Global Currency Hedging with Ambiguity

Urban Ulrych[a,b] Nikola Vasiljević[a,c]

urban.ulrych@bf.uzh.ch nikola.vasiljevic@bf.uzh.ch

[a]*Department of Banking and Finance, University of Zurich, Switzerland*

[b]*Swiss Finance Institute*

[c]*School of Computing, Union University, Belgrade, Serbia*

This Version: November 2022

Abstract

This paper addresses the problem of optimal currency allocation for a risk-and-ambiguity-averse international investor. A robust mean–variance model with smooth ambiguity preferences is used to derive the optimal currency exposure in closed form. In the theoretical part of the paper, we characterize the sample-efficient currency hedging demand as the solution to a generalized ridge regression. Through the lens of these results, we show that the investor's dislike for model uncertainty induces stronger currency hedging demand. The empirical analysis demonstrates how ambiguity leads to a larger estimation bias and simultaneously narrows the confidence interval of the sample efficient optimal currency exposure. The out-of-sample backtest illustrates that accounting for ambiguity enhances the stability of optimal currency allocation over time and significantly reduces portfolio volatility net of transaction costs.

Keywords: Ambiguity Aversion, Currency Hedging, Ridge Regression, International Asset Allocation.

JEL Classification: D81, D83, F31, G11, G15.

The authors are grateful to Yves Allkofer, Maxime Auberson, Marlon Azinović, Stjepan Begušić, John Y. Campbell, Jérôme Detemple, Deniz Erdemlioglu, Walter Farkas, Damir Filipović, Lisa Goldberg, Sean Hofland, Isabella Kooij, Felix Kübler, Markus Leippold, Marc Paolella, Paweł Polak, Andrej Srakar, Michal Svatoň, Serafeim Tsoukas, Alexandre Ziegler, Žan Žurič, and the participants of the SFI Research Days 2019 in Gerzensee, the 26th Forecasting Financial Markets Conference (2019) in Venice, the 2020 American Finance Association Annual Meeting in San Diego, and the 37th International Conference of the French Finance Association (2021) for helpful comments and suggestions. Earlier versions of the paper were circulated under the titles "Optimal currency exposure under risk and ambiguity aversion" and "Ambiguity and the home currency bias".

1 Introduction

Diversification is often said to be the only free lunch in finance. International asset allocation is a natural way to improve risk-adjusted portfolio performance—it opens access to a multitude of global investment opportunities and allows for diversification across asset classes, factors, styles, and geographies. A carefully tailored international portfolio can improve both growth potential and risk management. However, risks can only be reduced and not eliminated.

One of the main challenges for global asset allocation is the currency risk. Fluctuations of exchange rates are influenced by macroeconomic, financial, and political factors. Shifts in currency markets can have a profound impact on portfolio performance. Therefore, investors have to decide on the hedging policy in their portfolios, i.e., the amount of foreign currency exposure that should be hedged.

Increasing complexity of financial markets and products, structural changes caused by secular trends (such as the fourth industrial revolution, aging population, and climate change) or exogenous shocks (such as natural disasters, pandemics, or wars), as well as frequent short-term volatility bouts add to investors' ambiguity aversion (i.e., the aversion to difficult-to-quantify uncertainty) on top of the risk aversion.

In this paper, we study the optimal currency overlay in international portfolios for a risk-and-ambiguity averse investor. Building on the work of Maccheroni, Marinacci, and Ruffino (2013), we account for model uncertainty directly in the investor's ambiguity-adjusted preferences and solve a robust mean–variance optimization problem for currency overlay strategies.

Considering ambiguity in the context of optimal currency allocation is relevant for at least two reasons. First, there is a general dissatisfaction with the empirical and predictive performance of the standard rational expectations paradigm and expected utility theory. Second, many empirical stylized facts cannot be explained using the classical, risk-only models. Embedding ambiguity aversion into financial modeling is currently perceived as one of the most promising avenues for future research.[1]

Our main theoretical contribution is a closed-form expression for the optimal foreign currency exposure under both risk and ambiguity aversion, which can be further decomposed into ambiguity-adjusted hedging and speculative demand. We show that this result can be interpreted as the solution of a generalized ridge regression, where ambiguity aversion drives the intensity of regularization. Therefore, in our model, the optimal foreign currency exposure is inversely related to the ambiguity aversion. In the extreme case of an infinitely ambiguity-averse investor, we find that the optimal strategy is to fully hedge the currency risk.

Overall, the model is analytically tractable and offers an intuitive geometric interpretation of the optimal currency exposure shrinkage in the presence of ambiguity. It represents a generalization of the solution obtained via ordinary least squares regression in the risk-only framework considered in

[1] An excellent literature overview of the applications of ambiguity in asset pricing and portfolio selection is provided in Guidolin and Rinaldi (2013).

Campbell, Serfaty-De Medeiros, and Viceira (2010), which builds on the seminal paper by Britten-Jones (1999).

In the empirical part, we consider currency overlay strategies in a portfolio context. Such strategies are widely applied in the asset and wealth management industry. They are designed for a specific purpose—to separate asset allocation and portfolio construction from currency hedging decisions. Arguably, this approach improves the risk management process. This is particularly important for active managers and their positioning over shorter, tactical investment horizons. Therefore, we follow the common practice and assume that portfolio weights are predefined. This allows us to focus on currencies, and isolate the impact of ambiguity aversion on the optimal foreign currency exposure.

In the first step, we study the effect of ambiguity aversion on the optimal in-sample currency exposure. We demonstrate that ambiguity induces a statistical bias and reduces the standard deviation (i.e., shrinks the confidence interval) of the proposed optimal currency exposure estimator. The decrease in standard error from the risk-only to the ambiguity-adjusted case is the most prominent for highly correlated currencies. Moreover, the in-sample total portfolio volatility is statistically significantly reduced for both risk-only and ambiguity-adjusted investors while they are statistically indistinguishable from each other. In the second step, we show that the ambiguity-adjusted investor performs at least as well as the risk-only investor also out of sample and net of transaction costs. Both investors significantly reduce their portfolio volatility and achieve indistinguishable Sharpe ratios. However, the investor that acknowledges ambiguity achieves this with smaller and more stable currency positions that decrease the amount of hedging turnover over time. Therefore, besides its theoretical appeal, our model is of particular interest also for practical applications of currency hedging in the financial industry.

The paper is organized as follows. Section 2 reviews the literature on optimal currency hedging and asset allocation with ambiguity. Section 3 introduces a theoretical framework for international asset allocation based on a robust mean–variance model under smooth ambiguity preferences studied in Klibanoff, Marinacci, and Mukerji (2005) and Maccheroni, Marinacci, and Ruffino (2013). We provide a closed-form solution to the currency overlay optimization problem for a risk-and-ambiguity-averse investor and a geometric interpretation of our results. Section 4 presents an empirical study in which the impact of ambiguity aversion on the optimal currency exposure is investigated both in sample and out of sample. Section 5 concludes. Auxiliary technical results and descriptions, mathematical proofs, and supplementary figures are delegated to Appendices A and B.

2 Literature Review

Our paper is related to two strands of literature. This section provides an overview of the extant research that laid out the foundation for our study. The first group of papers is related to optimal currency hedging in the context of international asset allocation. The second strand addresses the estimation and model risk, specifically the impact of ambiguity on asset allocation.

2.1 Currency Hedging Policy

A number of hedging strategies have been presented in the literature over the past four decades. Some authors offered pragmatic but rather simplistic solutions to the optimal currency exposure problem. Others considered more nuanced approaches to currency hedging. Overall, there seems to be a broad consensus that (*a*) currency hedging tends to lower portfolio volatility, and (*b*) conditional hedging outperforms strategies that employ fixed hedge ratios. When combined, these results indicate that one should consider both risk minimization and speculative side when deciding about the optimal hedging policy. The empirical results presented in the literature are varied across reference currencies, asset classes, asset allocation mixes, and investment horizons.

Perold and Schulman (1988) proposed full hedging of the face value of a foreign investment. They argued that currency trading is a zero-sum game on average; therefore, hedging represents a "free lunch" because it reduces volatility without a loss of expected return over the long haul. Similar findings were presented in Madura and Reiff (1985) and Jorion (1989), among others.

These results stand in sharp contrast to Froot (1993), who argued that full hedging is beneficial only over short investment horizons. In the long run, full hedging can increase risk without an adequate return compensation. The main argument is that hedged returns are dominated by surprises in inflation and real interest rates. In particular, real interest rates are mean-reverting due to the purchasing power parity, and therefore provide a "natural hedge" over long periods. Since hedging currency exposure does not provide a protection against risk factors affecting long-term exchange rates, the optimal hedging ratio should be zero. However, these results are based on the currency regimes preceding the free-floating exchange rates. Using more recent data, Schmittmann (2010) found that the investment horizon is of limited importance for the hedging policy.

Black (1989, 1990) argued that all investors should apply a "universal hedging policy" irrespective of the portfolio composition and the reference currency. The hedging ratio is always less than one, as a consequence of Siegel (1972)'s paradox. Solnik (1993) challenged this view and showed that the equilibrium hedge ratio is specific for each investor, and it is a function of their risk preferences and relative wealth. Gardner and Wuilloud (1995) and Gastineau (1995) advocated the use of 50% hedge as a baseline currency position. Their "middle-road approach" is an attempt to find a compromise between the universal hedging policy and various practical considerations (e.g., benchmarking, active management, hedging costs, and certain behavioral aspects). Additionally, we note that Ang and Bekaert (2002) obtained a similar hedge ratio using a regime-switching framework.

Glen and Jorion (1993) argued that currency hedging reduces portfolio risk; however it is beneficial only if it does not materially impact the portfolio return. Their analysis indicates that conditional hedging strategies are crucial for the performance of internationally diversified portfolios. Levy and Lim (1994) demonstrated that currency hedging often reduces risk at the expense of lower return. They concluded that this can be explained by the currency forward rate being a biased predictor of the future spot rate. De Roon, Eiling, Gerard, and Hillion (2011) included currency positions as an additional asset class in internationally diversified portfolios. They showed that improvements in portfolio risk-return trade-off are primarily driven by speculative rather than risk hedging benefits. More recently, Boudoukh, Richardson, Thapar, and Wang (2019) formalized these ideas by introducing a modified portfolio mean–variance optimization approach and decomposing the optimal portfolio mix into three components: (a) an equity portfolio, (b) a risk-minimizing currency portfolio, and (c) an alpha-generating currency portfolio.

Jorion (1994) considered a global mean–variance optimization, where positions in assets and currencies can be determined jointly or separately. Joint optimization is studied also in Barroso, Reichenecker, and Reichenecker (2019) and Burkhardt and Ulrych (2022). Either way, the optimal currency exposure strongly depends on the portfolio reference currency. Haefliger, Wydler, and Waelchli (2002) proposed full hedging for fixed-income portfolios, while equity portfolios should be partially hedged or unhedged, depending on the correlations between equity and currency returns. Campbell, Serfaty-De Medeiros, and Viceira (2010) considered risk-minimizing currency overlay strategies for global equity and bond investors. Their analysis critically relies on the correlations between the currencies and core asset classes. For example, low currency–bond correlations indicate that international bond investors should fully hedge their currency exposure. This is consistent with the common practice of institutional investors. On the other hand, due to positive (negative) correlations between commodity (reserve) currencies and global equities, the optimal currency overlay for an international equity investor consists of a long position in reserve currencies and a short position in commodity currencies. A large body of literature on dynamic hedging concluded that currency strategies that are conditional on the interest rate spreads significantly improve portfolio performance compared to static hedging with forward contracts (e.g., Kroner and Sultan, 1993; Tong, 1996; De Roon, Nijman, and Werker, 2003; Brown, Dark, and Zhang, 2012; Caporin, Jimenez-Martin, and González Serrano, 2014; Cho, Min, and McDonald, 2020; Polak and Ulrych, 2021). Additionally, several recent papers which explore the currency predictability using currency risk factors (e.g., Lustig, Roussanov, and Verdelhan, 2011; Verdelhan, 2018; Opie and Riddiough, 2020) demonstrated that dynamic hedging represents a promising route to achieve competitive out-of-sample performance.

Finally, within the realm of currency hedging, a special case of interest is the home currency bias, which represents investors' tendency to disproportionately hold assets denominated in their home currency (e.g., Solnik, 1974; Campbell, Viceira, and White, 2003). Home currency bias is at least partially attributable to the home country bias. Burger, Warnock, and Warnock (2018) showed that this is particularly true in international bond portfolios. Maggiori, Neiman, and Schreger (2020) empirically show that the home currency bias has a substantial impact on global capital allocation. Finally, considering a more general setting, Bianchi and Tallon (2019) and Berger and Eeckhoudt (2021) argue that ambiguity aversion can result in a lower diversification and an elevated portfolio risk.

2.2 Ambiguity and Asset Allocation

Most studies about optimal currency exposure rest on the assumption that investors know perfectly the true probability law governing the asset and currency return dynamics. However, investors are often uncertain about the model validity. The two primary concerns are estimation and model risk. For example, Eun and Resnick (1988) demonstrated that low accuracy of estimated input parameters—in particular the mean returns—is the main driver of poor ex-ante performance of the joint optimization of asset allocations and currency hedge ratios.

A fundamental distinction between risk and uncertainty has been recognized and extensively studied since Knight (1921). Generally, the risk can be defined if all relevant events are associated with a unique probability measure. The (Knightian) uncertainty refers to situations in which some events do not have an obvious probability assignment. In other words, the Knightian uncertainty represents "non-probabilized" uncertainty, as opposed to the risk which corresponds to "probabilized" uncertainty. Ellsberg (1961) provided experimental evidence that agents are not always able to derive a unique probability distribution over the reference state space. This is also known as ambiguity, and a dislike of uncertainty is commonly referred to as ambiguity aversion—in accordance with the literature, we use uncertainty and ambiguity interchangeably. The two most prominent approaches to incorporate uncertainty into investment decision making are Bayesian portfolio analysis and ambiguity-averse preferences.

Bayesian portfolio analysis allows for an inclusion of prior information about various quantities of interest (e.g., asset prices and macroeconomic variables) while accounting for the estimation risk and model uncertainty. An excellent literature review is presented in Avramov and Zhou (2010), who classified Bayesian portfolio studies into three categories based on the assumptions regarding asset return dynamics. The first strand of literature focuses on independently and identically distributed asset returns (e.g., among others, Klein and Bawa, 1976; Jorion, 1986; Black and Litterman, 1992; Pástor, 2000; Pástor and Stambaugh, 2000; Kan and Zhou, 2007; Tu and Zhou, 2010). The priors used in these studies range from uninformative and data-driven to those based on asset pricing theories. The second group of papers considers the possibility that asset returns can be predicted by macroeconomic and fundamental financial variables such as growth, inflation, dividend yield, term, and credit spreads (e.g., Kandel and Stambaugh, 1996; Stambaugh, 1999; Barberis, 2000; Pástor and Stambaugh, 2002; Avramov, 2002, 2004; Avramov and Chordia, 2006; Wachter and Warusawitharana, 2009). The third category comprises alternative, more complex models of asset returns such as stochastic volatility and regime switching (e.g., Han, 2006; Pástor and Stambaugh, 2009; Tu, 2010).

Ambiguity-averse preferences take a leap from the traditional use of rational expectations and maximization of (subjective) expected utility. Uncertainty regarding the asset return predictions is captured directly in the utility function. Intuitively, when a decision maker has too little information to form a single prior, she may consider a set of probability distributions instead. This notion was formalized in

Schmeidler (1989), who stressed that the probability attached to an uncertain event may not reflect the heuristic amount of information that has led to that particular probability assignment. To enable the encoding of information that additive probabilities cannot represent, non-additive probabilities (i.e., capacities) were proposed. Gilboa and Schmeidler (1989) extended this work by introducing multiple prior preferences. Building on these seminal papers, Anderson, Hansen, and Sargent (1998, 2003) and Hansen and Sargent (2001) extended the use of multi-prior criteria to the robust control theory. In that framework, a set of probabilities is generated by statistical perturbations of an approximating model. This approach corresponds to the situations in which agents have a specific model of reference and, acknowledging the possibility of errors, seek robustness against misspecifications.

Klibanoff, Marinacci, and Mukerji (2005) proposed that the ambiguity of a risky event can be characterized by a set of subjectively plausible cumulative probability distributions for this event. The decision maker subjectively weights these distributions, and the resulting preference relation describes the investor's attitude toward ambiguity. Building on this, Maccheroni, Marinacci, and Ruffino (2013) derived the analogue of the classic Arrow–Pratt approximation of the certainty equivalent under model uncertainty as described by the smooth model of decision making under ambiguity. They study its scope by deriving a tractable mean–variance model adjusted for ambiguity and solving the corresponding portfolio allocation problem. We use a similar approach in this work and connect this ambiguity aversion adjusted mean–variance preferences to the optimal currency exposure framework. The analytical tractability of the enhanced Arrow–Pratt approximation renders this model especially well suited for calibration exercises aimed at exploring the consequences of model uncertainty on the optimal currency allocations.

