with four other conditionally mean–variance efficient strategies, MV_{CV}, MV_{CY}, MV_{CN}, and MV_{FS} (Section 2.2). Table I presents the OOS performance of these strategies. All strategies use information available at the end of month t to construct the portfolios that are held from end of t to end of $t+1$. Thus, there is no look-ahead bias. All reported statistics are annualized and include the average excess returns (mean), volatility (Vol), unconditional Sharpe ratio (SR), skewness (Skew), and kurtosis (Kurt). We use the test proposed by Ledoit and Wolf (2008), which is robust to heteroskedasticity and

Table I. Performance of mean–variance optimized portfolios

Statistics of monthly OOS excess returns of mean–variance optimized portfolios. MV maximizes the expected return minus a constant times the variance of the portfolio. MV_{CV}, MV_{CY}, and MV_{CN} impose a constant volatility, constant yield, and constant notional value in the construction of MV. MV_{FS} is the mean–variance portfolio suggested by Ferson and Siegel (2001) and reduces the risk exposure of MV when the conditional Sharpe ratio increases. ΔSR is the difference between the Sharpe ratio of MV and that of other strategies. The reported p-value is the test for $\Delta SR = 0$. Standard errors are robust to heteroskedasticity and autocorrelation (Ledoit and Wolf, 2008). All statistics are annualized.

	MV	MV_{CV}	MV_{CY}	MV_{CN}	MV_{FS}
Mean (in %)	20.30	12.04	8.59	8.01	9.93
Vol (in %)	18.06	13.93	21.04	12.79	9.29
SR	1.12	0.86	0.41	0.63	1.07
ΔSR		0.26***	0.72***	0.50***	0.06**
(p-value)		(0.0078)	(0.0026)	(0.0008)	(0.0234)
Skew	0.67	−0.48	−0.07	−0.66	0.09
Kurt	6.63	3.96	8.52	4.93	5.80
MDD (in %)	−36.07	−58.73	−132.13	−74.22	−22.60
|MDD|/Mean	1.78	4.88	15.39	9.26	2.28
Utility gamma = 1 (in %)		5.38	7.42	8.74	8.38
Utility gamma = 10 (in %)		2.81	3.88	3.08	5.50

autocorrelation, to test the hypothesis whether the differential Sharpe ratio (ΔSR) of MV versus other strategies is equal to zero, that is, the Sharpe ratios are equal. We indicate a rejection of the test at the 10%, 5%, and 1% significance level by one, two, and three stars.[4] We also report the MDD (in percentage points), which measures the maximum loss from peak to trough a strategy has experienced during the entire sample period, and the MDD per 1% expected excess return (|MDD|/Mean). |MDD|/Mean also measures the expected time in years to recover from the MDD. Since all strategies can be levered up or down, |MDD|/Mean is a more useful measure than MDD when we compare the downside or crash risks between strategies. We estimate the relative utility gain of MV with respect to other strategies (Utility) for investors with constant relative risk aversion of $\gamma = 1$ and $\gamma = 10$. That is, we follow Fleming, Kirby, and Ostdiek (2001) and define Utility as the annual percentage an investor is willing to pay to switch from a strategy to MV.

The main findings in Table I are as follows. The MV strategy outperforms other conditionally mean–variance efficient strategies across various performance measures. MV offers the highest OOS unconditional Sharpe ratio of 1.12 per annum, which is statistically significantly higher than that of other strategies. MV is the only strategy that features a return distribution with both a positive and a large (OOS and monthly) skewness of 0.67. That is, MV frequently has high returns (i.e., a large upside potential), whereas it's low returns are relatively close to the mean (i.e., a limited downside risk). A large literature argues that the high returns of currency strategies are a compensation for crash risk (Brunnermeier, Nagel, and Pedersen, 2008; Dobrynskaya, 2014; Galsband and Nitschka, 2014; Lettau, Maggiori, and Weber, 2014). The positive skewness and high average return of the MV indicate that it is not exposed to or explained by large crash risks. The MDD strengthens the finding that MV does not have much crash risk exposure. Among all reported strategies, MV has the least maximum loss of −36.07%. In comparison, the yield target mean–variance optimized strategy MV_{CY} has a MDD of −132.13%, that is, losing all of its investment at some point during our sample. The MV also has the lowest MDD per 1% expected excess return (|MDD|/Mean) at 1.78, indicating its shortest expected time to recover from the MDD. MV also offers a large positive utility gain. That is, for a reasonable range of relative risk aversions ($\gamma = 1$ and $\gamma = 10$), investors are willing to pay more than 5.38% per year to switch from any of the alternative conditionally mean–variance efficient strategies to MV. Finally, note that the next best strategy is MV_{FS}. In theory, MV_{FS} is constructed to optimize the unconditional Sharpe ratio. However, empirically it underperforms MV OOS.

In summary, as all strategies in Table I are conditionally mean–variance efficient, MV's unconditional outperformance is due to its market timing. The outperformance is statistically and economically significant. Our findings demonstrate the importance of the market timing of MV in real market data. It is costly to limit the leverage or risk of a portfolio (i.e., impose a constant notional value or volatility) or implement other suboptimal market timing policies. This finding is important as such limits are often imposed in practice.

4 We choose a block size of 10 for block-bootstrapping to estimate *p*-values. This is a conservative value and our results are stronger if we use a smaller value closer to what is used by Ledoit and Wolf (2008) in their examples.

3.2 Market Timing and FX Market Predictability

The outperformance of MV due to its market timing indicates that the strategy has predictive power in FX markets. It is important to understand whether properties of MV's OOS return distribution, such as the positive skewness and low relative MDD (|MDD|/Mean), can be explained by MV's ability to predict adverse events in FX markets. We examine this hypothesis by testing whether MV's market timing can predict future FX market returns or moments at either the portfolio level (Table II) or individual currency level (Table III).

We first test whether the notional value of MV $\sum_i |\theta_{i,t}^{\mathrm{MV}}|$ can predict returns and risks of the conditionally mean–variance efficient portfolio $\mathrm{MV}_{\mathrm{CN}}$ without market timing. Specifically, we predict the future (next-month) realized Sharpe ratio (i.e., realized excess return divided by the volatility), realized excess return, and volatility of strategy $\mathrm{MV}_{\mathrm{CN}}$. The volatility is estimated using daily returns within the month. Since the weights of MV are equal to the notional value of MV multiplied by the weights of $\mathrm{MV}_{\mathrm{CN}}$, successful market timing implies that the notional value predicts expected returns or risks of $\mathrm{MV}_{\mathrm{CN}}$. Table II reports the estimation results and R^2 of the predictive regressions. We control for the conditional expected excess return and volatility of $\mathrm{MV}_{\mathrm{CN}}$ in our regressions. This provides insights to what extent the first two moments of $\mathrm{MV}_{\mathrm{CN}}$ capture the information content of MV's notional value.