3 Theoretical Framework

In this section, we first derive the optimal currency overlay in the case of a risk-and-ambiguity-averse (RAA) investor. Additionally, we discuss several special cases to provide economic background and intuition. In the second step, we show that a generalized ridge regression framework can be used to recover sample-efficient currency exposures and that it offers an insightful geometric interpretation of the results.

3.1 Optimal Currency Overlay with Ambiguity

3.1.1 General Case: Risk-and-Ambiguity-Averse Investor

We consider an RAA investor who wants to optimize currency exposure in her portfolio by entering positions in currency forward contracts. The asset weights are assumed to be predefined and unaffected by the investor's decision regarding the currency overlay. To account for model uncertainty, we cast the investor's utility function in the Maccheroni, Marinacci, and Ruffino (2013)'s framework. Therefore, we consider a robust optimization that builds on the classical mean–variance expected utility and embeds the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005).

For a period from t to $t+1$, denote the fully hedged portfolio return with $\tilde{R}^{fh}_{\mathcal{P},t+1}$, the excess return in currency c with $e_{c,t+1} - f_{c,t}$, and the exposure to currency c with $\psi_{c,t}$. Then, the hedged portfolio return $\tilde{R}^{h}_{\mathcal{P},t+1}$ can be expressed as

$$\tilde{R}^{h}_{\mathcal{P},t+1} = \tilde{R}^{fh}_{\mathcal{P},t+1} + \sum_{c=2}^{K+1} \psi_{c,t}(e_{c,t+1} - f_{c,t}), \tag{1}$$

where K denotes the number of considered foreign currencies. Additionally, an exposure to a domestic currency, denoted by $c = 1$, is given by $\psi_{1,t} = -\sum_{c=2}^{M+1} \psi_{c,t}$ such that the currency overlay portfolio is a zero investment (i.e., hedging) portfolio. The model-free expression from Eq. (1) that connects the fully hedged portfolio return component and a net currency exposure component is derived in Appendix A.1.

Given the risk and ambiguity aversion parameters λ and θ, the investor's objective is to maximize the ambiguity-adjusted utility function with respect to the vector of net foreign currency exposures $\mathbf{\Psi}_t$:

$$\max_{\mathbf{\Psi}_t} \mathcal{U}(\tilde{R}^{h}_{\mathcal{P},t+1}) = \max_{\mathbf{\Psi}_t} \left\{ \mathrm{E}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{h}_{\mathcal{P},t+1}\,] - \frac{\lambda}{2}\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{h}_{\mathcal{P},t+1}\,] - \frac{\theta}{2}\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{h}_{\mathcal{P},t+1}\,]\,] \right\}, \tag{2}$$

where $\mathbb{Q}$ indexes the probability measures corresponding to candidate models, and μ is the investor's prior over all such probability measures. The probability $\bar{\mathbb{Q}}$ can be interpreted as a weighted average of all probabilities $\mathbb{Q}$ given the prior beliefs μ. Appendix A.2 provides a brief summary of the smooth ambiguity preferences model. A risky and ambiguous prospect in our robust mean–variance optimization problem is represented by the hedged portfolio returns $\tilde{R}^{h}_{\mathcal{P},t+1}$.

The optimal net foreign currency exposure is then obtained in the closed form:

$$\begin{aligned}\mathbf{\Psi}^*_t = &-\Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]\Big)^{-1} \\ &\cdot\Big(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Cov}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\Big),\end{aligned} \tag{3}$$

where the first term (bracket) on the right-hand side is constructed as a positive definite and hence invertible matrix.

The optimal currency exposure can be further broken down into two components which have clear and precise economic meaning: (a) The ambiguity-adjusted hedging demand $\mathbf{\Xi}^*_t$, and (b) The ambiguity-adjusted speculative demand (also known as the ambiguity-adjusted market price of currency risk) $\mathbf{\Lambda}^*_t$:

$$\begin{aligned}\mathbf{\Xi}^*_t := &-\Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]\Big)^{-1} \\ &\cdot\Big(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Cov}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big),\end{aligned} \tag{4a}$$

$$\mathbf{\Lambda}^*_t := \Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big)^{-1} \cdot \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]. \tag{4b}$$

Eqs. (3), (4a) and (4b) represent the key theoretical results of our paper. To gain further insights we examine three special cases of interest below, and provide an alternative representation of the above results in Appendix A.4.

3.1.2 Special Case: Infinite Risk Aversion

The objective of an infinitely risk-averse investor (i.e., $\lambda \to \infty$) is to minimize the variance of portfolio returns. The optimal net foreign currency exposure for such an investor is:

$$\mathbf{\Psi}^*_{t,\text{risk}} := \lim_{\lambda\to\infty} \mathbf{\Psi}^*_t = -\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]^{-1} \cdot \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,]. \tag{5}$$

Although risk aversion dominates in this case, the investor's ambiguity aversion is still reflected via the reduced probability $\bar{\mathbb{Q}}$. In the absence of model uncertainty, the reduced probability $\bar{\mathbb{Q}}$ collapses to a single probability measure (e.g., a historical measure $\mathbb{H}$), and we obtain the classical minimum–variance result as in Campbell, Serfaty-De Medeiros, and Viceira (2010).

The above result indicates that the optimal currency exposure for an infinitely risk-averse investor is reduced to the hedging demand component $\mathbf{\Xi}^*_t$ defined in Eq. (4a), which is driven by the correlation between the fully hedged portfolio returns and excess currency returns. An infinitely risk-averse investor prefers long (short) positions in the currencies which tend to appreciate (depreciate) when the fully hedged portfolio loses value. If the correlation is sufficiently negative (positive), the investor can reduce portfolio risk also by under-hedging (over-hedging). This can be accomplished by holding a long (short) position in excess of the total portfolio weight associated with the investor's base currency. If the correlation is zero, it is optimal to fully hedge the currency risk. Intuitively, an infinitely risk-averse investor is concerned only with risk minimization and does not account for an expected compensation for the additional (currency) risk in her portfolio. Hence, she is better off if the foreign currency exposure is fully hedged.

3.1.3 Special Case: Infinite Ambiguity Aversion

The optimal net foreign currency exposure for an infinitely ambiguity-averse investor (i.e., $\theta \to \infty$) is given by

$$\mathbf{\Psi}^*_{t,\text{amb}} := \lim_{\theta\to\infty} \mathbf{\Psi}^*_t = -\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]^{-1} \cdot \mathrm{Cov}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t]\,]. \tag{6}$$

The matrix $\mathrm{Var}_{\mu}(\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t])$ must be positive definite to ensure its invertibility. The uncertainty is captured by the investor's prior probability over the space of possible models μ.

An infinitely ambiguity-averse investor prefers long (short) positions in the currencies which exhibit a negative (positive) uncertainty-adjusted correlation with the fully hedged portfolio returns. The intuition is very similar to the case of an infinitely risk-averse investor, and the optimal currency exposure is again exclusively driven by the hedging demand.

3.1.4 Special Case: Risky Assets and Ambiguous Currencies

Following Maccheroni, Marinacci, and Ruffino (2013), we consider a setting in which one can distinguish between purely risky and ambiguous assets. We assume that cash represents a risk-free asset, other asset classes such as government bonds and equities are viewed as purely risky when fully hedged back to the domestic currency, whereas currencies are viewed as ambiguous assets. We further assume that the investor holds a fully hedged portfolio and seeks to optimize her foreign currency exposure, given a set of feasible currency overlay strategies across the risk–ambiguity–return spectrum.

Since the underlying portfolio has no foreign currency exposure, it can be treated as purely risky. Thus, $\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,]$ is constant for all models (indexed by $\mathbb{Q}$) and $\mathrm{Cov}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,] = 0$. Using Eq. (3), we compute the optimal currency exposure for a fully hedged investor:

$$\begin{aligned}\mathbf{\Psi}^*_{t,\mathrm{fh}} = -&\Big(\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \frac{\theta}{\lambda}\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]\Big)^{-1} \\ &\cdot\Big(\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \frac{1}{\lambda}\mathrm{E}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\Big).\end{aligned} \tag{7}$$

If the investor is infinitely risk-averse (i.e., $\lambda \to \infty$), the optimal currency exposure converges to the minimum variance case given in Eq. (5). In the case of an infinitely ambiguity-averse investor (i.e., $\theta \to \infty$), the optimal foreign currency exposure converges to zero. Hence, the currency risk remains fully hedged and the investor is maximally biased toward her base currency.

3.2 Generalized Ridge Regression Representation

A common approach in the literature is to investigate the role of sampling error when constructing *ex-post* (i.e., in-sample) efficient portfolio weights or currency exposures. For example, Britten-Jones (1999) showed that the ordinary least squares (OLS) regression of a constant onto excess asset returns—without an intercept term—results in an estimated coefficient vector that represents a set of risky-asset-only portfolio weights for a sample-efficient tangency portfolio. A similar idea is applied in Campbell, Serfaty-De Medeiros, and Viceira (2010) to derive the optimal currency exposure.

In this section, we demonstrate how our setting—which incorporates ambiguity—compares to the extant literature. We focus on the hedging demand to keep our analysis tractable and to facilitate an easier comparison with selected benchmark models. We present only the key theoretical results and the associated notation, whereas detailed derivations are delegated to Appendix A.5. For completeness, we extend our analysis to the general case that includes the speculative demand in Appendix A.6.

3.2.1 Sample Efficiency

Starting with Eq. (2), the optimal in-sample hedging demand (i.e., without the so-called speculative demand) can be computed by solving the following optimization problem:

$$\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H}} = \underset{\boldsymbol{\Psi}_t}{\arg\min} \left\{ \left\|\mathbf{y} - \mathbf{X}(-\boldsymbol{\Psi}_t)\right\|^2_{\mathbf{W}} + \left\|(-\boldsymbol{\Psi}_t) - (-\boldsymbol{\Psi}^*_{t,\text{amb}})\right\|^2_{\mathbf{Z}} \right\}, \tag{8}$$

where $\mathbb{H}$ represents the historical probability measure, $\mathbf{X}$ is the matrix of demeaned historical excess currency returns $\mathbf{e}_{t+1} - \mathbf{f}_t$, and $\mathbf{y}$ is the vector of demeaned historical fully hedged portfolio returns $\tilde{R}^{fh}_{\mathcal{P},t+1}$. The weighting matrix $\mathbf{W} := \frac{\lambda}{T}\mathbf{I}$ is the identity matrix scaled by the investor's risk aversion λ and the number of observations T, and the matrix $\mathbf{Z} := \theta \text{Var}_{\mu}[\,\text{E}_{\mathbb{Q}}[\tilde{R}^h_{\mathcal{P},t+1}]\,]$ measures the model uncertainty. By construction, both matrices are positive definite. Finally, for an arbitrary vector $\mathbf{v}$ and a positive definite matrix $\mathbf{D}$, we define the weighted L^2-norm as $\left\|\mathbf{v}\right\|^2_{\mathbf{D}} = \mathbf{v}'\,\mathbf{D}\mathbf{v}$. The derivation is provided in Appendix A.5.

Equation (8) represents an important theoretical result showing that the optimal sample-efficient currency exposure for an RAA agent can be found by running a generalized ridge regression of the demeaned fully hedged portfolio returns on the demeaned excess currency returns. The shrinkage toward the optimal hedging demand of an infinitely ambiguity-averse investor $\boldsymbol{\Psi}^*_{t,\text{amb}}$, derived in Eq. (6), is driven by the degree and structure of ambiguity. These effects are captured by the matrix $\mathbf{Z}$.

Since both terms in Eq. (8) are convex in $\boldsymbol{\Psi}_t$, a unique solution exists. After some algebra, we obtain the sample-efficient optimal currency exposure for an RAA investor in the following form:

$$\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H}} = -\big(\mathbf{X}'\,\mathbf{W}\mathbf{X} + \mathbf{Z}\big)^{-1}\big(\mathbf{X}'\,\mathbf{W}\mathbf{y} - \mathbf{Z}\boldsymbol{\Psi}^*_{t,\text{amb}}\big). \tag{9}$$

Admittedly, this expression is rather intricate and does not easily lend itself to intuition. For this reason, we consider two special cases of interest.

We start with the risk-only case, which represents the limiting case of Eq. (8) when ambiguity aversion converges to zero. In this setting, there is no utility loss from model uncertainty. The penalty term vanishes and the optimization problem becomes:

$$\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H},\text{risk}} := \widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H}}\Big|_{\theta\to 0} = \underset{\boldsymbol{\Psi}_t}{\arg\min} \left\{ \left\|\mathbf{y} - \mathbf{X}(-\boldsymbol{\Psi}_t)\right\|^2_{\mathbf{I}} \right\}. \tag{10}$$

The optimal in-sample currency exposure is estimated via an ordinary least squares (OLS) regression of the demeaned fully hedged portfolio returns $\mathbf{y}$ on the demeaned excess currency returns $\mathbf{X}$:

$$\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H},\text{risk}} = -(\mathbf{X}'\,\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{y}. \tag{11}$$

This result is in line with the findings presented in Campbell, Serfaty-De Medeiros, and Viceira (2010) and Schmittmann (2010).

The second special case includes ambiguity, but simplifies Eq. (9) and offers additional insights. In particular, we turn our attention to the setting introduced in Section 3.1.4, in which fully hedged assets are assumed to be risky and currencies are treated as ambiguous. The underlying portfolio becomes purely risky since it is fully hedged. Consequently, $\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t]\,] = 0$ and the optimal hedging demand of an infinitely ambiguity-averse investor from Eq. (6) is zero (i.e., $\mathbf{\Psi}^*_{t,\mathrm{amb}} = 0$).

To focus on the impact of ambiguity aversion on the sample-efficient currency exposure, we make two additional assumptions. First, we assume that the uncovered interest rate parity holds (i.e., the forward rate is an unbiased predictor of the future spot rate). Second, we consider a setting that allows for an in-sample analysis without explicitly specifying various predictive models. Without the loss of generality, we assume independence of investor's predictive models and impose the condition $\mathbf{Z} := \frac{\theta}{T^2}\mathbf{I}$. The scaling parameter $1/T^2$ is introduced so that the covariance matrices calculated under $\mathbb{H}$ and μ have the same order of magnitude.

Under these assumptions, our model is transformed into an ordinary (i.e., non-generalized) ridge regression:

$$\widehat{\mathbf{\Psi}}^*_{t,\mathbb{H}} = -(\mathbf{X}'\mathbf{X} + q\mathbf{I})^{-1}\mathbf{X}'\mathbf{y}, \tag{12}$$

where the ridge regularization parameter is given by $q = \frac{\theta}{\lambda T}$. This result shows that ambiguity aversion affects the optimal risk-only hedging demand in Eq. (11) by shrinking it toward zero. The shrinkage effect magnifies (fades away) with increasing ambiguity aversion θ (risk aversion λ). This is intuitive since a higher (lower) ratio of ambiguity to risk aversion (i.e., θ/λ) implies a stronger (weaker) pull toward zero. Additionally, an increasing number of observations improves the precision of sample estimates and reduces parameter and model uncertainty, ultimately resulting in weaker shrinkage effects.

3.2.2 Geometric Interpretation

The OLS regression has a geometric interpretation as an orthogonal projection of the fully hedged portfolio returns onto the space spanned by the excess currency returns. The estimated in-sample currency weights represent a net-zero currency overlay strategy that is the closest to the fully hedged portfolio returns (in the sense of the least-squares distance). The closer the two vectors are, the larger the risk reduction via the currency overlay.

In the presence of ambiguity, the second term in Eq. (8) plays an important role. The estimated optimal ambiguity-adjusted currency exposures still lie in the span of the regression predictors (i.e., the excess currency returns), while a shrinkage of the OLS-based solution is induced. The regularization magnitude θ and the structure of model uncertainty $\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{h}_{\mathcal{P},t+1}\,]\,]$ control the shrinkage toward the target $\mathbf{\Psi}^*_{t,\mathrm{amb}}$. Geometrically, this can be interpreted as an ellipsoidal constraint on the regression coefficients. In the general case, such a constraint is centered at $\mathbf{\Psi}^*_{t,\mathrm{amb}}$, which is not necessarily equal to zero.

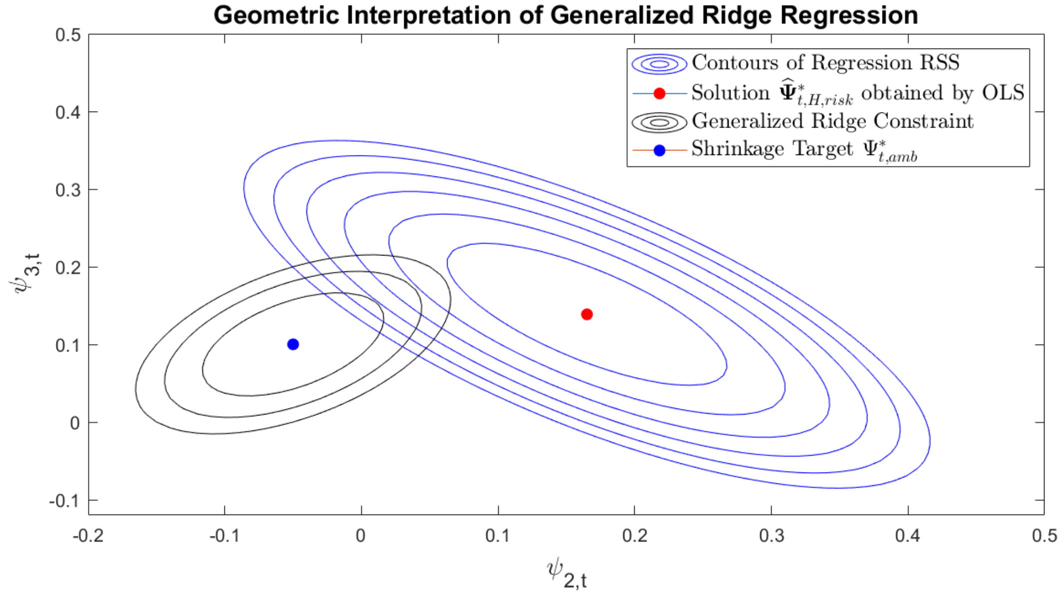

Figure 1 – A geometric interpretation of a generalized ridge regression in two dimensions. Namely, $\psi_{2,t}$ and $\psi_{3,t}$ denote two foreign currency exposures, while $\psi_{1,t}$ would denote the corresponding domestic exposure. The red point depicts the solution obtained by the ordinary least squares regression – the risk-only case $\widehat{\mathbf{\Psi}}^*_{t,\mathbb{H},\mathrm{risk}}$. The contours centered at the red point represent the sums of squared residuals in the regression. The blue point is the shrinkage target $\mathbf{\Psi}^*_{t,\mathrm{amb}}$ and the ellipsoid around that point represents the contours of feasible points given a particular choice of the penalization parameter θ. The set of feasible points shrinks toward the target for larger values of the penalization parameter and vice versa. The shape of the ellipsoid is determined by the positive definite matrix $\mathrm{Var}_\mu[\,\mathrm{E}_\mathbb{Q}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]\,]$ in the weighted L^2-norm. The optimal solution of this optimization problem is given as a tangent of the OLS contour sets and the set of feasible points.

Figure 1 provides an illustration of this effect in the case of two foreign currencies. The red point corresponds to the two-dimensional vector of optimal currency exposure obtained via an OLS regression, whereas the blue point represents the shrinkage target. The black ellipsoid contains all sets of feasible currency exposures. For smaller (larger) regularization coefficient θ, these sets becomes larger (smaller) due to the relaxation (tightening) of the constraint. The shape of the ellipsoid is determined by the model uncertainty matrix $\mathrm{Var}_\mu[\,\mathrm{E}_\mathbb{Q}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]\,]$. The optimal solution is given as the tangent of the contour sets corresponding to the OLS residual sums of squares (RSS) and the set of feasible solutions. Specifically, in the classical ridge regression in Eq. (12), the penalty is given by a scaled identity matrix and the target is set at the origin. This implies that the feasible set is represented by a sphere centered at the origin. In the generalized ridge regression in Eq. (9), the sphere becomes an ellipsoid centered at a particular shrinkage target.

The penalty term in Eq. (8) represents the utility loss due to model uncertainty. Therefore, the risk-minimizing sample-efficient currency exposure $\widehat{\mathbf{\Psi}}^*_{t,\mathbb{H},\mathrm{risk}}$ can be thought of as the "first-best" solution in the absence of ambiguity. Once the model uncertainty is accounted for via the ambiguity aversion θ, the currency overlay $\widehat{\mathbf{\Psi}}^*_{t,\mathbb{H}}$ becomes the "second-best" solution. The adjustment of the optimal currency exposure from the "first-best" to the "second-best" is illustrated in Fig. 1. This simple example demonstrates that there is an equivalence between the robust mean–variance utility and generalized ridge regression representations. As such, it establishes a direct link between the areas of financial economics (i.e., a problem of asset/currency allocation) with statistical learning (i.e., arising from regularization methods).

4 Empirical Analysis

In this section, we turn our attention to the empirical analysis. First, we present the summary data statistics. Second, we study the effects of ambiguity aversion on the optimal in-sample currency exposure. Finally, we design a currency hedging strategy and present the out-of-sample backtesting results. This empirical analysis allows us to gauge the impact of ambiguity-adjusted global currency hedging on the out-of-sample performance of international portfolios.

4.1 Data

Following Campbell, Serfaty-De Medeiros, and Viceira (2010), we assembled a data set covering seven major developed economies: Australia, Canada, Switzerland, Eurozone, the United Kingdom, Japan, and the United States. We have collected daily time series of spot and forward exchange rates, short-term interest rates, broad equity indices, and government bonds indices (with maturities from 3 to 5 years) from January 1999 (i.e., since the dawn of the euro) until December 2019. The data used in our analysis are Refinitiv Datastream exchange rates, short-term rates, and total return indices.[2] This sample allows for an investigation of possible shifts in the optimal currency exposures during or after the major financial and economic crises, e.g., the Dot-Com Bubble 2000–2002, the Global Financial Crisis 2008–2009, and the European Sovereign Debt Crisis 2009–2011.

Summary Statistics

Panel A: Equities	Australia	Canada	Switzerland	Eurozone	UK	Japan	US
Return (%)	8.85	8.75	5.69	7.00	6.94	6.15	7.66
Volatility (%)	15.04	16.22	16.58	21.01	17.63	20.73	18.85
Panel B: Bonds	Australia	Canada	Switzerland	Eurozone	UK	Japan	US
Return (%)	4.93	3.78	2.05	3.55	4.04	1.07	3.79
Volatility (%)	3.26	2.57	1.73	2.44	2.55	1.30	3.47
Panel C: Currencies	AUD	CAD	CHF	EUR	GBP	JPY	USD
Return (%)	1.71	1.17	2.19	0.44	−0.72	0.61	–
Volatility (%)	12.53	9.00	10.79	9.71	9.26	10.23	–

Table 1 – A summary table comprising the annualized means and the standard deviations of equity, bond, and currency returns for a fully hedged investor based in the US dollar over the period from January 1999 until December 2019.

Table 1 reports the annualized mean return and volatility of major equity and government bond indices in local currencies, as well as currency returns for the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), pound sterling (GBP), and Japanese yen (JPY) with respect

[2] "Broad equity index" refers to an index that seeks to be the benchmark for the entire equity market. "Total return index" is an index that assume reinvestment of any cash distributions such as dividends and coupons.

Cross-country Return Correlations

Panel A: Equities	Australia	Canada	Switzerland	Eurozone	UK	Japan	US
Australia	1.00						
Canada	0.22	1.00					
Switzerland	0.33	0.47	1.00				
Eurozone	0.33	0.53	0.73	1.00			
UK	0.35	0.54	0.81	0.81	1.00		
Japan	0.55	0.21	0.31	0.27	0.30	1.00	
US	0.14	0.73	0.50	0.54	0.53	0.12	1.00
Panel B: Bonds	Australia	Canada	Switzerland	Eurozone	UK	Japan	US
Australia	1.00						
Canada	0.15	1.00					
Switzerland	0.27	0.40	1.00				
Eurozone	0.12	0.40	0.49	1.00			
UK	0.25	0.49	0.53	0.55	1.00		
Japan	0.22	0.09	0.15	0.10	0.13	1.00	
US	0.08	0.77	0.34	0.38	0.45	0.08	1.00
Panel C: Currencies	AUD	CAD	CHF	EUR	GBP	JPY	USD
AUD	1.00						
CAD	0.53	1.00					
CHF	0.25	0.30	1.00				
EUR	0.43	0.44	0.68	1.00			
GBP	0.41	0.47	0.41	0.57	1.00		
JPY	0.09	0.24	0.41	0.35	0.28	1.00	
USD	0.29	0.54	0.37	0.44	0.50	0.53	1.00

Table 2 – Cross-country correlations for equity, bond, and currency returns from January 1999 until December 2019. Panels A and B report the correlations of local-currency equity and bond returns, respectively. Panel C reports the correlations of currency returns, averaged across all possible base currencies.

to the US dollar (USD) from January 1999 until December 2019. Table 2 reports the correlations of equity returns (Panel A), government bond returns (Panel B), and currency returns averaged across all possible base currencies (Panel C). Across all three panels, we observe elevated correlations between the US and Canada, as well as among Swiss, Eurozone, and UK markets, reflecting their respective economic integration due to the geographic proximity. All correlation coefficients are positive but relatively limited, hence indicating likely international diversification benefits during the observation period.

4.2 In-Sample Analysis: The Impact of Ambiguity Aversion

We turn our attention now to the in-sample analysis of the impact of ambiguity aversion on the optimal hedging demand. To keep our analysis tractable, we consider the case of risky assets and ambiguous currencies discussed in Sections 3.1.4 and 3.2.1. We leverage Eq. (12) to calculate the optimal currency exposure as a function of the risk and ambiguity aversion parameters.

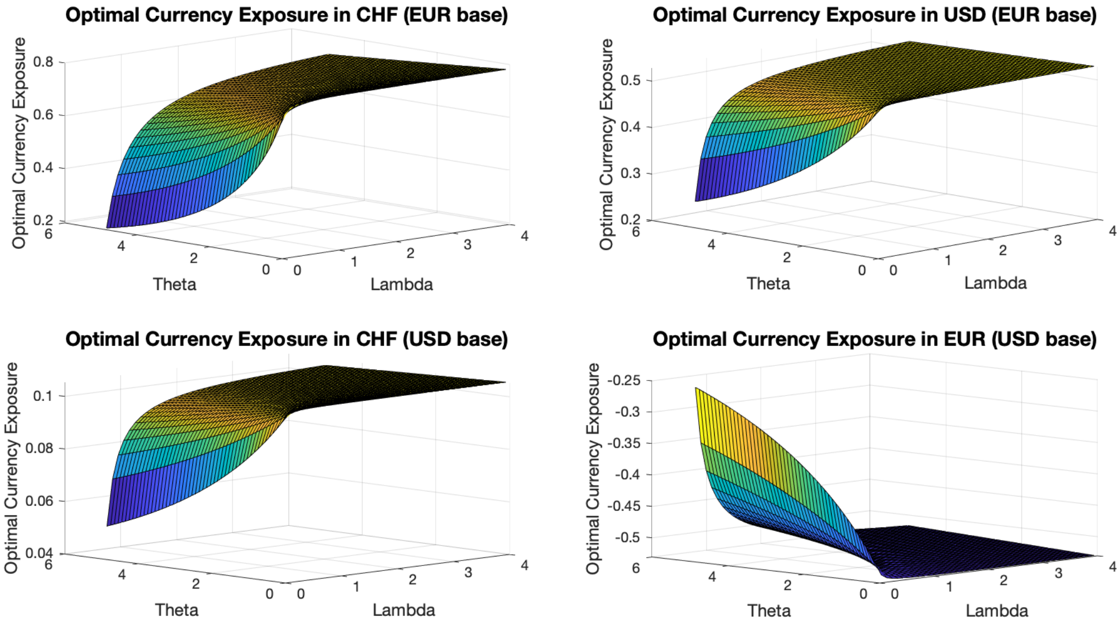

Figure 2 – The optimal in-sample currency exposure in CHF, EUR, and USD for a EUR- and USD-based investor as a function of the risk and ambiguity aversion parameters λ and θ. We assume that (a) the uncovered interest rate parity condition is satisfied, and (b) the prediction models are independent, and it holds that $\widehat{\text{Var}}_\mu[\,\widehat{\text{E}}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] = \frac{1}{T^2}\mathbf{I}$. The sample spans the period from January 1999 until December 2019.

4.2.1 A Sensitivity Analysis of the Optimal Hedging Demand

We start our analysis by assuming that the investor's portfolio consists of two assets: equities and government bonds. Furthermore, the portfolio is exposed only to a single foreign currency. The underlying assets are mixed in a ratio of three (equities) to one (bonds) and are equally weighted between the two analyzed countries (i.e., home and foreign) within each asset class. This simplified example allows us to visualize the optimal currency exposure in dependence of the risk and ambiguity aversion parameters in Fig. 2. The upper (lower) charts provide the optimal currency exposure in CHF and USD (CHF and EUR) for a EUR-based (USD-based) RAA investors.

In the case of an extremely high risk aversion (i.e., when $\lambda \to \infty$), the optimization is reduced to the minimum variance problem. This is equivalent to the solution obtained via an ordinary least squares regression. This case is often considered in the literature (e.g., see Campbell, Serfaty-De Medeiros, and Viceira, 2010). However, for finite values of the risk-aversion parameter λ, ambiguity becomes important. The investor's confidence in the (subjective) predictive models is getting weaker as the ambiguity aversion parameter θ increases. Therefore, an ambiguity-averse investor prefers to hold smaller exposure to foreign currencies. In the limiting case when $\theta \to \infty$, the optimal currency exposure converges to zero, which corresponds to full hedging of the currency risk. We also observe that the convergence rate is inversely related to the level of risk aversion. All these results are fully in line with our theoretical findings in Sections 3.1.2 and 3.2.1.

Interestingly, we observe a particularly pronounced effect of ambiguity aversion on the optimal exposure in Swiss francs for a euro-based investor. For such an investor, the ambiguity-induced shrinkage from the ridge regression is stronger for the Swiss franc than for the US dollar because of a high correlation between the Swiss and Eurozone economies.[3]

Being a function of data (i.e., finite samples), the optimal currency exposure estimator in Eq. (12) is actually a random variable whose properties can be studied as well. To gain further insights, we investigate how different values of λ and θ affect the confidence intervals of the optimal in-sample currency exposure. In the first step, we construct time series of monthly non-overlapping returns. Then, assuming independent and identically distributed (i.i.d.) returns, we are able to estimate the confidence intervals of the optimal sample-efficient currency exposure by the means of non-parametric bootstrapping (i.e., a random sampling with replacement).

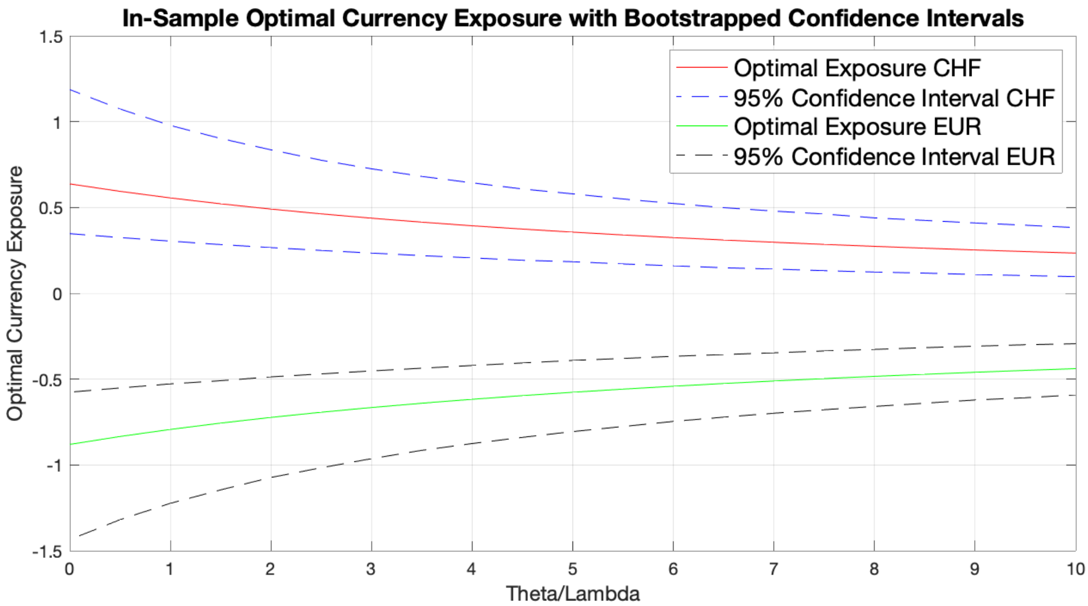

Figure 3 – Optimal currency exposure in CHF and EUR and the corresponding 95% confidence intervals for a USD-based investor, as a function of the risk and ambiguity aversion parameters λ and θ. We assume that (a) the uncovered interest rate parity condition is satisfied, and (b) the prediction models are independent, and it holds that $\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] = \frac{1}{T^2}\mathbf{I}$. The sample spans the period from January 1999 until December 2019.

The key message is that ambiguity simultaneously induces a bias and shrinkage of the confidence intervals for the proposed estimator. This result gives rise to a possible bias-variance trade-off. For example, let us consider a USD-based investor whose portfolio comprises Swiss, Eurozone, and US equities and bonds. The underlying asset mix introduced at the beginning of this section is assumed here as well (with one additional foreign currency). Figure 3 displays the corresponding optimal exposure to the Swiss franc and the euro, jointly with the bootstrapped 95% confidence intervals as a function of the ratio of the ambiguity-to-risk-aversion parameter ratio θ/λ. Using statistical learning terminology, the plotted functions can be interpreted as ridge regularization paths of the optimal currency exposures.