We find that MV's notional value $\sum_i |\theta_{i,t}^{\mathrm{MV}}|$ by itself is able to predict the future realized excess return, volatility, and Sharpe ratio of strategy $\mathrm{MV}_{\mathrm{CN}}$. The regression slope coefficients are significant at the 1% level and the adjusted R^2 is 2.67%, 19.57%, and 5.72%, respectively. These results confirm MV's ability to time the market. The relatively high R^2 in the prediction of the volatility indicates that there is particularly much value to time the market based on changes in risks. The signs of slope coefficients are as expected. An increase in MV's current exposure to FX markets is associated with a higher return, lower volatility, and higher Sharpe ratio of the strategy $\mathrm{MV}_{\mathrm{CN}}$ in the subsequent month. That is, MV trades more aggressively when future returns are higher, risks are lower, and the expected return-risk trade-off is more attractive.

Panel A shows that the conditional mean $\mu_t^{\mathrm{MV_{CN}}}$ of strategy $\mathrm{MV}_{\mathrm{CN}}$ by itself does not have much predictive power. When we include both $\sum_i |\theta_{i,t}^{\mathrm{MV}}|$ and $\mu_t^{\mathrm{MV_{CN}}}$ as predictors in the predictive regression (Panel A), MV's notional value remains a statistically significant predictor of $\mathrm{MV}_{\mathrm{CN}}$'s future realized return, volatility, and Sharpe ratio. Panel B shows that the conditional volatility $\sigma_t^{\mathrm{MV_{CN}}}$ of strategy $\mathrm{MV}_{\mathrm{CN}}$ by itself is able to predict the future performance of $\mathrm{MV}_{\mathrm{CN}}$ at the 1% level of significance. When we include both $\sum_i |\theta_{i,t}^{\mathrm{MV}}|$ and $\sigma_t^{\mathrm{MV_{CN}}}$ as predictors in the predictive regression (Panel B), MV's notional value also remains a statistically significant predictor of $\mathrm{MV}_{\mathrm{CN}}$'s volatility and Sharpe ratio. These results show that $\mathrm{MV}_{\mathrm{CN}}$'s conditional mean and conditional volatility do not fully subsume the predictive power of MV's notional value.

In summary, Table II provides evidence in favor of MV's ability to time the market. In particular, we find that the notional value of MV increases when $\mathrm{MV}_{\mathrm{CN}}$'s realized returns are high, risk is low, and its expected return-risk trade-off is more attractive in the subsequent month.

Next, we consider changes in MV's portfolio weights as a predictor of returns and risks of individual currencies (Table III). These more granular regressions shed light on both the market timing ability of MV and the conditional efficiency of mean–variance optimized portfolio in general. That is, changes in the portfolio weights of MV should reflect a time-

Table II. Portfolio level predictive regressions

One-month ahead predictive regressions: $Y_{t,t+1} = c_{\text{const}} + c_{\text{trend}}\,t + \sum_i c_i X_{i,t} + \epsilon_t$. $Y_{t,t+1}$: realized Sharpe ratio (SR(MV_{CN})), realized return (MV_{CN}, in %), or realized volatility (Vol(MV_{CN}), in %) of strategy MV_{CN}. $X_{i,t}$: MV dollar risk exposure ($|\theta_{i,t}^{MV}|$), conditional mean ($\mu_t^{MV_{CN}}$, in %), conditional risk ($\sigma_t^{MV_{CN}}$, in %). c_{const} is a constant term, $c_{\text{trend}}\,t$ controls for any time-series trend in $Y_{t,t+1}$, c_i are the slope coefficients associated with the predictors $X_{i,t}$; their corresponding *t*-statistics are reported in parentheses. Significance at the 1%, 5%, or 10% level is indicated by ***, **, or *, where standard errors are calculated using Newey and West (1987) to account for heteroskedasticity and auto-correlation.

Panel A		MV_{CN}			Vol(MV_{CN})			SR(MV_{CN})	
$\sum_i \lvert\theta_{i,t}^{MV}\rvert$	0.126***		0.122**	−0.430***		−0.450***	0.060***		0.059***
(*t*-stat)	(2.649)		(2.281)	(−6.651)		(−6.839)	(3.512)		(3.342)
$\mu_t^{MV_{CN}}$		0.057	0.036		0.082	0.158**		0.016	0.007
(*t*-stat)		(0.750)	(0.447)		(1.090)	(2.388)		(0.922)	(0.346)
Adj. R^2 (%)	2.67	1.01	2.53	19.57	6.66	20.70	5.72	1.43	5.51
Panel B		MV_{CN}			Vol(MV_{CN})			SR(MV_{CN})	
$\sum_i \lvert\theta_{i,t}^{MV}\rvert$	0.126***		0.028	−0.430***		−0.126**	0.060***		0.041**
(*t*-stat)	(2.649)		(0.581)	(−6.651)		(−2.330)	(3.512)		(2.000)
$\sigma_t^{MV_{CN}}$		−0.203***	−0.180*		0.658***	0.557***		−0.067***	−0.034
(*t*-stat)		(−2.720)	(−1.939)		(9.058)	(5.786)		(−3.661)	(−1.498)
Adj. R^2 (%)	2.67	4.12	3.93	19.57	27.38	27.84	5.72	5.10	6.05