[3] Since ridge regression projects the predictor variables onto the principal components of input data points and then shrinks the coefficients of the low–variance components more than the high-variance components, it effectively addresses some of the shortcomings of the OLS regression in the presence of correlated regressors. We refer an interested reader to Hastie, Tibshirani, and Friedman (2009) for additional information.

This result can be connected to Fig. 1 presented in Section 3.2.2; a given regularization path captures the shift of the red dot (the OLS solution) toward the blue dot (the optimal currency exposure for an infinitely ambiguity-averse investor). Since the ambiguity aversion parameter corresponds to the shrinkage intensity parameter in the ridge regression, an elevated ambiguity aversion leads to a larger estimation bias. Simultaneously, the confidence intervals of the optimal exposure estimator are shrunk. In the limiting case of $\theta \to \infty$, the optimal currency exposures and the confidence intervals for both CHF and EUR converge to zero.

Another important observation is that an elevated correlation between the home and foreign markets and currencies directly impacts the precision of the optimal currency exposure estimate by inflating its standard errors. For example, as mentioned above, the Swiss and Eurozone economies are highly integrated. This is reflected, among other things, in a large and positive correlation coefficient between the Swiss franc and the euro (see Panel C in Table 2). Indeed, we observe wide confidence intervals for $\theta = 0$ in Fig. 3. Moreover, a large positive coefficient for the Swiss franc is partially neutralized by a negative coefficient for the euro. This issue is mitigated by imposing a constraint on the coefficients in the ridge regression via the regularization term. An additional chart—which plots the whole distribution of the optimal exposure estimates corresponding to Fig. 3—is provided in Appendix B.1.

4.2.2 A Comparative Analysis of the Optimal Hedging Demand

In this section, we study the difference in optimal in-sample currency exposures between risk-only (RO) and ambiguity-adjusted (AA) investors. The former (latter) corresponds to the case $\theta = 0$ ($\theta/\lambda = 2$) in Eq. (12). By analogy with Panel B of Table III in Campbell, Serfaty-De Medeiros, and Viceira (2010), Table 3 reports optimal currency exposures for an investor that is fully invested in a single-country equity portfolio and considers all available currencies for hedging.[4] We assume a quarterly hedging horizon and employ a non-parametric bootstrap to estimate the standard deviations of the optimal exposures for both RO and AA investors. The home currency exposure—represented by the diagonal elements in the table—is chosen such that the hedging portfolio is a net-zero-value portfolio.[5]

Table 3 shows that single-country stock investors considering all currencies simultaneously almost always choose positive exposures to the Swiss franc, euro, Japanese yen, and US dollar, and negative exposures to the Australian dollar, Canadian dollar, and British pound. The currencies falling into the first (second) group are often perceived as safe-haven (commodity) currencies. These results are consistent between both investor types. However, we observe a stark contrast between the RO and AA investors in terms of the standard deviations of their respective optimal currency exposures. In particular, the standard deviations are considerably lower for the AA investor, as a consequence of the shrinkage from the ridge regression. The decrease in standard deviations from RO to AA investor is

[4] An equivalent version of Panel A of Table III in Campbell, Serfaty-De Medeiros, and Viceira (2010) studying a difference between RO and AA investors when an exposure to only a single foreign currency is possible is presented in Appendix B.1.

[5] See Eq. (23) in Appendix A.1.3 for further details.

Optimal Currency Exposure:
Single-Country and Multiple-Currency Portfolios

Stock Market	Model	Currency: AUD	CAD	CHF	EUR	GBP	JPY	USD
Australia	RO	-0.45	-0.33	0.32	0.25	-0.54	0.37	0.38
		(0.15)	(0.24)	(0.35)	(0.40)	(0.24)	(0.10)	(0.20)
	AA	-0.52	-0.20	0.28	0.16	-0.35	0.38	0.24
		(0.13)	(0.14)	(0.17)	(0.18)	(0.14)	(0.09)	(0.13)
Canada	RO	-0.85	-0.39	0.27	0.52	-0.35	0.27	0.53
		(0.22)	(0.28)	(0.34)	(0.37)	(0.22)	(0.14)	(0.22)
	AA	-0.70	-0.44	0.27	0.33	-0.19	0.30	0.43
		(0.18)	(0.24)	(0.17)	(0.18)	(0.14)	(0.13)	(0.16)
Switzerland	RO	-0.58	-0.23	0.68	0.08	-0.44	0.29	0.21
		(0.17)	(0.24)	(0.30)	(0.29)	(0.25)	(0.13)	(0.23)
	AA	-0.52	-0.22	0.69	-0.02	-0.33	0.28	0.13
		(0.12)	(0.15)	(0.23)	(0.16)	(0.16)	(0.11)	(0.14)
Eurozone	RO	-0.67	-0.31	0.35	-0.29	-0.81	0.28	1.45
		(0.22)	(0.32)	(0.46)	(0.51)	(0.28)	(0.14)	(0.26)
	AA	-0.65	-0.17	0.28	-0.40	-0.50	0.39	1.05
		(0.16)	(0.19)	(0.27)	(0.34)	(0.19)	(0.12)	(0.17)
UK	RO	-0.78	-0.39	0.25	0.33	-0.02	0.31	0.28
		(0.14)	(0.20)	(0.29)	(0.31)	(0.20)	(0.10)	(0.17)
	AA	-0.68	-0.33	0.24	0.23	0.01	0.32	0.22
		(0.10)	(0.12)	(0.14)	(0.15)	(0.19)	(0.09)	(0.11)
Japan	RO	-0.48	-0.61	0.14	1.04	-0.87	0.81	-0.02
		(0.29)	(0.33)	(0.40)	(0.42)	(0.32)	(0.25)	(0.34)
	AA	-0.45	-0.50	0.24	0.60	-0.56	0.83	-0.16
		(0.20)	(0.19)	(0.22)	(0.20)	(0.19)	(0.24)	(0.22)
USA	RO	-0.67	-0.54	0.38	0.36	-0.43	0.33	0.58
		(0.18)	(0.22)	(0.30)	(0.31)	(0.23)	(0.14)	(0.20)
	AA	-0.61	-0.46	0.32	0.23	-0.31	0.32	0.52
		(0.13)	(0.13)	(0.16)	(0.14)	(0.15)	(0.12)	(0.17)

Table 3 – This table presents optimal currency exposures and their standard deviations for a domestic equity investor who chooses positions from all available currencies to optimize her portfolio. The rows of the table indicate the base country (which corresponds also to the base currency) and the columns represent the optimal positions in currencies for both risk-only (RO) and ambiguity-adjusted (AA) investors. The currency exposures are computed by Eq. (12), where $\theta = 0$ by construction (i.e., ordinary least squares regression) for the RO investor and $\theta/\lambda = 2$ is assumed for the AA investor (i.e., ridge regression). We assume a quarterly hedging horizon and employ a non-parametric bootstrap to estimate the standard deviations of the optimal exposures for both RO and AA investors.

the most prominent for highly correlated currencies (and markets) such as the CAD and USD, as well as the CHF, EUR, and GBP.

The second question that naturally arises in our comparative study is how large the in-sample total risk reductions for global RO and AA investors are. By analogy with Table VII in Campbell, Serfaty-De Medeiros, and Viceira (2010), Table 4 reports the in-sample total portfolio volatilities for the RO and AA investors. These numbers are calculated as the standard deviations of portfolios comprising

Standard Deviations of Hedged Global Equity and Bond Portfolios

Base Currency	Constant Hedge			Optimal Hedge		Tests of Significance									
						RO vs. No Hedge		AA vs. No Hedge		RO vs. Full Hedge		AA vs. Full Hedge		RO vs. AA	
	No Hedge	Half Hedge	Full Hedge	RO	AA	F-Stat	p-value	F-Stat	p-value	F-Stat	p-value	F-Stat	p-value	F-Stat	p-value
Equity															
AUD	11.78	13.06	15.38	9.78	9.91	1.45	9.99	1.42	12.51	2.47	0.01	2.41	0.01	0.97	91.05
CAD	13.70	14.10	15.02	9.85	9.98	1.93	0.38	1.89	0.53	2.32	0.02	2.27	0.04	0.97	90.92
CHF	16.75	15.50	14.65	9.67	9.80	3.00	0.00	2.92	0.00	2.29	0.03	2.24	0.04	0.97	90.95
EUR	15.96	15.08	14.67	9.60	9.70	2.77	0.00	2.71	0.00	2.34	0.02	2.29	0.03	0.98	92.42
GBP	14.68	14.56	15.02	9.71	9.80	2.29	0.03	2.24	0.04	2.39	0.01	2.35	0.02	0.98	93.47
JPY	19.93	16.73	14.21	9.95	10.08	4.01	0.00	3.91	0.00	2.04	0.18	1.99	0.26	0.97	90.57
USD	17.58	15.86	14.62	9.83	9.93	3.20	0.00	3.13	0.00	2.21	0.05	2.17	0.07	0.98	92.88
Bonds															
AUD	10.29	6.65	4.10	3.50	3.52	8.62	0.00	8.56	0.00	1.37	16.34	1.36	17.40	0.99	97.24
CAD	7.87	5.55	3.90	3.37	3.38	5.44	0.00	5.42	0.00	1.34	20.13	1.33	20.68	1.00	98.77
CHF	6.11	4.35	3.91	3.41	3.42	3.22	0.00	3.20	0.00	1.32	22.47	1.31	23.53	0.99	97.81
EUR	6.87	4.88	4.00	3.47	3.48	3.92	0.00	3.89	0.00	1.33	20.42	1.32	21.52	0.99	97.58
GBP	7.97	5.58	4.00	3.46	3.47	5.30	0.00	5.28	0.00	1.34	19.66	1.33	20.20	1.00	98.76
JPY	8.30	4.70	3.84	3.39	3.40	6.00	0.00	5.96	0.00	1.29	26.70	1.28	28.09	0.99	97.46
USD	7.16	4.82	3.83	3.36	3.37	4.54	0.00	4.52	0.00	1.30	24.53	1.29	25.63	0.99	97.88

Table 4 – The in-sample total portfolio volatilities of different constant or model-based hedging strategies. The results for equally-weighted global equity and bond portfolios are studied. Within the equity and bond panels, rows represent the base currencies and columns represent the hedging strategy. “No Hedge”, “Half Hedge”, and “Full Hedge” refer to the constant hedging strategies where zero, half, and all implicit currency exposure is neutralized, respectively. The optimal hedge strategies—the risk-only (RO) and the ambiguity-adjusted (AA) strategy—refer to Eq. (12), where $\theta = 0$ by construction (i.e., ordinary least squares regression) for the RO investor and $\theta/\lambda = 2$ is assumed for the AA investor (i.e., ridge regression). Reported standard deviations are annualized and given in percentage points. The p-values from the tests of significance of portfolio variances are also given in percentage points. All presented results are computed with returns (and hedging) at a quarterly horizon.

the underlying assets (i.e., equities and bonds) and the respective currency overlay strategies. We consider the case of equally-weighted global equity and bond portfolios. For each home currency, we compute the in-sample attained annualized portfolio volatility of quarterly returns for several hedging strategies. The sample frequency is matched with the hedging frequency. The first three columns report the annualized volatilities for three naïve/constant-hedging strategies: no-, half-, and full-hedge. The subsequent two columns report the annualized volatilities for the RO and AA investors, obtained via OLS and ridge regressions, respectively.[6] The remaining numbers in the table are the F-statistics and the corresponding p-values based on our tests of statistical significance of the volatility reductions of both model-based strategies against zero and full hedging, as well as of the RO and AA models against each other.

Based on our comparative analysis, we infer that the in-sample benefit of full hedging for global equity portfolios depends on the investor's base currency. Investors with risk-reducing base currencies such as the JPY or USD considerably reduce the total portfolio volatility by hedging foreign currency exposures. The opposite conclusion holds for investors based in commodity currencies such as AUD and CAD, which are positively correlated with equities. For such base currencies, the total portfolio volatility increases with the hedging ratio.

Next, we conclude that the RO-based hedging reduces the in-sample total portfolio volatility irrespectively of the investor's base currency. A similar magnitude of risk reduction is observed for the AA-based hedging. The estimated risk reduction is statistically significant at the 1% level for RO and AA equity investors against the no-hedge and full-hedge strategies (with an exception of the investors based in AUD, when RO and AA are tested against the naïve no-hedge strategy). When comparing the risk reductions between the RO and AA investors, we see that they are statistically indistinguishable with extremely high p-values of about 90%. While statistically indistinguishable, the AA investor estimates the optimal currency positions with mitigated parameter uncertainty and generally smaller optimal currency positions due to the shrinkage from the ridge regression.

For global bond investors, the in-sample risk reductions from currency hedging are substantial for all base currencies. This is not surprising as most of the risk in international bond portfolios arises rather from fluctuating currencies than from interest rates. Consequently, almost all risk reduction can be attained by fully hedging the currency exposure. Our tests of significance for global bond investors show that the total portfolio volatility reduction stemming from the model-based hedging strategies is not statistically significantly different from the naïve full-hedge strategy. Moreover, the RO and AA models produce almost identical in-sample risk reductions, thus being statistically indistinguishable from each other.

[6] We note that Campbell, Serfaty-De Medeiros, and Viceira (2010) found that the in-sample total portfolio volatilities for the naïve full-hedge and RO-based hedging strategy do not depend on the base currency. However, these authors work with the approximate log portfolio excess returns. Our model is more general and does not include any approximations. Consequently, the resulting in-sample total portfolio volatilities for the naïve full-hedge and RO-based hedging strategy are quite similar but do not match perfectly.

4.3 Out-of-Sample Backtest

In this section, we consider an out-of-sample (OOS) backtest of risk-only (RO) and ambiguity-adjusted (AA) currency overlay strategies for a global portfolio consisting of equities and government bonds, mixed in proportion three-to-one in favor of equities. Within each of the two asset classes, we assume an equally-weighted sub-portfolio consisting of the developed market indices described in Section 4.1. As in the in-sample empirical analysis, we also here study the ambiguity-adjusted currency hedging demand. Appendix B.2 extends this by including also the ambiguity-adjusted speculative currency demand. The backtest is performed for each of the seven base currencies considered in the in-sample analysis above. The OOS performance of the RO and AA currency overlay strategies is compared to three naïve benchmarks, which are specified in terms of rule-based constant-hedging strategies: no-, half-, and full-hedge. Additionally, we investigate the performance of the RO and AA models also in the presence of leverage constraints. These restricted models are labeled RO-Con and AA-Con, respectively.[7] By construction, the risk aversion coefficient vanishes in the RO strategies ($\theta = 0$). In line with our hyperparameter choices in Section 4.2.2, we set $\theta/\lambda = 2$ for the AA strategies.

The currency hedging is implemented using forward contracts which are rolled over quarterly. We assume that the asset positions are also rebalanced quarterly. When a currency forward contract expires, the resulting P&L is reinvested in the assets such that the equal portfolio weighting is preserved over the backtest. A new hedge is formed immediately thereafter and the process is repeated over time. All models are calibrated on a two-year rolling window of daily historical data. Thus, the first two years of data are used for pure calibration and the rest of the data set corresponds to the out-of-sample (rolling-window-based) empirical exercise.

Given the equally weighted sub-portfolios comprising the seven markets introduced above, the strategies RO-Con and AA-Con restrict the currency exposures to the interval $[-\frac{1}{7}, \frac{2}{7}]$. The lower and upper bounds are expressed in terms of relative value compared to the total portfolio value. This choice reflects a symmetric treatment of over- and under-hedging (i.e., no a priori preference for either direction). All results in this section are net of hedging-induced transaction costs, which are assumed to be ten basis points (relative to the notional of a currency forward contract). Moreover, we assume that the asset positions are rebalanced without transaction costs in order to isolate the effects of different currency overlay strategies on the portfolio performance. Our findings are robust to alternative assumptions regarding the size of transaction costs, hedging horizon/rebalancing frequency, and portfolio composition. The corresponding results are available on request.

Table 5 reports the annualized average returns, volatilities, and risk-adjusted performance metrics (i.e., the Sharpe and Sortino ratios) of daily portfolio returns, as well as the average hedging turnover.[8]

---

[7] The unconstrained models correspond either to the OLS or ridge regression framework introduced in Section 3.2. The constrained models are solved via quadratic programming, which is discussed in Section A.4.

[8] We approximate the values of the currency forward contracts used for hedging on a daily basis, which allows us to work with daily returns irrespective of the hedging maturity.