Table III. Individual currencies predictive panel regression

One-month ahead predictive regressions: $Y_{i,t+1} = c_{\text{const}} + c_{\text{trend}}\,t + c_\theta X_{i,t} + c_\mu \mu_{i,t} + c_\sigma \sigma_{i,t} + c_{\text{lag}} Y_{i,t} + \epsilon_{i,t}$. c_{const} is a constant term, $c_{\text{trend}}\,t$ controls for any time-series trend, and c_{lag} controls for the lag of the LHS variable. $Y_{i,t+1} = \Delta rx_{i,t+1} = (rx_{i,t+1} - rx_{i,t})$ is the change in realized monthly return of currency i against the USD. $Y_{i,t+1} = \Delta\text{Vol}(rx_{i,t+1}) = (\text{Volatility of } rx_{i,t+1} - \text{Volatility of } rx_{i,t})$ is the change in realized monthly volatility (i.e., volatility estimated from daily data within the month) of the returns of currency i against the USD. $Y_{i,t+1} = \Delta\text{Skew}(rx_{i,t+1}) = (\text{Skewness of } rx_{i,t+1} - \text{Skewness of } rx_{i,t})$ is the change in realized monthly skewness of the returns of currency i against the USD. $\Delta\theta^{MV}_{i,t} = (\theta^{MV}_{i,t} - \theta^{MV}_{i,t-1})$ is the change in portfolio weight of MV in currency i. $\Delta|\theta^{MV}_{i,t}| = (|\theta^{MV}_{i,t}| - |\theta^{MV}_{i,t-1}|)$ is the change in the exposure of MV to currency i. The panel regressions include currency-fixed effects. Significance at the 1%, 5%, or 10% level is indicated by ***, **, or *. Standard errors are calculated to account for cross-currency correlations, heteroskedasticity, and within currency autocorrelations (Newey and West, 1987).

$Y_{i,t+1}$	$\Delta rx_{i,t+1}$	$\Delta\text{Vol}(rx_{i,t+1})$	$\Delta\text{Skew}(rx_{i,t+1})$
$\Delta\theta^{MV}_{i,t}$	0.0006***		0.0035***
(t-stat)	(3.08)		(2.67)
$\Delta\lvert\theta^{MV}_{i,t}\rvert$		−0.0012***	
(t-stat)		(−6.93)	
$\mu_{i,t}$	0.0680**	−0.0353***	0.0111
(t-stat)	(2.12)	(−3.77)	(0.08)
$\sigma_{i,t}$	−0.5609	2.2488***	0.9527
(t-stat)	(−0.52)	(5.56)	(0.28)
R^2 (in %)	26.46	28.33	27.00

variation in the expected return-risk trade-off of the portfolio as well as of the individual currency relative to the other currencies. Table III reports the estimated regression coefficients and R^2 of predictive panel regressions. We predict changes in future returns, volatility, and skewness of all individual currencies. All predictive panel regressions of Table III include the currency-specific forward discount $\mu_{i,t} = f_{i,t}$, the conditional volatility $\sigma_{i,t} = \sqrt{\Omega_{i,i,t}}$, the time lag of the predicted quantity, and a time trend as controls.

In the panel data of individual currencies, the change $\Delta\theta^{\text{MV}}_{i,t}$ in MV's portfolio weight in currency i over past months is able to predict changes in the return and skewness of currency i's return against the USD in the subsequent month. The regression slope coefficients are significant at the 1% level.[5] Similarly, a change in MV's exposure to currency i in the past, $\Delta|\theta^{\text{MV}}_{i,t}|$, is able to predict changes in the volatility of individual currency returns in the subsequent month. The regression slope coefficient is significant at the 1% level. This confirms that MV invests more in a currency when the future excess return is higher and the crash risk lower. Moreover, MV reduces its risk exposure to currencies which are more volatile in the subsequent month. That is, MV can time currency markets and increases loadings on desirable (high return, low volatility, and positive skewness) currencies.

5 Because both return and skewness are signed quantities (i.e., both their sign and magnitude are meaningful), we employ the original (but not absolute) change $\Delta\theta^{\text{MV}}_{i,t}$ in portfolio weight as a signed predictor. For unsigned volatility prediction, we employ the unsigned exposure $\Delta|\theta^{\text{MV}}_{i,t}|$ as an unsigned predictor.

Compared with Table II, the current predictive regressions tend to offer higher R^2. These may be due to several reasons. First, we control for the conditional mean $\mu_{i,t}$ and conditional volatility $\sigma_{i,t}$ of individual currency returns. Second, the panel regression does not only focus on market timing but also on the relative importance of each currency in the composition of the MV portfolio.

In summary, Table III strengthens the evidence for MV's market timing and efficiency. We show that changes in MV's portfolio weights and exposures predict future changes in the performance of individual currencies, even after controlling for the currency-specific forward discount and volatility.

For completeness, the Supplementary (Online) Appendix presents the slope coefficient estimates and R^2 of predictive regressions for all individual currencies, concerning the results of the regressions (i) of the change $\Delta rx_{i,t+1}$ in currency i's realized return on the change $\Delta\theta_{i,t}^{\mathrm{MV}}$ in MV's portfolio weight in currency i, (ii) of the change $\Delta \mathrm{Vol}(rx_{i,t+1})$ in currency i's realized return volatility on the change $\Delta|\theta_{i,t}^{\mathrm{MV}}|$ in MV's exposure to currency i, and (iii) of the change $\Delta \mathrm{Skew}(rx_{i,t+1})$ in currency i's realized return skewness on $\Delta\theta_{i,t}^{\mathrm{MV}}$.

3.3 Market Timing and Changes in Economic Conditions

Having demonstrated the effect of the market timing on the performance and predictive power of MV, we now examine an important testable economic implication of MV's market timing. As suggested by the intuitive analysis in Section 2.1, the outperformance of MV over other conditionally mean–variance efficient strategies is facilitated by the time variation in economic conditions. That is, market timing delivers value in a market environment in which the timing is relevant. To test this implication, we consider various conditioning variables that proxy for the state of the economy and financial markets. We examine whether periods of large versus small absolute changes in these variables coincide with periods of significant versus less significant outperformance of strategy MV over other conditionally mean–variance efficient strategies.

Table IV reports the outperformance of MV over the strategies $\mathrm{MV_{CV}}$, $\mathrm{MV_{CY}}$, $\mathrm{MV_{CN}}$, and $\mathrm{MV_{FS}}$ while conditioning on small versus large changes in five alternative state variables. The five conditioning variables are the conditional expected excess return of MV (μ_t^{MV}; Panel A), conditional volatility of MV (σ_t^{MV}; Panel B), 3-month T-bill rate (Panel C), term spread (the difference between 10-year and 3-month Treasury yields, Panel D), and default spread (the difference between Moody's Aaa corporate bond yield and 3-month Treasury yield, Panel E). These variables either characterize MV's expected return-risk trade-off, or are known indicators of the economy and financial markets. The relative performance is measured by the difference in Sharpe ratios ΔSR of MV versus an alternative strategy. We divide our data into two (quantile) subsamples, which contain months of, respectively, small (below median) and large (above median) absolute changes in the five alternative conditioning variables. For each conditioning variable, we report the difference in Sharpe ratios ΔSR for the two subsamples.

Across all conditioning variables, the Sharpe ratio differential between the MV and all alternative strategies is consistently higher in subsamples of large absolute changes in the conditioning variables compared with subsamples of small absolute changes. In fact, when changes in the conditioning variables are large, then the Sharpe ratio of MV is always statistically significantly higher than the Sharpe ratio of all other strategies. In contrast, when changes in the conditioning variables are small, then the difference in Sharpe ratios is in

Table IV. Economic conditions and MV

Differences between the Sharpe ratio of MV and those of other conditionally efficient mean–variance portfolios conditional on small (below median) versus large (above median) changes in state variables. Significance at the 1%, 5%, or 10% level is indicated by ***, **, or *. Each panel reports the results for one conditioning variable.

	$\Delta SR(MV_{CV})$	$\Delta SR(MV_{CY})$	$\Delta SR(MV_{CN})$	$\Delta SR(MV_{FS})$
	Panel A. Conditional mean μ_t^{MV}			
Small changes	0.214	0.429*	0.407*	0.034
Large changes	0.231**	0.755**	0.472**	0.068*
	Panel B. Conditional volatility σ_t^{MV}			
Small changes	0.222	0.587*	0.437*	0.032
Large changes	0.256**	0.721**	0.519**	0.069*
	Panel C. 3-month T-bill			
Small changes	0.168	0.504*	0.273	0.036
Large changes	0.340***	0.935*	0.698***Z***	0.070**
	Panel D. Term spread			
Small changes	0.128	0.462	0.242	0.028
Large changes	0.375***	0.941**	0.711***	0.077**
	Panel E. Default spread			
Small changes	0.159	0.547**	0.257	0.032
Large changes	0.349**	0.865**	0.676***	0.075**

most cases statistically insignificant. As absolute changes in these conditioning variables represent the time variation in market conditions, this finding supports the implication that MV's market timing is relatively more beneficial when return moments are more volatile.

Furthermore, in the cross section of all strategies and for both subsamples of small and large changes in the conditioning variables, MV always offers the highest Sharpe ratio among all conditionally efficient mean–variance strategies. This is reflected by $\Delta SR > 0$ everywhere in Table IV. This emphasizes MV's steady outperformance over other strategies across different economic conditions. The Sharpe ratio differential (i.e., outperformance of MV) ranges from 0.941 for MV versus MV_{CY} in the subsample of large changes in the Term Spread, to 0.028 for MV versus MV_{FS} in the subsample of small changes in the Term Spread. Among the alternative strategies, MV_{FS} is closest to MV, and MV_{CY} consistently performs the worst. This finding fits the fact that MV_{FS} features market timing to improve the unconditional Sharpe ratio, while MV_{CY} basically follows the opposite timing pattern of MV.

In summary, Table IV provides evidence to support the relative advantage of MV's market timing over other conditionally mean–variance efficient strategies. Market timing is most beneficial when market conditions are rapidly changing from month to month. In turn, it is especially costly to impose leverage or risk (i.e., conditional volatility) limits during these times.

3.4 Market Timing Based on Expected Returns versus Volatility

To further evaluate the economic value of timing with respect to the conditional expected excess return and the conditional volatility, we construct alternative strategies that focus on the expected excess return, the return volatility, and combinations of the two. We compare

Table V. Market timing: Expected returns versus volatility

Statistics of monthly OOS excess returns. MV maximizes the expected return minus a constant times the variance of the portfolio. HML is an equally weighted carry trade strategy. SW is the spread-weighted carry trade strategy. HML_{VM} and SW_{VM} are volatility managed portfolio of HML and SW. ΔSR is the difference between the Sharpe ratio of MV and that of other strategies. The reported *p*-value is the test for $\Delta SR = 0$. Standard errors are robust to heteroskedasticity and autocorrelation (Ledoit and Wolf, 2008). All statistics are annualized.

	MV	HML	SW	HML_{VM}	SW_{VM}
Mean	20.30	5.23	5.67	7.93	11.11
Vol	18.06	9.66	9.00	10.57	13.20
SR	1.12	0.54	0.63	0.75	0.84
ΔSR		0.59**	0.49**	0.37**	0.28
(*p*-value)		(0.012)	(0.030)	(0.043)	(0.202)
Skew	0.67	−0.89	−0.61	2.56	−0.02
Kurt	6.63	5.59	11.92	26.21	13.66
MDD	−36.07	−42.78	−43.78	−22.09	−39.90
|MDD|/Mean	1.78	8.18	7.72	2.79	3.59

MV to four alternative strategies: HML, SW, HML_{VM}, and SW_{VM}. Unlike the strategies in Table I, these four strategies are (in general) not conditionally mean–variance efficient. HML is the standard high-minus-low interest rate sorted currency carry strategy (Lustig and Verdelhan, 2007). SW is a spread weighted currency carry strategy with portfolio weights $\theta_t^{SW} = f_t$. The subscript VM (volatility managed) indicates a scaling of the portfolio weights based on the current return volatility of the strategy (see Section 2.2 for details). Table V presents the OOS performance of these strategies. The performance metrics are the same as in Table I.

First, we make comparisons between the four alternative strategies to get insights into the value of market timing based on expected returns versus volatility. In contrast to HML, SW is not equally weighted and its notional value increases (decreases) when the average magnitude of forward discounts increases (decreases). Thus, similar to MV, SW features some fine-tuning in its portfolio weights with respect to $\mu_t = f_t$ and it uses market timing based on the magnitude of conditional expected excess returns. SW has a higher Sharpe ratio, less negative skewness, and smaller MDD per 1% expected return compared with HML. This suggests that there is value in fine-tuning the portfolio weights and timing the market based on the magnitude of expected excess returns. However, the improvement in OOS performance is relatively moderate. That is, SW still delivers less attractive OOS returns than MV.

The comparison between HML and HML_{VM} as well as SW and SW_{VM} provides insights about the benefit of market timing with respect to volatility while keeping other characteristics of the strategy constant. The volatility adjusted portfolios are in the spirit of Fleming, Kirby, and Ostdiek (2001) and Moreira and Muir (2017). The volatility managed strategies have higher Sharpe ratios (an increase of about 0.2). More strikingly there is a substantial improvement with respect to the downside risk. The skewness turns from negative to positive (in the case of HML) or zero (in the case of SW). Moreover, the MDD per 1% expected return is less than half after market timing based on volatility. Therefore, market timing based on volatility is important, in particular in terms of reducing crash risks.