Out-of-Sample Backtest of Currency Overlay Strategies

	Constant Hedge			Model Hedge			
Panel A: AUD	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	5.43	6.35	7.26	7.54	7.07	7.06	7.05
Volatility (%)	9.09	8.05	9.26	7.47	7.25	7.32	7.31
Sharpe Ratio	0.60	0.79	0.78	1.00	0.98	0.96	0.96
Sortino Ratio	0.86	1.10	1.08	1.48	1.38	1.36	1.36
Turnover	0	0.43	0.86	1.68	1.20	1.13	1.09
Panel B: CAD	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	6.04	5.75	5.46	5.54	5.39	5.27	5.40
Volatility (%)	9.47	8.74	9.25	7.47	7.39	7.37	7.45
Sharpe Ratio	0.64	0.66	0.59	0.74	0.73	0.71	0.72
Sortino Ratio	0.90	0.92	0.81	1.07	1.03	1.00	1.02
Turnover	0	0.43	0.86	1.55	1.02	0.93	0.89
Panel C: CHF	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	4.30	4.24	4.18	3.83	3.79	3.49	3.85
Volatility (%)	12.73	10.60	9.24	7.44	7.37	7.33	7.45
Sharpe Ratio	0.34	0.40	0.45	0.52	0.51	0.48	0.52
Sortino Ratio	0.47	0.55	0.62	0.74	0.72	0.66	0.72
Turnover	0	0.43	0.86	1.60	1.29	1.24	1.19
Panel D: EUR	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	5.66	5.34	5.01	5.00	5.20	4.80	5.14
Volatility (%)	10.32	9.47	9.24	7.48	7.35	7.34	7.39
Sharpe Ratio	0.55	0.56	0.54	0.67	0.71	0.65	0.70
Sortino Ratio	0.77	0.78	0.75	0.97	1.00	0.91	0.97
Turnover	0	0.43	0.86	1.39	1.08	1.01	0.97
Panel E: GBP	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	7.27	6.57	5.87	5.78	5.96	5.66	6.04
Volatility (%)	10.42	9.34	9.25	7.47	7.40	7.37	7.48
Sharpe Ratio	0.70	0.70	0.63	0.77	0.81	0.77	0.81
Sortino Ratio	1.00	0.99	0.88	1.12	1.14	1.08	1.14
Turnover	0	0.43	0.86	1.70	1.12	1.06	1.01
Panel F: JPY	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	7.48	5.79	4.10	3.55	3.63	3.37	3.76
Volatility (%)	15.67	12.04	9.23	7.48	7.40	7.34	7.46
Sharpe Ratio	0.48	0.48	0.44	0.48	0.49	0.46	0.50
Sortino Ratio	0.65	0.66	0.61	0.68	0.68	0.64	0.70
Turnover	0	0.43	0.86	1.72	1.22	1.16	1.13
Panel G: USD	Zero	Half	Full	RO	RO-Con	AA	AA-Con
Mean (%)	7.11	6.37	5.64	5.34	5.58	5.20	5.59
Volatility (%)	12.32	10.44	9.24	7.47	7.40	7.35	7.48
Sharpe Ratio	0.58	0.61	0.61	0.72	0.75	0.71	0.75
Sortino Ratio	0.81	0.85	0.84	1.03	1.06	0.99	1.04
Turnover	0	0.43	0.86	1.68	1.20	1.13	1.09

Table 5 – We report the out-of-sample (annualized) average return, volatility, Sharpe ratio, Sortino ratio, and turnover for seven currency overlay strategies (their definitions are provided at the beginning of Section 4.3) and across seven base currencies. The constrained (namely, RA-Con and AA-Con) strategies allow for currency exposures to lie within the interval $[-\frac{1}{7}, \frac{2}{7}]$, expressed in terms of relative value compared to the total portfolio value. We assume $\theta/\lambda = 2$ for the AA strategies, while $\theta = 0$ holds by construction for the RO strategies. The underlying global portfolio consists of equities and bonds mixed in proportion three (equities) to one (bonds). Within each asset class, we assume an equally-weighted sub-portfolio that consists of the corresponding developed market indices.

The statistics are presented for all considered base currencies and across different hedging strategies described above.

First, we compare the currency overlay strategies in terms of the total portfolio volatility. The naïve full-hedge strategy—in comparison to the no-hedge strategy—reduces volatility for all base currencies. This is particularly pronounced for Swiss, Japanese, and US investors, which is intuitive because CHF, JPY, and USD are considered to be safe-haven currencies (i.e., they tend to be less correlated with equities in down markets than other currencies). Next, we stress that all currency overlay strategies outperform constant-hedging benchmarks in terms of the risk-reduction potential. The volatility reduction is the largest for Japanese investors—it is about a half of the volatility generated by the zero-hedging strategy. After consolidating the results across all base currencies, we see that all model-based strategies perform very similarly in terms of volatility reduction. The robust hypothesis test for the difference in variances from Ledoit and Wolf (2011) shows that all models produce a highly statistically significant reduction in variance compared to both zero and full hedging. On the other hand, the null hypothesis of equal variances cannot be rejected between any of the models.

Second, we compare the models in terms of the realized average portfolio returns. Full hedging as well as model-based hedging tend to lower portfolio returns compared to the rule-based no-hedge strategy, with the exception of the AUD-based investor. Note that these results depend also on the level of transaction costs. Importantly, as the sample average portfolio returns carry substantial parameter uncertainty, these results should be interpreted with caution. Nonetheless, this is not surprising, given that the mean of the returns is not entering the utility function of the models analyzed in this section. As an extension, Appendix B.2 provides additional results where the speculative demand for currencies is incorporated in the models.

Third, we investigate the risk-adjusted portfolio performance as measured by the Sharpe and Sortino ratios. Among the constant-hedging strategies, the two performance metrics vary across the base currencies. However, the model-based strategies tend to outperform their rule-based benchmarks. The largest relative increase in the Sharpe and Sortino ratios is observed for Australian and Swiss investors. The robust hypothesis test for the difference in Sharpe ratios from Ledoit and Wolf (2008) shows that these differences are not statistically significant. The only statistically significant case is the zero hedging AUD-based investor in comparison to any of the model-based strategies. Note that statistically significantly distinguishing Sharpe ratios is difficult since Sharpe ratios sensitively depend on the realized average portfolio returns that are noisy estimates of the true mean portfolio returns.

Finally, we study the turnover of our currency overlay strategies. We define the average hedging turnover of a strategy s as:

$$\overline{\mathrm{HT}}^{s} = \frac{1}{T} \sum_{t=1}^{T} \sum_{c=2}^{K+1} |\phi_{c,t}^{s}|, \tag{13}$$

where T is the number of trading days, K is the number of foreign currencies, and $\phi_{c,t}^{s}$ is the notional value of the hedging position for currency c at time t relative to the total portfolio value. We note that this is an average turnover given a rebalancing frequency, which is quarterly in our analysis.

Table 5 demonstrates that the turnover of the constant-hedging strategies remains unchanged across all base currencies. This is because we rebalance the underlying portfolio on the frequency at which the currency hedge is rolled over. By construction, the currency overlay strategies are more reactive to changing market conditions than the constant-hedging strategies, hence generating a higher average turnover. The RO model, based on the OLS estimation, exhibits the highest turnover among all analyzed models. This is expected since no shrinkage takes place in the estimation phase. The turnover of the RO model can be decreased by imposing constraints on the allowed currency exposure, which can be seen in the RO-Con model. Such constraints can be interpreted as a form of shrinkage, see Jagannathan and Ma (2003). On the other hand, the AA model across all base currencies displays decreased turnovers that are smaller even than the ones from the RO-Con model. This decrease occurs because ridge regression shrinks the estimated optimal exposure toward zero, stabilizes the optimal exposure estimates over time, and consequently reduces the hedging turnover. It is interesting to note that the AA-Con model only slightly decreases its turnover compared to the unconstrained AA model. This shows that the ridge regression by itself already shrinks the optimal exposure estimates in or close to the imposed constraint region.

To conclude the out-of-sample backtest analysis, we explore how the estimated optimal currency exposure—in USD for a EUR-based investor and in CHF for a USD-based investor from Table 5—varies over time. Our findings are presented in Fig. 4. The optimal exposure is computed relative to the total portfolio value. To gain further insights into the impact of the ambiguity on the dynamics of the optimal currency exposure, we consider $\theta/\lambda = 2$ for the AA and AA-Con strategies and $\theta/\lambda = 100$ for the AA Extreme strategy. The two constrained models impose optimal exposure to lie on the interval between -1/7 and 2/7 (while each country weight in the portfolio corresponds to 1/7). One can observe the increased stability driven by shrinkage in the AA model compared to its OLS-based RO counterpart. Even in the unconstrained AA case, the optimal currency positions are substantially reduced in comparison to the RO case and most of the time lie close to the AA Constraint counterpart. This reduction is driven by the ambiguity aversion of the investor. In the case of extreme ambiguity aversion, extreme shrinkage is applied and the optimal exposure converges close to full hedging.

The results presented in Table 5 and Fig. 4 show that the ambiguity-adjusted model performs at least as well as the risk-only model out of sample. Both models significantly reduce risk measured as the total portfolio volatility and achieve similar Sharpe ratios that are statistically indistinguishable from each other. However, by acknowledging ambiguity, investors can reduce the risk of their international portfolios by being exposed to smaller currency positions compared to the risk-only case. Moreover, shrinkage that is induced by ambiguity stabilizes the currency positions over time and reduces the amount of hedging turnover out of sample.

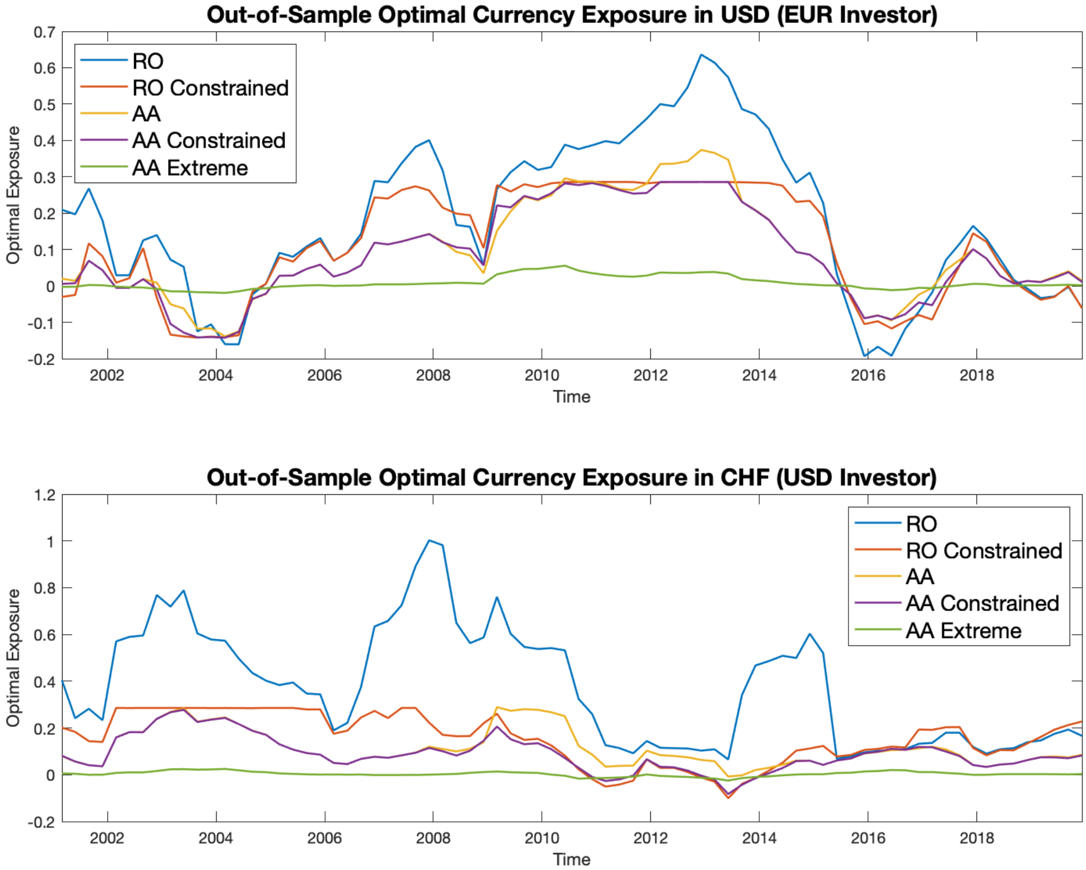

Figure 4 – This figure depicts the evolution of optimal exposures of portfolios analyzed in Table 5 over time. The upper sub-plot shows the optimal exposure in USD for a EUR-based investor and the lower sub-plot portrays the optimal exposure in CHF for a USD-based investor. The RO models correspond to the case of $\theta = 0$, we assume $\theta/\lambda = 2$ for the AA and AA-Con strategies, and $\theta/\lambda = 100$ for the AA Extreme strategy. The latter demonstrates the convergence of optimal currency exposure toward full hedging in the case of extreme ambiguity aversion. RA-Con and AA-Con strategies allow for currency exposures in $[-\frac{1}{7}, \frac{2}{7}]$, expressed in terms of relative value compared to the total portfolio value.

5 Conclusions

The objective of this paper is to study the optimal currency allocation for a risk-and-ambiguity-averse global investor. Utilizing a robust mean–variance utility representation, we derive a closed-form expression for the optimal foreign currency exposure under both risk and ambiguity aversion. Moreover, we show that the sample-efficient currency exposure for a risk-and-ambiguity-averse investor can be found by a generalized ridge regression. This characterization nests the OLS-based approaches commonly presented in the existing literature as a special case without ambiguity aversion. The generalized penalty term of the ridge regression corresponds to the utility loss arising from the model uncertainty and induces regularization. This regularization, biasing of the estimator that stabilizes the inference, is here not assumed a priori. It originates as a solution to the robust mean–variance maximization problem. This shows that accounting for ambiguity enables the formal relation between the areas of financial economics and statistical learning. Furthermore, in a setting where currencies are treated as ambiguous assets whereas other assets are risky, an extreme ambiguity aversion implies the optimality of full hedging for the risk-and-ambiguity-averse investor. This shows that elevated ambiguity aversion can lead to insufficient currency diversification of global investors.

In the empirical part of the work, we explore the impact of ambiguity aversion on optimal currency allocations in and out of sample. We show that acknowledging ambiguity corresponds to an increase in bias and a simultaneous shrinkage of the confidence intervals of the sample efficient optimal currency exposure. Moreover, we show that both risk-only and ambiguity-adjusted investors significantly reduce total portfolio volatility in sample. However, the ambiguity-adjusted investor achieves this with reduced optimal exposure standard errors. The out-of-sample empirical exercise demonstrates that accounting for ambiguity stabilizes the optimal exposure estimates over time. While being statistically indistinguishable from the risk-only investor, the ambiguity-adjusted investor significantly reduces total portfolio risk also out-of-sample and net of transaction costs. This improvement compared to the risk-only case is obtained with smaller currency positions and reduced hedging turnover driven by regularization arising from acknowledging ambiguity in the hedging decision.

This paper introduces a benchmark model that combines the risk and the ambiguity aversion of international investors. The model is intuitive since it allows for closed-form solutions and a geometric interpretation of obtained results but is involved enough to study the impact of investor's ambiguity on the optimal currency allocation in global portfolios. Given its underlying assumption of independent and identically distributed returns, our model does not capture empirical stylized facts present in the financial data. Therefore, this paper allows for possible generalizations in several research directions that can later be benchmarked to our model. For example, further implications of risk and ambiguity in higher moments can be pursued. Thereby, robust preferences capturing the coskewness and cokurtosis of asset and currency returns can be investigated. One can study their effect on the optimal currency exposure and home currency bias, potentially in closed form. Another important consideration can also be the study of ambiguity around higher moments of the currency return distribution. Moreover, one can include the effect of transaction costs directly in the optimized utility function. This would correspond to an investor seeking to improve the risk-ambiguity-return characteristics in a cost-efficient manner. The transaction costs would, in this case, correspond to the L^1-normed penalty, such as in Brodie, Daubechies, De Mol, Giannone, and Loris (2009). The generalized ridge regression arising in our model can then potentially be extended to the generalized elastic net regularized regression–a method that linearly combines the L^1 and L^2 penalties of the lasso and ridge methods.

From the empirical point of view, it might be valuable to investigate emerging market currencies and indices jointly with the already explored developed markets. It is also possible to extend an empirical analysis by using different asset classes, such as corporate bonds and commodities. For example, possible differences in optimal currency exposure of a portfolio consisting of investment grade vs. high yield bonds of developed and emerging markets can be investigated. Another potential prospect is studying specific sectors, such as energy, agriculture, and precious metals, and observing if particular sectors in various market environments co-vary with currencies in specific ways. This relates to the growing academic interest in sector rotation investment strategies, where one can examine how currencies relate to the corresponding macro variables used for forecasting and sector rotation decisions. One can as well analyze specific investment styles, such as value, size, momentum, and low volatility, and investigate the relation between currency exposure and different types of factor models.

References

E. W. Anderson, L. P. Hansen, and T. J. Sargent. Risk and robustness in general equilibrium. *Preprint University of Chicago*, 1998.

E. W. Anderson, L. P. Hansen, and T. J. Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1): 68–123, 2003.

R. W. Anderson and J.-P. Danthine. Cross hedging. *Journal of Political Economy*, 89(6):1182–1196, 1981.

A. Ang and G. Bekaert. International asset allocation with regime shifts. *The review of financial studies*, 15(4):1137–1187, 2002.

D. Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64 (3):423–458, 2002.