Finally, we emphasize that MV outperforms all four strategies across various performance measures. MV offers an OOS unconditional Sharpe ratio of 1.12, which is statistically significantly higher than that of HML (0.54), SW (0.63), and HML_{VM} (0.75). It also dominates the Sharpe ratio of SW_{VM} (0.84), though the difference is not statistically significant. Furthermore, MV features an unconditional return distribution with a positive skewness and it has the most desirable MDD per 1% expected excess return, $|MDD|/Mean = 1.78$. In comparison, |MDD|/Mean for HML_{VM} is 2.79, for SW_{VM} is 3.59, and for HML and SW they are 8.18 and 7.72. Accordingly, MV features substantially less crash risks.

In summary, we find that market timing based on volatility delivers relatively more value than market timing based on expected returns. However, the two are not mutually exclusive and both dimensions are important. A combination of both dimensions significantly improves the OOS performance.

4. MV versus Common Benchmarks in the Literature

In Section 4.1, we first compare the performance of MV and various benchmark portfolios. In Section 4.2, we show that MV is not spanned by common FX risk factors.

4.1 Performance of MV versus Benchmarks

We now compare MV to popular FX strategies in the literature, namely, DOL, DDOL, HML, MOM, and VAL. In our analysis, we consider these five strategies individually as benchmarks and further construct various portfolios of these strategies. This empirical exercise aims to position MV against a broad set of common benchmarks in the FX literature. Panel A in Table VI reports the OOS performance of MV, while Panel B provides statistics of the five benchmark strategies. Panel C considers equally weighted portfolios of the five benchmark strategies, and Panels D and E construct global minimum variance and mean–variance optimized portfolios.[6] For the global minimum variance and the mean–variance optimized portfolios, we require the covariance matrix and the expected excess returns of the five strategies. For simplicity, we use sample estimates using our entire sample. In that sense, these are static and unconditionally optimized portfolios. Moreover, there is a look-ahead bias and the performance is not OOS. Accordingly, an outperformance of MV (which does not suffer from a look ahead bias and its performance is OOS) is even more impressive. The portfolios are constructed using diverse subsets of the five strategies.[7]

The main findings of Table VI are as follows. MV offers returns with a higher OOS unconditional Sharpe ratio than all benchmarks. In many cases, the difference in Sharpe ratios is statistically significant. In addition, MV has a positive skewness, while the skewness is negative for all other strategies, except for MOM. In the same spirit, MV has the most desirable MDD per 1% expected excess return |MDD|/Mean. Thus, MV is more attractive in terms of the Sharpe ratio as well as in terms of crash risks. This finding emphasizes the importance of MV in comparison to popular benchmarks in the literature.

6 Global-minimum-variance portfolio is the left-most portfolio on the efficient frontier in the variance–mean coordinate plane.

7 The subsets are {HML, DDOL}, {HML, DDOL, MOM}, {HML, DDOL, VAL}, {HML, DDOL, DOL}, {HML, DDOL, MOM, VAL}, and {HML, DDOL, MOM, VAL, DOL}.

Table VI. Performance of MV and popular currency portfolios in the literature

Panel A shows the performance of MV. Panel B shows the performance of popular long-short currency portfolios. Panels C, D, and E show the performance of equally weighted, global minimum variance, and mean–variance optimized portfolios of subsets of the strategies in Panel B. All statistics are annualized.

	Mean (%)	Vol (%)	SR	ΔSR	Skew	Kurt	$\frac{\lvert MDD \rvert}{Mean}$
Panel A. MV	20.30	18.06	1.12		0.67	6.63	1.78
Panel B. Long-short							
DOL	1.55	8.63	0.18	0.94***	−0.17	3.60	27.21
DDOL	5.15	8.52	0.61	0.52**	−0.16	3.71	4.68
HML	5.23	9.66	0.54	0.58**	−0.89	5.59	8.18
MOM	2.34	13.79	0.17	0.95***	0.40	6.91	12.98
VAL	4.55	8.68	0.52	0.60**	−0.01	4.18	5.01
Panel C. 1/*N*							
HML, DDOL	5.19	6.85	0.76	0.37*	−0.61	4.44	4.80
HML, DDOL, MOM	4.24	6.63	0.64	0.49**	−0.58	4.03	3.67
HML, DDOL, VAL	4.98	5.64	0.88	0.24	−0.49	5.00	2.91
HML, DDOL, DOL	3.98	6.23	0.64	0.49**	−0.45	5.15	5.56
HML, DDOL, MOM, VAL	4.32	5.36	0.81	0.32	−0.41	4.45	2.50
HML, DDOL, MOM, VAL, DOL	3.76	4.70	0.80	0.32*	−0.56	4.19	2.57
Panel D. Global minimum variance							
HML, DDOL	5.19	6.80	0.76	0.36*	−0.55	4.33	4.30
HML, DDOL, MOM	4.67	6.23	0.75	0.38*	−0.70	3.87	3.62
HML, DDOL, VAL	4.93	5.45	0.91	0.22	−0.28	4.72	2.17
HML, DDOL, DOL	4.03	6.23	0.65	0.48**	−0.45	5.14	5.41
HML, DDOL, MOM, VAL	4.54	5.01	0.91	0.22	−0.16	5.08	1.82
HML, DDOL, MOM, VAL, DOL	3.43	4.26	0.81	0.32	−0.19	4.57	3.21
Panel E. mean–variance optimized							
HML, DDOL	5.29	6.94	0.76	0.36*	−0.56	4.34	4.34
HML, DDOL, MOM	5.43	7.03	0.77	0.35*	−0.64	4.02	3.90
HML, DDOL, VAL	7.48	8.25	0.91	0.22	−0.31	4.67	2.14
HML, DDOL, DOL	5.45	7.04	0.77	0.35*	−0.54	4.15	4.06
HML, DDOL, MOM, VAL	7.81	8.43	0.93	0.20	−0.29	4.42	1.97
HML, DDOL, MOM, VAL, DOL	7.85	8.45	0.93	0.20	−0.29	4.40	1.94

In summary, Table VI provides additional evidence for the importance of MV and its outperformance over a broad mix of well-known benchmarks.