D. Avramov. Stock return predictability and asset pricing models. *The Review of Financial Studies*, 17(3):699–738, 2004.

D. Avramov and T. Chordia. Asset pricing models and financial market anomalies. *The Review of Financial Studies*, 19(3):1001–1040, 2006.

D. Avramov and G. Zhou. Bayesian portfolio analysis. *Annu. Rev. Financ. Econ.*, 2(1):25–47, 2010.

N. Barberis. Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1): 225–264, 2000.

P. Barroso, J. A. Reichenecker, and M. R. Reichenecker. Let the Parametric Phoenix Fly. *Available at SSRN 3421686*, 2019.

L. Berger and L. Eeckhoudt. Risk, ambiguity, and the value of diversification. *Management Science*, 67(3):1639–1647, 2021.

M. Bianchi and J.-M. Tallon. Ambiguity preferences and portfolio choices: Evidence from the field. *Management Science*, 65(4):1486–1501, 2019.

F. Black. Universal hedging: Optimizing currency risk and reward in international equity portfolios. *Financial Analysts Journal*, 45(4):16–22, 1989.

F. Black. Equilibrium exchange rate hedging. *The Journal of Finance*, 45(3):899–907, 1990.

F. Black and R. Litterman. Global portfolio optimization. *Financial analysts journal*, 48(5):28–43, 1992.

J. Boudoukh, M. Richardson, A. Thapar, and F. Wang. Optimal currency hedging for international equity portfolios. *Financial Analysts Journal*, 75(4):65–83, 2019.

M. Britten-Jones. The sampling error in estimates of mean-variance efficient portfolio weights. *The Journal of Finance*, 54(2):655–671, 1999.

J. Brodie, I. Daubechies, C. De Mol, D. Giannone, and I. Loris. Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences*, 106(30):12267–12272, 2009.

C. Brown, J. Dark, and W. Zhang. Dynamic currency hedging for international stock portfolios. *Review of futures markets*, 20:419–455, 2012.

J. D. Burger, F. E. Warnock, and V. C. Warnock. Currency matters: Analyzing international bond portfolios. *Journal of International Economics*, 114:376–388, 2018.

R. Burkhardt and U. Ulrych. Sparse and Stable International Portfolio Optimization and Currency Risk Management. *Swiss Finance Institute Research Paper No. 22-07*, 2022.

J. Y. Campbell, L. M. Viceira, and J. S. White. Foreign currency for long-term investors. *The Economic Journal*, 113(486):C1–C25, 2003.

J. Y. Campbell, K. Serfaty-De Medeiros, and L. M. Viceira. Global currency hedging. *The Journal of Finance*, 65(1):87–121, 2010.

M. Caporin, J.-A. Jimenez-Martin, and L. González Serrano. Currency hedging strategies in strategic benchmarks and the global and euro sovereign financial crises. *Journal of International Financial Markets, Institutions and Money*, 31:159–177, 2014.

Y.-W. Cheung, M. D. Chinn, A. G. Pascual, and Y. Zhang. Exchange rate prediction redux: New models, new data, new currencies. *Journal of International Money and Finance*, 95:332–362, 2019.

J.-B. Cho, H.-G. Min, and J. A. McDonald. Volatility and dynamic currency hedging. *Journal of International Financial Markets, Institutions and Money*, 64:101163, 2020.

C. Conflitti, C. De Mol, and D. Giannone. Optimal combination of survey forecasts. *International Journal of Forecasting*, 31(4):1096–1103, 2015.

F. De Roon, E. Eiling, B. Gerard, and P. Hillion. Speculative profits or hedging benefits? Currency investing in global portfolios. *Working Paper*, 2011.

F. A. De Roon, T. E. Nijman, and B. J. Werker. Currency hedging for international stock portfolios: The usefulness of mean–variance analysis. *Journal of Banking & Finance*, 27(2):327–349, 2003.

W. Du, A. Tepper, and A. Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018.

M. R. Eaker and D. M. Grant. Cross-hedging foreign currency risk. *Journal of International Money and Finance*, 6(1):85–105, 1987.

D. Ellsberg. Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, pages 643–669, 1961.

C. S. Eun and B. G. Resnick. Exchange rate uncertainty, forward contracts, and international portfolio selection. *The Journal of Finance*, 43(1):197–215, 1988.

K. A. Froot. Currency hedging over long horizons. Technical report, National Bureau of Economic Research, 1993.

L. Gagnon, G. J. Lypny, and T. H. McCurdy. Hedging foreign currency portfolios. *Journal of Empirical Finance*, 5(3):197–220, 1998.

G. W. Gardner and T. Wuilloud. Currency risk in international portfolios: How satisfying is optimal hedging. *Journal of Portfolio Management*, 21(3):59, 1995.

G. L. Gastineau. The currency hedging decision: A search for synthesis in asset allocation. *Financial Analysts Journal*, 51(3):8–17, 1995.

I. Gilboa and D. Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

J. Glen and P. Jorion. Currency hedging for international portfolios. *The Journal of Finance*, 48(5): 1865–1886, 1993.

M. Guidolin and F. Rinaldi. Ambiguity in asset pricing and portfolio choice: A review of the literature. *Theory and Decision*, 74(2):183–217, 2013.

T. Haefliger, D. Wydler, and U. Waelchli. Hedging currency risk: Does it have to be so complicated? *Available at SSRN 363080*, 2002.

Y. Han. Asset allocation with a high dimensional latent factor stochastic volatility model. *The Review of Financial Studies*, 19(1):237–271, 2006.

L. Hansen and T. J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91 (2):60–66, 2001.

T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.

R. Jagannathan and T. Ma. Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps. *The Journal of Finance*, 58(4):1651–1683, 2003.

P. Jorion. Bayes-Stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis*, pages 279–292, 1986.

P. Jorion. Asset allocation with hedged and unhedged foreign stocks and bonds. *Journal of Portfolio Management*, 15(4):49, 1989.

P. Jorion. Mean/variance analysis of currency overlays. *Financial Analysts Journal*, 50(3):48–56, 1994.

R. Kan and G. Zhou. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, pages 621–656, 2007.

S. Kandel and R. F. Stambaugh. On the predictability of stock returns: an asset-allocation perspective. *The Journal of Finance*, 51(2):385–424, 1996.

R. W. Klein and V. S. Bawa. The effect of estimation risk on optimal portfolio choice. *Journal of financial economics*, 3(3):215–231, 1976.

P. Klibanoff, M. Marinacci, and S. Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

F. H. Knight. *Risk, uncertainty and profit*. Houghton Mifflin, 1921.

K. F. Kroner and J. Sultan. Time-varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis*, pages 535–551, 1993.

O. Ledoit and M. Wolf. Robust Performance Hypothesis Testing with the Sharpe Ratio. *Journal of Empirical Finance*, 15(5):850–859, 2008.

O. Ledoit and M. Wolf. Robust performances hypothesis testing with the variance. *Wilmott*, 55:86–89, 2011.

H. Levy and K. C. Lim. Forward exchange bias, hedging and the gains from international diversification of investment portfolios. *Journal of International Money and Finance*, 13(2):159–170, 1994.

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

F. Maccheroni, M. Marinacci, and D. Ruffino. Alpha as ambiguity: Robust mean–variance portfolio analysis. *Econometrica*, 81(3):1075–1113, 2013.

J. Madura and W. Reiff. A hedge strategy for international portfolios. *The Journal of Portfolio Management*, 12(1):70–74, 1985.

M. Maggiori, B. Neiman, and J. Schreger. International currencies and capital allocation. *Journal of Political Economy*, 128(6):2019–2066, 2020.

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

L. Pástor. Portfolio selection and asset pricing models. *The Journal of Finance*, 55(1):179–223, 2000.

L. Pástor and R. F. Stambaugh. Comparing asset pricing models: an investment perspective. *Journal of Financial Economics*, 56(3):335–381, 2000.

L. Pástor and R. F. Stambaugh. Mutual fund performance and seemingly unrelated assets. *Journal of Financial Economics*, 63(3):315–349, 2002.

L. Pástor and R. F. Stambaugh. Predictive systems: Living with imperfect predictors. *The Journal of Finance*, 64(4):1583–1628, 2009.

A. F. Perold and E. C. Schulman. The free lunch in currency hedging: Implications for investment policy and peformance standards. *Financial Analysts Journal*, 44(3):45–50, 1988.

P. Polak and U. Ulrych. Dynamic Currency Hedging with Non-Gaussianity and Ambiguity. *Swiss Finance Institute Research Paper No. 21-60*, 2021.

D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica: Journal of the Econometric Society*, pages 571–587, 1989.

J. M. Schmittmann. Currency hedging for international portfolios. *International Monetary Fund Working Paper*, 10/151, 2010.

J. J. Siegel. Risk, interest rates and the forward exchange. *The Quarterly Journal of Economics*, pages 303–309, 1972.

B. Solnik. Currency hedging and Siegel's paradox: On Black's universal hedging rule. *Review of International Economics*, 1(2):180–187, 1993.

B. H. Solnik. An equilibrium model of the international capital market. *Journal of economic theory*, 8 (4):500–524, 1974.

R. F. Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54(3):375–421, 1999.

A. Timmermann. *Chapter 4: Forecast Combinations.* Handbook of Economic Forecasting, Vol. 1. Elsevier, 2006.

W. H. Tong. An examination of dynamic hedging. *Journal of International Money and Finance*, 15 (1):19–35, 1996.

J. Tu. Is regime switching in stock returns important in portfolio decisions? *Management Science*, 56 (7):1198–1215, 2010.

J. Tu and G. Zhou. Incorporating economic objectives into Bayesian priors: Portfolio choice under parameter uncertainty. *Journal of Financial and Quantitative Analysis*, pages 959–986, 2010.

W. N. van Wieringen. Lecture notes on ridge regression, Version May 2021. Vrije Universiteit Amsterdam.

A. Verdelhan. The share of systematic variation in bilateral exchange rates. *The Journal of Finance*, 73(1):375–418, 2018.

J. A. Wachter and M. Warusawitharana. Predictable returns and asset allocation: Should a skeptical investor time the market? *Journal of Econometrics*, 148(2):162–178, 2009.

Appendix A Additional Theoretical Results

A.1 Foreign Currency Exposure and Portfolio Returns

In this section, we consider currency hedging with forward contracts in an international portfolio context. Building on the work of Eun and Resnick (1988), closed-form expressions for hedged and unhedged portfolio returns are derived. These results are model-free in the sense that no specific dynamics for the asset and currency returns are assumed.

A.1.1 Single-Asset Case

Our starting point is a simple example of an unhedged foreign currency-denominated asset. More specifically, we consider an asset i whose price in a given currency at time t is denoted by $P_{i,t}$. The simple return on this asset from t to $t+1$ is given by $R_{i,t+1}$.[9] We introduce $S_{c,t}$ as the spot exchange rate in the home/domestic/reference currency (HC) per unit of foreign/local currency (LC) at time t, where c denotes the local currency of asset i. Alternatively, for financial assets that have embedded multi-currency exposure (e.g., global equity indices), c_i can be interpreted as a single currency in which prices of asset i are quoted on markets. The exchange rate return from t to $t+1$ is $e_{c,t+1}$. The unhedged return on asset i expressed in the home currency is

$$\tilde{R}^u_{i,t+1} = \frac{P_{i,t+1}S_{c,t+1}}{P_{i,t}S_{c,t}} - 1 = \underbrace{R_{i,t+1}}_{\text{LC asset return}} + \underbrace{e_{c,t+1}}_{\text{FX return}} + \underbrace{R_{i,t+1}e_{c,t+1}}_{\text{cross-product}}. \tag{14}$$

The above equation demonstrates that the unhedged asset return in home currency is driven by three components: (a) The asset return in local currency, (b) The exchange rate (FX) return, and (c) The second-order cross-product between the first two terms.

A.1.2 Multi-Asset Case: Unhedged Portfolio Returns

The above result can be generalized from a single-asset case to a portfolio context. We consider an investor—with an arbitrary home currency—who holds a portfolio $\mathcal{P}$ which consists of N assets. The fraction of wealth invested in asset $i = 1, 2, \ldots, N$ is defined as $x_{i,t}$. Let us further assume that this portfolio has a direct exposure to K foreign currencies. For simplicity, we label the home currency by $c = 1$, whereas the foreign currencies are denoted by $c = 2, 3, \ldots, K+1$. The assets can be classified and grouped by their local currency. The collection of all assets denominated in currency c held in a portfolio at time t is denominated by $\mathcal{A}_{c,t}$. The fraction of wealth directly exposed to currency c

[9] The following notation is applied throughout the paper. Local currency returns are denoted by $R_{\star,\star}$, where the first (second) subscript indicates the investment position (time). For example, $R_{i,t}$ is a return on asset i in period t in local currency. Home currency returns are denoted by $\tilde{R}^{\star}_{\star,\star}$. The subscripts bear the same meaning as for the local currency returns, whereas the superscript indicates currency hedging. For example, $\tilde{R}^h_{i,t}$ and $\tilde{R}^u_{i,t}$ represent hedged and unhedged home currency returns, respectively. To ease notation, we drop the indexation of the home currency in mathematical expressions. All results derived in this paper hinge on the assumption that a home currency is prespecified.

is defined as $w_{c,t} := \sum_{j \in \mathcal{A}_{c,t}} x_{j,t} \neq 0$. We introduce $R_{\mathcal{P}_c,t+1} := \sum_{j \in \mathcal{A}_{c,t}} \frac{x_{j,t}}{w_{c,t}} R_{j,t+1}$ as the return on sub-portfolio $\mathcal{P}_c$ that consists of all assets denominated in currency c. From Eq. (14) it follows that the unhedged return on sub-portfolio $\mathcal{P}_c$ expressed in an investor's home currency can be computed as

$$\tilde{R}^u_{\mathcal{P}_c,t+1} = R_{\mathcal{P}_c,t+1} + e_{c,t+1} + R_{\mathcal{P}_c,t+1} e_{c,t+1}. \tag{15}$$

Equations (14) and (15) usher in two representations for the unhedged return on portfolio $\mathcal{P}$:

$$\tilde{R}^u_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t} \tilde{R}^u_{i,t+1} = \sum_{c=1}^{K+1} w_{c,t} \tilde{R}^u_{\mathcal{P}_c,t+1}, \tag{16}$$

with $\sum_{i=1}^{N} x_{i,t} = \sum_{c=1}^{K+1} w_{c,t} = 1$, for every t. Note that this setting permits short selling as the only condition is that the asset weights sum up to 1. By substituting Eq. (15) in the second representation in Eq. (16) we obtain the following result:

$$\tilde{R}^u_{\mathcal{P},t+1} = \underbrace{\sum_{c=1}^{K+1} w_{c,t} R_{\mathcal{P}_c,t+1}}_{\text{LC sub-portfolio returns}} + \underbrace{\sum_{c=2}^{K+1} w_{c,t} e_{c,t+1}}_{\text{FX returns}} + \underbrace{\sum_{c=2}^{K+1} w_{c,t} R_{\mathcal{P}_c,t+1} e_{c,t+1}}_{\text{cross-products}}. \tag{17}$$

A.1.3 Multi-Asset Case: Hedged Portfolio Returns

Having derived the expression for unhedged portfolio returns, we turn our attention to hedging via currency overlay strategies that can be implemented using forward exchange contracts. The forward exchange rate in home currency per unit of foreign currency c at time t is denoted by $F_{c,t}$. The price of the forward contract is assumed to be zero at the inception. We consider the forward contract with delivery date $t+1$. The forward premium, i.e., the return on the (short) forward contract, is defined as $f_{c,t} := (F_{c,t} - S_{c,t})/S_{c,t}$. At time t, this quantity is known and represents the cost of carry. Note that $F_{1,t} = 1$ and $f_{1,t} = 0$ trivially hold for all t. We denote by $\phi_{c,t}$ the amount invested at time t in a forward exchange contract for currency c, expressed in the home currency as a fraction of total portfolio value. Similarly, $\phi_{c,t}/S_{c,t}$ represents the relative notional value of the forward contract in local currency c at time t. Hedging of foreign currency exposure can be achieved by selling a forward exchange contract (i.e., $\phi_{c,t} > 0$). By the no-arbitrage principle, this is analogous to shorting foreign bonds and holding domestic bonds (or alternatively, borrowing funds in foreign currency and lending the same amount in home currency). The pay-off at time $t+1$ is $F_{c,t} - S_{c,t+1}$.

Furthermore, we extend the opportunity set by assuming that the total number of foreign currencies in which investors can trade is $M \geq K$. Therefore, an investor can enter positions in currencies to which her portfolio is not directly exposed. This optionality allows investors to expand the investment universe and implement various cross-hedging strategies. Anderson and Danthine (1981) and Eaker and Grant (1987) showed that cross-hedging can reduce portfolio risk, albeit to a lesser extent than direct hedging. However, in the absence of direct hedging instruments, cross-hedging strategies provide a valuable alternative. Accounting for portfolio effects due to cross-hedging substantially improves efficiency and utility gains, see Gagnon, Lypny, and McCurdy (1998). Moreover, cross-hedging portfolio effects hold even in the case when the hedging opportunity set is restricted to the currencies to which investor's portfolio is directly exposed. For these reasons, cross-hedging is commonly applied in the wealth and asset management industry.

A hedged portfolio return is then equal to

$$\tilde{R}^{h}_{\mathcal{P},t+1} = \tilde{R}^{u}_{\mathcal{P},t+1} + \sum_{c=2}^{M+1} \phi_{c,t}(f_{c,t} - e_{c,t+1}), \tag{18}$$

where $f_{c,t} - e_{c,t+1} = (F_{c,t} - S_{c,t+1})/S_{c,t}$ represents the normalized payoff of a short (for $\phi_{c,t} > 0$) forward contract on currency c at time $t+1$. Since $S_{1,t} = F_{1,t} = 1$ for all t, the choice of $\phi_{1,t}$ is arbitrary. Nevertheless, we aim to keep the interpretation of $\phi_{c,t}$ as a fraction of the total portfolio value corresponding to the notional of the forward contract on currency c. Therefore, we impose the following condition:

$$\phi_{1,t} = 1 - \sum_{c=2}^{M+1} \phi_{c,t}. \tag{19}$$

Consequently, all (net) currency exposures sum up to zero, i.e., the currency portfolio is a zero investment portfolio. For a portfolio that is directly exposed to currency c (i.e., meaning $w_{c,t} \neq 0$), the hedge ratio can be defined as $h_{c,t} \coloneqq \phi_{c,t}/w_{c,t}$. If $\phi_{c,t} = h_{c,t} = 0$, the assets denominated in currency c are unhedged. Conversely, if $\phi_{c,t} = w_{c,t}$ or $h_{c,t} = 1$, the assets are fully hedged.

Fully hedged returns can be computed by setting $\phi_{c,t} = w_{c,t}$ for $c = 1, 2, \ldots, K+1$, and $\phi_{c,t} = 0$ for $c = K+2, K+3, \ldots, M+1$ in Eq. (18). Cross-hedging with additional currencies is excluded in this case. Therefore, the fully hedged portfolio return can be expressed as

$$\tilde{R}^{fh}_{\mathcal{P},t+1} = \underbrace{\sum_{c=1}^{K+1} w_{c,t} R_{\mathcal{P}_c,t+1}}_{\text{LC returns}} + \underbrace{\sum_{c=2}^{K+1} w_{c,t} f_{c,t}}_{\text{FX forward premia}} + \underbrace{\sum_{c=2}^{K+1} w_{c,t} R_{\mathcal{P}_c,t+1} e_{c,t+1}}_{\text{cross-products}}. \tag{20}$$

Comparison of this result with the expression for unhedged portfolio return in Eq. (17) reveals that currency returns are replaced by forward premia in the case of a fully hedged portfolio. Therefore, currency hedging eliminates randomness stemming from the exchange rate fluctuations. This is due to the fact that—for any foreign currency c—the forward premium $f_{c,t}$ is known at time t. However, the exchange rate risk is not completely eliminated since the second-order cross-product term remains. The hedging positions are entered at time t, and the exact currency exposure at time $t+1$ is unknown in advance.[10]

[10] Over tactical investment horizons and in most market environments, the cross-product term is relatively small and can be neglected. Hence, nearly perfect (full) currency hedging can typically be achieved in practice. However, over longer investment horizons (e.g., beyond one year) or in the case of a sudden market crash, these terms could significantly impact the portfolio returns. To address this issue, one possible solution is to review currency hedging

Unhedged and fully hedged portfolio returns represent two special cases of interest, which lie on the opposite sides of the currency hedging spectrum. To accommodate partial hedging of direct currency exposure and cross-hedging with additional currencies, we define the net exposure to currency c as

$$\psi_{c,t} \coloneqq w_{c,t} - \phi_{c,t}. \tag{21}$$

Here, $w_{c,t}$ represents the direct currency exposure and $\phi_{c,t}$ reflects the position in a forward contract. We distinguish among four cases of currency hedging: (a) No hedging ($\psi_{c,t} = w_{c,t}$), (b) Full hedging ($\psi_{c,t} = 0$), (c) Partial hedging ($0 < \psi_{c,t} \leq w_{c,t}$), and (d) Over-hedging and under-hedging ($\psi_{c,t} < 0$ and $\psi_{c,t} > w_{c,t}$, respectively).

Cross-hedging strategies with the remaining $M - K$ currencies can be implemented through forward contracts as well. For $c = K + 2, K + 3, \ldots, M + 1$, the direct currency exposure is zero (i.e., $w_{c,t} = 0$), and the net exposure is $\psi_{c,t} = -\phi_{c,t}$. Therefore, Eq. (18) can be rewritten in the following form:

$$\begin{aligned} \tilde{R}^{h}_{\mathcal{P},t+1} &= \tilde{R}^{fh}_{\mathcal{P},t+1} + \sum_{c=2}^{K+1} \psi_{c,t}(e_{c,t+1} - f_{c,t}) - \sum_{c=K+2}^{M+1} \phi_{c,t}(e_{c,t+1} - f_{c,t}) \\ &= \tilde{R}^{fh}_{\mathcal{P},t+1} + \sum_{c=2}^{M+1} \psi_{c,t}(e_{c,t+1} - f_{c,t}). \end{aligned} \tag{22}$$

Additionally, from Eq. (19) we compute the net home currency exposure:

$$\psi_{1,t} = - \sum_{c=2}^{M+1} \psi_{c,t}. \tag{23}$$

Eqs. (18) and (22) are mathematically equivalent, however, they offer different economic interpretations. The former (latter) equation decomposes portfolio returns into unhedged (fully hedged) returns and a currency hedging (net currency exposure) component.

A.1.4 Embedding a Currency Model: The Covered Interest Rate Parity

The results derived in Appendix A.1.1–A.1.3 are model-free. As such, they represent a cornerstone on top of which a modeling framework can be superposed. To illustrate this, in this section we consider a particularly relevant case in point—the covered interest rate parity model.

Let us denote the nominal risk-free interest rate in currency c by $r_{c,t}$. Based on the notation introduced in Section 3, $r_{1,t}$ represents the risk-free interest rate in the home currency, whereas other risk-free interest rates correspond to the foreign currencies. The covered interest rate parity (CIRP) asserts that $F_{c,t}/S_{c,t} = (1 + r_{1,t})/(1 + r_{c,t})$.[11] Consequently, the forward premium can be computed as

policies when market conditions change substantially.

[11] Among others, Du, Tepper, and Verdelhan (2018) showed that CIRP condition has been systematically violated after the Global Financial Crisis of 2008–2009. The CIRP deviations increase (a) at the quarter ends following the crisis, (b) when banks' balance sheet costs rise, (c) when other risk-free spreads increase, and (d) when nominal interest rates rise.

$f_{c,t} = (r_{1,t} - r_{c,t})/(1 + r_{c,t}) \approx r_{1,t} - r_{c,t}$. Plugging this approximation into Eq. (22) yields the following result for hedged excess portfolio returns:

$$\tilde{R}^h_{\mathcal{P},t+1} - r_{1,t} \approx \underbrace{\sum_{c=1}^{K+1} w_{c,t}(R_{\mathcal{P}_c,t+1} - r_{c,t})}_{\text{LC excess sub-portfolio returns}} + \underbrace{\sum_{c=2}^{M+1} \psi_{c,t}(e_{c,t+1} - r_{1,t} + r_{c,t})}_{\text{excess exchange rate return}} + \underbrace{\sum_{c=2}^{K+1} w_{c,t} R_{\mathcal{P}_c,t+1} e_{c,t+1}}_{\text{cross-products}}.$$

Therefore, the total return can be decomposed into three components: (a) Allocation-weighted average of the excess local-currency returns on sub-portfolios $\mathcal{P}_c$, (b) Net-exposure-weighted average of CIRP payoffs, and (c) Allocation-weighted average of cross-products between the local-currency returns on sub-portfolios $\mathcal{P}_c$ and the corresponding exchange rate returns. ♣

A.2 Smooth Ambiguity Preferences and Robust Mean–Variance Optimization

This section provides an overview of the workhorse model in our framework—Maccheroni, Marinacci, and Ruffino (2013)'s robust mean–variance optimization, which is rooted in the Arrow–Pratt approximation of the certainty equivalent in the case of a risk-and-ambiguity-averse (RAA) agent who maximizes von Neumann–Morgenstern expected utility.

First, let us define a state space Ω that consists of all possible realizations of uncertainty. Sets of states of nature are called events ω, and the outcome space represents a σ-algebra $\mathcal{F}$, which contains random payoffs of an investor's decisions. A preference relation is defined over the mapping from Ω to $\mathcal{F}$. In a risk-only setting, all agents agree on the probability measure $\mathbb{P}$. To capture model uncertainty we introduce a space Δ of possible models $\mathcal{Q}$. We further assume that an agent's prior over all probability measures $\mathbb{Q}$ corresponding to the models in $\mathcal{Q}$ is given by μ. Then, we can compute the reduced probability $\bar{\mathbb{Q}} := \int_\Delta \mathbb{Q}\, d\mu(\mathbb{Q})$ induced by the prior μ—also called the barycenter of μ—which plays an important role in our setting.[12]

To account for model uncertainty, Maccheroni, Marinacci, and Ruffino (2013) (MMR) proposed a robust optimization that builds on the classical mean–variance expected utility framework and embeds the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005) (KMM). The utility function is given by

$$\mathcal{U}(\ell) = \mathrm{E}_{\bar{\mathbb{Q}}}[\,\ell\,] - \frac{\lambda}{2}\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\ell\,] - \frac{\theta}{2}\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\ell\,]\,], \tag{24}$$

where ℓ is an uncertain prospect, and positive coefficients λ and θ represent the risk and ambiguity aversion, respectively. Therefore, the MMR model provides additional flexibility in a computationally tractable and economically meaningful way. Terms $\mathrm{E}_{\bar{\mathbb{Q}}}[\,\cdot\,]$ and $\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\cdot\,]$ represent the reduced-probability estimators of the mean and variance obtained by combining predictions from different $\mathcal{Q}$-models. The underlying weighting scheme is captured by the agent's prior μ. The estimator $\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\cdot\,]$ measures

[12] Two remarks are due. First, the probability measure $\bar{\mathbb{Q}}$ is called the reduction of μ on Ω since it can be interpreted in terms of reduction of compound lotteries. Second, if $\mathrm{supp}(\mu) = \{\mathbb{Q}_1, \mathbb{Q}_2, \ldots, \mathbb{Q}_n\}$ is finite and $\mu(\mathbb{Q}_i) = \mu_i$ for $i = 1, 2, \ldots, n$, then $\bar{\mathbb{Q}}(A) = \mu_1\mathbb{Q}_1(A) + \mu_2\mathbb{Q}_2(A) + \cdots + \mu_n\mathbb{Q}_n(A)$, for any event (i.e., set of outcomes) $A \in \mathcal{F}$.

risk under the reduced probability measure, whereas the estimator $\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\cdot\,]\,]$ quantifies the model uncertainty via the dispersion of different $\mathcal{Q}$-predictions with respect to μ. The latter can be computed as

$$\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\ell\,]\,] = \int_\Delta \left(\int_\Omega \ell(\omega)\, d\mathbb{Q}(\omega) \right)^2 d\mu(\mathbb{Q}) - \left(\int_\Delta \left(\int_\Omega \ell(\omega)\, d\mathbb{Q}(\omega) \right) d\mu(\mathbb{Q}) \right)^2. \tag{25}$$

Higher estimates of the model uncertainty are indicative of investor's lower confidence in a single model. Conversely, if the investor's prior μ is a singleton (i.e., when represented by the Dirac function), the prospect is regarded as purely risky. In that case, all probability measures induced by the set of models $\mathcal{Q}$ are then mapped to the probability measure $\mathbb{P}$, which corresponds to the risk-only setting. Furthermore, the third term in Eq. (24) vanishes in the absence of ambiguity, and the model collapses into the classical mean–variance utility. Investors with exceptionally high or low confidence in a single model occupy the two ends of the ambiguity spectrum. Ambiguity-neutral investors fall somewhere in-between these two extremes, and their expectations are formed based on the reduced probability $\bar{\mathbb{Q}}$ instead of a single probabilistic model. ♣

A.3 Optimal Currency Overlay with Ambiguity: The Main Result

To ease our notation, we introduce additional conventions. First, the vector of exchange rate returns is given by $\mathbf{e}_{t+1} = (e_{2,t+1}, e_{3,t+1}, \ldots, e_{M+1,t+1})'$. Second, the vector of forward exchange premia is $\mathbf{f}_t = (f_{2,t}, f_{3,t}, \ldots, f_{M+1,t})'$. Using the linearity of expectations and the variance sum law, we obtain the following results for the expected return, risk, and model uncertainty for the investor's portfolio:

$$\mathrm{E}_{\bar{\mathbb{Q}}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,] = \mathrm{E}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,] + \mathbf{\Psi}_t'\,\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,], \tag{26a}$$

$$\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,] = \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,] + \mathbf{\Psi}_t'\,\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\mathbf{\Psi}_t + 2\mathbf{\Psi}_t'\,\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,], \tag{26b}$$

$$\begin{aligned}\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]\,] &= \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,]\,] + \mathbf{\Psi}_t'\,\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\mathbf{\Psi}_t \\ &\quad + 2\mathbf{\Psi}_t'\,\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,].\end{aligned} \tag{26c}$$

We note that these equations[13] prominently feature the normalized forward payoff, i.e., the difference between $\mathbf{e}_{t+1}$ and $\mathbf{f}_t$. Under the CIRP, the forward premium $f_{c,t}$ is the difference between the risk-free interest rates in foreign and home currency. Therefore, the random variable $e_{c,t+1} - f_{c,t}$ can be interpreted as the currency excess return. This highlights the importance of hedging decision trade-off between entering a forward contract and retaining net currency exposure.

By plugging the results from Eqs. (26a) to (26c) into Eq. (2), rearranging the terms, and dropping those which are not a function of $\mathbf{\Psi}_t$, we obtain the loss function in a quadratic programming format:

$$\mathbf{\Psi}_t^* = \underset{\mathbf{\Psi}_t}{\arg\min} \quad \left\{ \frac{1}{2}\mathbf{\Psi}_t'\mathbf{A}\mathbf{\Psi_t} + \mathbf{b}'\mathbf{\Psi_t} \right\}, \tag{27}$$

[13] Observe that the term $\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,]$ denotes the $(M \times 1)$-dimensional vector of covariances between $\tilde{R}^{fh}_{\mathcal{P},t+1}$ and $e_{c,t+1} - f_{c,t}$, for $c = 2, \ldots, M+1$, and equivalently for $\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]$.

where $\mathbf{A}$ is an $(M \times M)$-dimensional symmetric matrix and $\mathbf{b}$ is an $(M \times 1)$-dimensional vector given by

$$\mathbf{A} = \lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,] + \theta \mathrm{Var}_{\mu}[\, \mathrm{E}_{\mathbb{Q}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]\,], \tag{28a}$$

$$\mathbf{b} = \lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\, \tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t \,] + \theta \mathrm{Cov}_{\mu}[\, \mathrm{E}_{\mathbb{Q}}[\, \tilde{R}^{fh}_{\mathcal{P},t+1} \,], \mathrm{E}_{\mathbb{Q}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]. \tag{28b}$$

Note that an arbitrary set of linear constraints on the currency exposures could be added to the loss function in Eq. (27), leading to a linearly constrained quadratic program. Linear constraints can be relevant for practical applications because they are often based on regulatory requirements. For example, pension funds use such constraints to prevent excessive currency positions and therefore avoid unnecessary risks. However, adding restrictions comes at a cost because an analytical solution to a constrained optimization problem is not attainable. Therefore, in a general case, the solution is obtained numerically through a quadratic programming procedure.