4.2 Pricing

This section empirically studies MV's performance from an asset pricing perspective. Specifically, we address the question whether MV's excess return and Sharpe ratio can be explained by, and hence is a compensation for MV's loadings on well-known risk factors in the FX literature.

Table VII. Abnormal returns of the MV

Abnormal returns of MV according to linear factor models. Time-series regression: $MV_t = \alpha_{MV} + \sum_i \beta_{MV,i} F_{i,t} + \varepsilon_t$. Cross-sectional relationship: $E[MV_t] = \alpha^*_{MV} + \sum_i \beta_{MV,i}\gamma_i$. $E[x_t]$ is the time-series average of x_t, α_{MV} the abnormal return in the time-series equation, α^*_{MV} the abnormal return in the cross-sectional equation, $\beta_{MV,i}$ is the factor loading of the MV on factor $F_{i,t}$, and γ_i is the risk premium (estimated in the cross-section) of factor $F_{i,t}$. Set A contains factors DOL, HML, DDOL, MOM, VAL, SW, HML_{VM}, and SW_{VM}. Set B contains all factors in Set A plus VOL, SKEW, and ILL. Set C contains all factors in Set B plus MKT and INT. The reported α_{MV} and α^*_{MV} are annualized. R^2 measures the fit of the time-series relationship. Newey and West (1987) robust t-statistics are reported in parentheses next to the coefficient estimates. Significance at the 1%, 5%, or 10% level is indicated by ***, **, or *.

Factors	α^*_{MV} (in %)	(t-stat α^*)	α_{MV} (in %)	(t-stat α)	R^2 (%)
DOL, HML	16.361***	(4.637)	15.880***	(4.743)	20
DOL, DDOL	18.453***	(4.718)	18.117***	(4.768)	5
DOL, SW	16.236***	(4.724)	14.690***	(4.859)	24
DOL, HML_{VM}	15.840***	(4.505)	12.510***	(4.073)	30
DOL, SW_{VM}	15.116***	(4.289)	12.268***	(4.005)	27
DOL, VOL	19.068***	(5.040)			4
DOL, SKEW	17.412***	(4.819)			3
DOL, ILL	19.120***	(4.941)			3
MKT, INT	19.575***	(5.033)	19.945***	(5.049)	0
Set A (eight factors)	12.588***	(3.909)	9.163***	(3.579)	43
Set B (eleven factors)	12.068***	(3.879)			44
Set C (thirteen factors)	9.546***	(3.308)			45

Table VII reports MV's abnormal return estimates in twelve different linear factor models. We consider both traded factors (DOL, HML, DDOL, MOM, VAL, SW, HML_{VM}, SW_{VM}, MKT, and INT) and non-traded factors (VOL, ILL, and SKEW). For each model, the pricing of MV's returns is evaluated based on its model-implied abnormal returns under two alternative approaches. The factor risk premia are estimated as the time-series average of the factor returns in the first approach (provided that the factors are traded portfolios) and as the implied premia in a cross-sectional pricing regression in the second approach.[8] The abnormal returns are referred to as α_{MV} (in the time-series estimation approach) and α^*_{MV} (in the cross-sectional estimation approach) in Table VII. For completeness, the corresponding estimates of MV's loadings (i.e., betas) on various pricing factors are reported in the Supplementary (Online) Appendix.

Table VII shows that across all linear factor models, strategy MV earns large positive and statistically significant abnormal returns, ranging from 9% to 20% for both time-series and cross-sectional estimations. The positive signs of the abnormal returns indicate that MV's excess returns are not explained by popular FX market risks. In other

8 Risk premia of factors that are not traded need to be estimated in the cross-sectional pricing regression using the generalized method of moments (GMM). To estimate risk premia in the cross-sectional pricing regression, we use the following test assets: five interest sorted, five momentum sorted, five value-sorted portfolios, DDOL, SW, SW_{CN}, HML_{VM}, SW_{VM}, and various mean–variance optimized portfolios.

words, MV is not spanned by well-known currency strategies or factors. In summary, Table VII establishes that MV's return is an anomaly with respect to popular FX risk factors and indicates the existence of new risk sources to rationalize MV's return.

5. Robustness

This section carries out various robustness checks to strengthen our empirical analysis. We briefly describe these robustness checks here and relegate details to the Supplementary (Online) Appendix. MV's construction is based on exchange rate forward discounts as proxies for the conditional expected excess returns and a robust PCA-based estimation of the covariance matrix (Section 2.2). We verify the robustness of these estimations.

For alternative estimates of the conditional expected excess returns, we either use a simple historical average of returns, or forward discounts together with an elastic net prediction of exchange rate changes $\Delta s_{i,t+1}$. For the exchange rate predictions, we use eighteen predictors related to past exchange rate changes, volatility, liquidity, skewness, and intermediary capital factors. The OOS R^2 to predict 1-month ahead exchange rate changes are positive but relatively small (0.36% on average). We find that the historical average of returns is a bad proxy. Moreover, the elastic net predictions also do not provide any economic value. That is, the OOS performance of MV is strongest in our baseline construction using the random walk assumption for the exchange rates (further evidence is presented in the Supplementary (Online) Appendix). The reasons are that the predictions are close to zero and the prediction noise is relatively large. The noise has a negative effect on the OOS performance of MV.

For the exchange rate covariance matrix, we employ thirty alternative estimates, input them into MV, and compare the resulting performances with the original MV performance that is based on the robust estimation of the covariance matrix using PCA. The results suggest that our specification of MV in the main text remains the best choice (see the Supplementary (Online) Appendix).

In addition, we test the sensitivity of MV to changes in the threshold value $\overline{\lambda}$ when we decide how many PCs to remove in our PCA. In the main text, we remove all PCs that explain less than $\overline{\lambda} = 1\%$ of the common variation in exchange rate changes. We construct MV and report the unconditional Sharpe ratio of MV for various threshold values $\overline{\lambda}$ in the Supplementary (Online) Appendix. We further provide a brief intuitive motivation of our PCA and argue that our results can be motivated by limits to arbitrage explanation.

Furthermore, to examine the sensitivity of MV's performance to the set of currencies, we repeat our analysis for various subsets of currencies. We show that MV's outperformance is resilient with respect to different subsets of currencies (see the Supplementary (Online) Appendix).

Finally, we extend the robustness checks of our analysis to a larger set of twenty-nine developed and emerging currencies and taking into account transaction costs. The results of these robustness checks broadly uphold our main findings that MV outperforms alternative strategies and market timing is economically important. The results for twenty-nine developed and emerging currencies are repetitive and we not reported them. These results are available upon request.

6. Conclusion

We study the importance of market timing in FX markets. We find that MV (mean–variance optimized portfolio that maximizes unconstraint mean–variance preferences) outperforms alternative mean–variance optimized currency portfolios with leverage, risk (i.e., conditional volatility), or yield targets in our OOS analysis. We conclude that it is costly to impose leverage and risk (i.e., conditional volatility) limits in the portfolio construction, or follow portfolio constructions with suboptimal market timing policies, even if these strategies are conditionally mean–variance efficient. Such suboptimal policies deliver is a significant reduction in the unconditional Sharpe ratio and increase in crash risks.

We further document that the market timing of MV has useful information to predict returns and risks in FX markets. Moreover, MV is not spanned by well-known currency pricing factors, and thus, it poses a new challenge for popular pricing models in the literature.

Appendix: Description of Additional Strategies

This Appendix briefly discusses FX strategies whose definitions are omitted in the main text but are used in the paper.

DOL, DDOL, HML

The Dollar DOL is a traded factor that invests equally in all currencies (Lustig, Roussanov, and Verdelhan, 2011), that is, $\theta_{i,t}^{\mathrm{DOL}} = \frac{1}{N}$. The dynamic or carry Dollar DDOL takes a long (short) position in the DOL when the median forward discount (against the USD) across all exchange rates is positive (negative) (Lustig, Roussanov, and Verdelhan, 2014), $\theta_t^{\mathrm{DDOL}} = \mathrm{sign}(\mathrm{median}(\{f_{j,t}\}_{j=1}^N))\theta_t^{\mathrm{DOL}}$. HML is an equally weighted high-minus-low forward discount carry portfolio introduced in Lustig and Verdelhan (2007). Currencies are sorted into quintiles based on their forward discount $f_{i,t}$ (against the USD). The HML takes a long position in the equally weighted portfolio of currencies in the top quintile and a short position in the equally weighted portfolio of currencies in the bottom quintile.

SW, $\mathrm{SW_{CN}}$

Portfolio SW has weights equal to the forward discounts (against the USD) $\theta_t^{\mathrm{SW}} = f_t$, which refines the strategy HML by taking into account the size and time variation in the forward discounts. Thus, the SW is dynamic and uses market timing based on the absolute size of the forward discounts. Portfolio $\mathrm{SW_{CN}}$ is a rescaled version of SW so that the notional value (i.e., total exposure to the USD) does not change through time.

$\mathrm{HML_{VM}}$, $\mathrm{SW_{VM}}$

These are, respectively, the volatility managed portfolios of HML and SW, following Fleming, Kirby, and Ostdiek (2001) and Moreira and Muir (2017). We compute the conditional volatility of the HML and SW denoted by σ_t^{HML} and σ_t^{SW} using daily returns of these factors over the past month and define the volatility managed factors as $\theta_t^{\mathrm{HML_{VM}}} = \frac{\theta_t^{\mathrm{HML}}}{(\sigma_t^{\mathrm{HML}})^2}$ and $\theta_t^{\mathrm{SW_{VM}}} = \frac{\theta_t^{\mathrm{SW}}}{(\sigma_t^{\mathrm{SW}})^2}$.

MOM

Momentum MOM strategies are popular in equity and FX markets (Burnside, Eichenbaum, and Rebelo, 2011; Menkhoff *et al.*, 2012b). Currencies are sorted on their past 12-month performance into quintiles. The top quintile contains the winner currencies and the bottom quintile the loser currencies. We build equally weighted currency portfolios for each quintile. We denote these five portfolios by $\mathrm{Mom}P_i \; \forall i \in \{1, \ldots, 5\}$. MOM takes a long position in the equally weighted portfolio of winner currencies and a short position in the equally weighted portfolio of loser currencies.

VAL

The value VAL strategy assumes that in the long run undervalued currencies (with low real exchange rates) appreciate against overvalued currencies (with high real exchange rates) (Bilson, 1984; Menkhoff *et al.*, 2017). Currencies are sorted on their past real exchange rates against the USD into quintiles (the top quintile contains overvalued and the bottom quintile undervalued currencies). We proxy the real exchange rate of currency i against USD using the PPP at time t (i.e., the ratio of the prices in currency i and the USD of a representative consumption bundle) multiplied by the nominal exchange rate $X_{i,t}$. VAL takes a long position in the equally weighted portfolio of currencies in the bottom quintile and a short position in the equally weighted portfolio of currencies in the top quintile.

VOL, ILL, SKEW

First, the global FX market volatility at the end of month t is constructed as

$$\widehat{\mathrm{VOL}}_t = \frac{1}{T_t \times N} \sum_{\tau=1}^{T_t} \sum_{I=1}^{N} |\Delta s_{d,i,\tau}|,$$

where $\Delta s_{d,i,\tau}$ is the daily exchange rate growth of currency i against the USD on day τ in month t and T_t is the number of trading days in month t. The VOL_t index is the time series of residuals after estimating an AR(1) process for the $\widehat{\mathrm{VOL}}_t$, that is, $\widehat{\mathrm{VOL}}_t = \rho_v \widehat{\mathrm{VOL}}_{t-1} + \mathrm{VOL}_t$. Thus, VOL captures unexpected changes in the volatility (Menkhoff *et al.*, 2012a). Second, a monthly systematic FX illiquidity measure $\widehat{\mathrm{ILL}}$ is constructed as the average of standardized daily relative bid–ask spread and Corwin and Schultz (2012) liquidity estimate within a month and across all currencies. The ILL_t index is the time series of residuals after estimating an AR(1) process for the $\widehat{\mathrm{ILL}}_t$, that is, $\widehat{\mathrm{ILL}}_t = \rho_{\mathrm{ILL}} \widehat{\mathrm{ILL}}_{t-1} + \mathrm{ILL}_t$. Thus, ILL captures unexpected changes in the illiquidity (Karnaukh, Ranaldo, and Soederlind, 2015). Third, SKEW is constructed to capture the average skewness of exchange rate growths of investment (i.e., positive forward discount) net of funding (i.e., negative forward discount) currencies (Rafferty, 2012),

$$\mathrm{SKEW}_t = \frac{1}{N} \sum_i \mathrm{sign}(f_{i,t-1}) \frac{\frac{1}{T_t} \sum_{\tau}^{T_t} \left(\Delta s_{d,i,\tau} - \overline{\Delta s}_{d,i,\tau} \right)^3}{\left(\frac{1}{T_t} \sum_{\tau}^{T_t} \left(\Delta s_{d,i,\tau} - \overline{\Delta s}_{d,i,\tau} \right)^2 \right)^{\frac{3}{2}}},$$

where $\Delta s_{d,i,\tau}$ is the daily exchange rate growth of currency i against the USD on day τ in month t, $\overline{\Delta s}_{d,i,\tau} = \frac{1}{T_t} \sum_{h=1}^{T_t} \Delta s_{d,i,h}$ is the sample average of the daily exchange rate growth $\Delta s_{d,i,\tau}$ in month t, and T_t is the number of trading days in month t.

Finally, we use two stock market factors, the value weighted US stock market index MKT and the intermediary capital risk factor INT of He, Kelly, and Manela (2017).

Supplementary Material

Supplementary data are available at *Review of Finance* online.

Acknowledgments

We are grateful to the Editor Amit Goyal, an anonymous referee, Anna Battauz, Pedro Barroso, Geert Bekaert, Tony Berrada, Ines Chaieb, Hoyong Choi, Pasquale Della Corte, Carlo Favero, Ilias Filippou, Nicola Gennaioli, Amit Goyal, Bruce Grundy, Regina Hammerschmid, Harald Hau, Thorsten Hens, Alexandra Janssen, Hanno Lustig, Loriano Mancini, Chris Neely, George Panayotov, Joshua Pollet, Peter Puehringer, Angelo Ranaldo, Matt Ringgenberg, Jean-Charles Rochet, Stefano Rossi, Nick Roussanov, Andreas Schrimpf, Fabio Trojani, Andrea Vedolin, Adrien Verdelhan, Guofu Zhou, seminar participants at Bocconi U, St Louis Fed, UIUC, EPFL/U Lausanne, U Geneva, and conference participants at the HKUST finance symposium 2016, FIRN 2016, AFBC 2016, VICIF 2017, MFA 2017, SGF 2017, and EFA 2017 for numerous valuable comments and suggestions.

References

Ackermann, F., Pohl, W., and Schmedders, K. (2016): Optimal and naive diversification in currency markets, *Management Science* 63, 3347–3360.

Baz, J., Breedon, F., Naik, V., and Peress, J. (2001): Optimal portfolios of foreign currencies, *The Journal of Portfolio Management* 28, 102–111.

Bekaert, G. and Panayotov, G. (2018): Good carry, bad carry, *Journal of Financial and Quantitative Analysis*, 55, 1063–1094.

Bilson, J. F. O. (1984): Purchasing power parity as a trading strategy, *Journal of Finance* 39, 715–724.

Brunnermeier, M. K., Nagel, S., and Pedersen, L. H. (2008): Carry trades and currency crashes, in: K. Rogoff, M. Woodford, and D. Acemoglu (eds.), *NBER Macroeconomics Annual 2008*, Vol. 23, The University of Chicago Press, pp. 313–347.

Burnside, C., Eichenbaum, M., and Rebelo, S. (2011): Carry trade and momentum in currency markets, *Annual Review of Financial Economics* 3, 511–535.

Cederburg, S., O'Doherty, M., Wang, F., and Yan, X. (2020): On the performance of volatility-managed portfolios, *Journal of Financial Economics* 138, 95–117.

Chernov, M., Graveline, J., and Zviadadze, I. (2018): Crash risk in currency returns, *Journal of Financial and Quantitative Analysis* 53, 137–170.

Chernov, M., Dahlquist, M., and Lochstoer, L. (2021): Pricing currency risks. Working paper, UCLA and Stockholm School of Economics.

Corwin, S. A. and Schultz, P. (2012): A simple way to estimate bid–ask spreads from daily high and low prices, *Journal of Finance* 67, 719–759.

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

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

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

Dobrynskaya, V. (2014): Downside market risk of carry trades, *Review of Finance* 18, 1885–1913.

Dobrynskaya, V. (2015): Currency exposure to downside risk: Which fundamentals matter?, *Review of International Economics* 23, 245–360.

Ferson, W. and Siegel, A. (2001): The efficient use of conditioning information in portfolios, *Journal of Finance* 56, 967–982.

Fleming, J., Kirby, C., and Ostdiek, B. (2001): The economic value of volatility timing, *Journal of Finance* 56, 329–352.

Galsband, V. and Nitschka, T. (2014): Currency excess returns and global downside risk, *Journal of International Money and Finance* 47, 268–285.

He, Z., Kelly, B., and Manela, A. (2017): Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* 126, 1–35.

Karnaukh, N., Ranaldo, A., and Soederlind, P. (2015): Understanding FX liquidity, *The Review of Financial Studies* 28, 3073–3108.

Kozak, S., Nagel, S., and Santosh, S. (2018): Interpreting factor models, *Journal of Finance* 73, 1183–1223.

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

Lettau, M., Maggiori, M., and Weber, M. (2014): Conditional risk premia in currency markets and other asset classes, *Journal of Financial Economics* 114, 197–225.

Lustig, H. and Verdelhan, A. (2007): The cross-section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.

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

Lustig, H., Roussanov, N., and Verdelhan, A. (2014): Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.

Maurer, T. A., To, T.-D., and Tran, N.-K. (2022): Pricing implications of covariances and spreads in currency markets, *Review of Asset Pricing Studies* 12, 336–388.

Meese, R. and Rogoff, K. (1983): Empirical exchange rate models of the seventies: Do they fit out of sample?, *Journal of International Economics* 14, 3–24.

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

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2012b): Currency momentum strategies, *Journal of Financial Economics* 106, 620–684.

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

Moreira, A. and Muir, T. (2017): Volatility managed portfolios, *Journal of Finance* 72, 1611–1644.

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

Penaranda, F. (2016): Understanding portfolio efficiency with conditioning information, *Journal of Financial and Quantitative Analysis* 51, 985–1011.

Penaranda, F. and Wu, L. (2020): Targets, predictability, and performance. Working paper, CUNY.

Rafferty, B. (2012): Currency returns, skewness and crash risk. Working paper, Duke University.

Ross, S. A. (1976): The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.