Taking a derivative of Eq. (27) with respect to $\mathbf{\Psi}_t$ yields the first-order condition

$$-\mathbf{A}\mathbf{\Psi}_t - \mathbf{b} = \mathbf{0}, \tag{29}$$

where $\mathbf{0}$ is a $(M \times 1)$ vector of zeros. The Hessian matrix of second-order derivatives is $\mathrm{Hess}(\mathbf{\Psi}_t) = -A$. Since components of the random vector $(\mathbf{e}_{t+1} - \mathbf{f}_t)$ are linearly independent, the corresponding covariance matrix $\mathrm{Var}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]$ is positive definite. Furthermore, risk and ambiguity aversion coefficients (i.e., λ and θ) are positive scalars, and $\mathrm{Var}_{\mu}[\, \mathrm{E}_{\mathbb{Q}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]\,]$ is a positive (semi-) definite matrix. This implies that the Hessian matrix is negative definite. Equivalently, matrix $\mathbf{A}$ defined in Eq. (28a) is positive definite and therefore invertible. We conclude that the first-order condition from Eq. (29) is a necessary and sufficient condition to characterize the (unconstrained) maximum of the currency overlay optimization problem from Eq. (2) with respect to the vector of (unconstrained) net foreign currency exposures $\mathbf{\Psi}_t$. ■

A.4 Optimal Currency Overlay with Ambiguity: An Alternative Representation

Let us consider an ambiguity-neutral agent (i.e., the case $\theta \to 0$). This case is equivalent to the mean–variance setting where all uncertainty is captured by the reduced probability $\bar{\mathbb{Q}}$. The optimal net foreign currency exposure can be expressed as the sum of the optimal minimum–variance currency portfolio from Eq. (5) (i.e., hedging demand) and the market price of currency risk (i.e., speculative demand), adjusted for the risk aversion λ, i.e.,

$$\begin{aligned} \mathbf{\Psi}^*_{t,\mathrm{mv}} := \lim_{\theta \to 0} \mathbf{\Psi}^*_t &= -\mathrm{Var}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]^{-1} \cdot \Big(\mathrm{Cov}_{\bar{\mathbb{Q}}}[\, \tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t \,] - \frac{1}{\lambda} \mathrm{E}_{\bar{\mathbb{Q}}}[\mathbf{e}_{t+1} - \mathbf{f}_t] \Big) \\ &= \mathbf{\Psi}^*_{t,\mathrm{risk}} + \frac{1}{\lambda} \mathrm{Var}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]^{-1} \cdot \mathrm{E}_{\bar{\mathbb{Q}}}[\, \mathbf{e}_{t+1} - \mathbf{f}_t \,]. \end{aligned} \tag{30}$$

The market price of currency risk represents the trade-off between the expected excess return and risk associated with the respective currency exposure. Its effect on the optimal mean–variance currency exposure vanishes as $\lambda \to \infty$. Moreover, the optimal mean–variance currency exposure can be interpreted as an unambiguous currency overlay strategy. While the minimum variance strategy only seeks to reduce risk, the mean–variance currency exposure introduces an additional component that aims to generate an excess return (i.e., the "alpha"), hence representing a speculative currency demand.[14]

Assuming positive definiteness of $\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]$, the optimal currency exposure in Eq. (3) can be rearranged as follows:

$$
\begin{aligned}
\mathbf{\Psi}_t^* = \quad & -\Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big)^{-1} \\
& \cdot \Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,](\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,](\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& - \theta \mathrm{Var}_\mu(\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& + \theta \mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big) \\
= \quad & -\Big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big)^{-1} \qquad (31) \\
& \cdot \Big((\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]) \\
& \quad \cdot (\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& \quad - \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,](\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}(\lambda \mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& \quad + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\mathrm{Var}_\mu(\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,])^{-1}\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big) \\
= \quad & -\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]^{-1}(\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \frac{1}{\lambda}\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& + \big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\big)^{-1}\theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] \\
& \quad \cdot \big(\mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]^{-1}(\mathrm{Cov}_{\bar{\mathbb{Q}}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] - \frac{1}{\lambda}\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]) \\
& \quad - \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]^{-1}\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\big) =,
\end{aligned}
$$

Finally, we obtain the following representation of the optimal foreign currency exposure:

$$
\mathbf{\Psi}_t^* = \mathbf{\Psi}_{t,\mathrm{mv}}^* + \big(\lambda \mathrm{Var}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]\big)^{-1} \theta \mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,] \cdot (\mathbf{\Psi}_{t,\mathrm{amb}}^* - \mathbf{\Psi}_{t,\mathrm{mv}}^*), \quad (32)
$$

This result represents an explicit correction of the optimal mean–variance currency exposure for an RAA investor, relative to a purely risk-averse investor. The term $(\mathbf{\Psi}_{t,\mathrm{amb}}^* - \mathbf{\Psi}_{t,\mathrm{mv}}^*)$ measures the difference between the optimal currency exposures for an infinitely ambiguity-averse and an unambiguous investor. The multiplier of this term accounts for the size and direction of the parameter and model uncertainty

[14] We reiterate that, in our framework, the decision regarding the optimal currency risk-return trade-off is independent of the asset allocation policy for the underlying portfolio. Moreover, we stress that currency overlay strategies do not represent direct investments. They are merely risk management strategies—in our case implemented with forward contracts—that do not require any portfolio rebalancing or additional investments.

correction. This term vanishes for $\theta \to 0$ and it can be interpreted as the share of the total excess currency return variance which can be attributed to ambiguity (up to the scaling factors λ and θ). Alternatively, Eq. (32) can be interpreted as a convex combination of the two limiting cases $\mathbf{\Psi}^*_{t,\text{mv}}$ and $\mathbf{\Psi}^*_{t,\text{amb}}$. ■

A.5 The Sample-Efficient Currency Exposure for an RAA Investor

Our empirical loss function is defined as a negative of the robust mean–variance utility function, i.e., $\mathcal{L}_{\mathbb{H}}(\tilde{R}^h_{\mathcal{P},t+1}) := -\mathcal{U}(\tilde{R}^h_{\mathcal{P},t+1})$. Therefore, the minimization of the loss function is equivalent to the maximization of the robust mean–variance utility with respect to the in-sample currency exposure

$$\widehat{\mathbf{\Psi}}^*_{t,\mathbb{H}} := \arg\min_{\mathbf{\Psi}_t} \mathcal{L}_{\mathbb{H}}(\tilde{R}^h_{\mathcal{P},t+1}) = \arg\min_{\mathbf{\Psi}_t} \left\{ \frac{\lambda}{2} \widehat{\text{Var}}_{\mathbb{H}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,] + \frac{\theta}{2} \text{Var}_{\mu}[\,\text{E}_{\mathbb{Q}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]\,] - \widehat{\text{E}}_{\mathbb{H}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,] \right\}, \quad (33)$$

where $\widehat{\text{Var}}_{\mathbb{H}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]$ and $\widehat{\text{E}}_{\mathbb{H}}[\,\tilde{R}^h_{\mathcal{P},t+1}\,]$ are the sample covariance matrix and the vector of sample averages of hedged portfolio returns, respectively. The index $\mathbb{Q}$ denotes probability measures that capture the parameter and model uncertainty (as perceived by the investor during the sample period).

As discussed in Section 3.1, a total currency demand is a sum of the (ambiguity-adjusted) hedging and speculative demands. Therefore, we can look at these two cases separately. The sample-efficient regression interpretation of the ambiguity-adjusted hedging demand $\widehat{\mathbf{\Xi}}^*_{t,\mathbb{H}}$ is provided in this section, whereas Appendix A.6 provides the equivalent representation also for the ambiguity-adjusted speculative demand $\widehat{\mathbf{\Lambda}}^*_{t,\mathbb{H}}$.

Consider the demeaned historical returns which come from the historical measure $\mathbb{H}$. We denote by $\mathbf{X}$ and $\mathbf{y}$ the $(T \times M)$ matrix of demeaned historical excess currency returns $\mathbf{e}_{t+1} - \mathbf{f}_t$, and the $(T \times 1)$ vector of demeaned historical fully hedged portfolio return $\tilde{R}^{fh}_{\mathcal{P},t+1}$, respectively, where T is the total number of observations. Thus, the sample covariance matrix of excess currency returns and the vector of covariances between the fully hedged portfolio and excess currency returns are given by

$$\widehat{\text{Var}}_{\mathbb{H}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] = \frac{1}{T}\mathbf{X}'\,\mathbf{X},$$
$$\widehat{\text{Cov}}_{\mathbb{H}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}, \mathbf{e}_{t+1} - \mathbf{f}_t\,] = \frac{1}{T}\mathbf{X}'\,\mathbf{y}.$$

Next, we define $\mathbf{W} := \frac{\lambda}{T}\mathbf{I}$, where $\mathbf{I}$ is a $(T \times T)$ identity matrix, $\mathbf{Z} := \theta \text{Var}_{\mu}[\,\text{E}_{\mathbb{Q}}[\tilde{R}^h_{\mathcal{P},t+1}]\,]$, and $\mathbf{z_0} := -\mathbf{\Psi}^*_{t,\text{amb}}$, assuming that $\text{Var}_{\mu}[\,\text{E}_{\mathbb{Q}}[\mathbf{e}_{t+1} - \mathbf{f}_t]\,]$ is positive definite and its inverse exists. Consequently, the loss function can be expressed as

$$\mathcal{L}_{\mathbb{H}}(\tilde{R}^h_{\mathcal{P},t+1}) = \frac{1}{2}\mathbf{\Psi}_t{}'\,\mathbf{X}'\,\mathbf{W}\mathbf{X}\mathbf{\Psi}_t + \mathbf{\Psi}_t{}'\,\mathbf{X}'\,\mathbf{W}\mathbf{y} + \frac{1}{2}\mathbf{\Psi}_t{}'\,\mathbf{Z}\mathbf{\Psi}_t + \mathbf{\Psi}_t{}'\,\mathbf{Z}\mathbf{z_0} + \text{“remaining terms”}, \quad (34)$$

where we explicitly write the terms which depend on $\mathbf{\Psi}_t$, and we capture the terms which do not affect the optimization under “remaining terms”. Using Eq. (34), the minimization problem can be rewritten as

$$
\begin{aligned}
\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H}} &:= \arg\min_{\boldsymbol{\Psi}_t} \mathcal{L}_{\mathbb{H}}(\tilde{R}^h_{\mathcal{P},t+1}) \\
&= \arg\min_{\boldsymbol{\Psi}_t} \left\{ \frac{1}{2}\Big(\mathbf{y} + \mathbf{X}\boldsymbol{\Psi}_t\Big)' \mathbf{W} \Big(\mathbf{y} + \mathbf{X}\boldsymbol{\Psi}_t\Big) + \frac{1}{2}(\boldsymbol{\Psi}_t + \mathbf{z_0})' \mathbf{Z}(\boldsymbol{\Psi}_t + \mathbf{z_0}) + \text{“remaining terms”} \right\} \\
&= \arg\min_{\boldsymbol{\Psi}_t} \left\{ \big\|\mathbf{y} - \mathbf{X}(-\boldsymbol{\Psi}_t)\big\|^2_{\mathbf{W}} + \big\|(-\boldsymbol{\Psi}_t) - (-\boldsymbol{\Psi}^*_{t,\text{amb}})\big\|^2_{\mathbf{Z}} \right\},
\end{aligned}
\tag{35}
$$

where $\|\boldsymbol{\Psi}_t\|^2_{\mathbf{D}} = \boldsymbol{\Psi}_t' \mathbf{D}\boldsymbol{\Psi}_t$ stands for the weighted L^2-norm given a positive definite matrix $\mathbf{D}$.

Two technical remarks are due. First, we note that if a constant term was included in the regression, i.e., an unambiguous prospect, the generalized penalty matrix $\mathbf{Z}$ would not have been positive definite. In fact, it would not generate a proper norm, and the shrinkage target would not exist. Hastie, Tibshirani, and Friedman (2009) argue that an intercept term should always be left out of the ridge penalty term. Otherwise, the penalization procedure would be dependent on the origin chosen for the predictor variable. This is precisely the reason why we use demeaned returns and exclude the intercept term from the regression. Second, we stress that the generalized ridge regression representation is a consequence of (a) The linear relationship between the portfolio returns and the underlying currency exposure, and (b) The choice of robust mean–variance preferences. In (a), the linearity is inherited from forward contracts which are used for the hedging of currency risk in our framework. In (b), the ambiguity is measured in a quadratic manner, i.e., see Eq. (25), which opens the door to a generalized ridge regression representation. Hence, the differentiability and smoothness are sustained also in the presence of ambiguity, and a unique closed-form solution exists.

To derive the expression from Eq. (9), we have to solve the generalized ridge regression given in Eq. (35). Using the notation introduced in Section 3.2.1, and by taking a vector derivative with respect to $\boldsymbol{\Psi}_t$, we obtain the first-order optimality condition:

$$
\mathbf{X}' \mathbf{W}\mathbf{X}\boldsymbol{\Psi}_t + \mathbf{X}' \mathbf{W}\mathbf{y} + \mathbf{Z}\boldsymbol{\Psi}_t + \mathbf{Z}\mathbf{z_0} \overset{!}{=} 0. \tag{36}
$$

After some matrix algebra, we obtain the following expression:

$$
\widehat{\boldsymbol{\Psi}}^*_{t,\mathbb{H}} = -\big(\mathbf{X}' \mathbf{W}\mathbf{X} + \mathbf{Z}\big)^{-1}\big(\mathbf{X}' \mathbf{W}\mathbf{y} + \mathbf{Z}\mathbf{z_0}\big), \tag{37}
$$

which is equivalent to Eq. (9). We note that the second-order vector derivative is positive definite, which ensures that the unique global minimum of the loss function is attained.

The above result can be further transformed into the following expression:

$$
\begin{aligned}
\widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H}} &= -\big(\mathbf{X}'\,\mathbf{W}\mathbf{X}+\mathbf{Z}\big)^{-1}\big(\mathbf{X}'\,\mathbf{W}\mathbf{X}(\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y}+ \\
&\quad +\mathbf{Z}(\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y}-\mathbf{Z}(\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y}+\mathbf{Z}\mathbf{z_0}\big) \\
&= -(\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y}+(\mathbf{X}'\,\mathbf{W}\mathbf{X}+\mathbf{Z})^{-1}\mathbf{Z}((\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y}-\mathbf{z_0}) \\
&= \widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H},\mathrm{risk}}+\big(\mathbf{X}'\,\mathbf{W}\mathbf{X}+\mathbf{Z}\big)^{-1}\mathbf{Z}\,\big(\mathbf{\Psi}^{*}_{t,\mathrm{amb}}-\widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H},\mathrm{risk}}\big),
\end{aligned}
\tag{38}
$$

where $\widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H},\mathrm{risk}} = -(\mathbf{X}'\,\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{W}\mathbf{y} = -(\mathbf{X}'\,\mathbf{X})^{-1}\mathbf{X}'\,\mathbf{y}$ is the sample-efficient currency exposure in the absence of ambiguity. The correction term arising from model uncertainty vanishes when $\theta \to 0$, or when $\mathbf{\Psi}^{*}_{t,\mathrm{amb}} = \widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H},\mathrm{risk}}$. In the latter case, the shrinkage target matches the OLS solution, which geometrically corresponds to an overlap of the red and the blue point in Fig. 1. The risk-only sample-efficient solution is then attainable for any θ and the model uncertainty does not affect the OLS-based solution. Furthermore, assuming that the cash is risk-free, government bonds and equities are purely risky, and currencies are ambiguous assets, such as in Section 3.1.4, the ridge target (i.e., the blue point in Fig. 1) becomes fixed in the origin. In the limit $\theta \to \infty$, the set of attainable currency exposures is shrunk to a single point in the origin, implying the optimality of full hedging for the extreme ambiguity-averse investor. This concludes the derivation of Eq. (9). ■

A.6 The Ambiguity-Adjusted Market Price of Currency Risk

Starting from Eq. (3), the expression for the optimal in-sample currency exposure can be written as

$$
\begin{aligned}
\widehat{\mathbf{\Psi}}^{*}_{t,\mathbb{H}} &= -\Big(\lambda\widehat{\mathrm{Var}}_{\mathbb{H}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]+\theta\widehat{\mathrm{Var}}_{\mu}[\,\widehat{\mathrm{E}}_{\mathbb{Q}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]\,]\Big)^{-1} \\
&\quad \cdot\Big(\lambda\widehat{\mathrm{Cov}}_{\mathbb{H}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1},\mathbf{e}_{t+1}-\mathbf{f}_t\,]+\theta\widehat{\mathrm{Cov}}_{\mu}[\,\widehat{\mathrm{E}}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,],\widehat{\mathrm{E}}_{\mathbb{Q}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]\,]\Big) \\
&\quad +\Big(\lambda\widehat{\mathrm{Var}}_{\mathbb{H}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]+\theta\widehat{\mathrm{Var}}_{\mu}[\,\widehat{\mathrm{E}}_{\mathbb{Q}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]\,]\Big)^{-1}\widehat{\mathrm{E}}_{\mathbb{H}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,].
\end{aligned}
\tag{39}
$$

As discussed in Section 3.1, the first term in the final expression represents the ambiguity-adjusted hedging demand $\widehat{\mathbf{\Xi}}^{*}_{t,\mathbb{H}}$. The second term is the ambiguity-adjusted market price of currency risk $\widehat{\mathbf{\Lambda}}^{*}_{t,\mathbb{H}}$, which can be interpreted as a currency "alpha" demand. The ambiguity-adjusted hedging demand and market price of currency risk are defined in Eqs. (4a) and (4b). We demonstrate below that the latter term can be calculated as the solution to a generalized ridge regression. More specifically, we consider the following problem:

$$
\mathbf{x}^{*} = \arg\min_{\mathbf{x}} \Big\{ \big\|\mathbf{A}\mathbf{x}-\mathbf{b}\big\|^{2}_{\mathbf{P}} + \big\|\mathbf{x}\big\|^{2}_{\mathbf{Q}} \Big\}.
$$

This is a generalized ridge regression with the shrinkage target at the origin. Following van Wieringen (2021), it can be solved in a closed form:

$$
\mathbf{x}^{*} = (\mathbf{A}'\mathbf{P}\mathbf{A}+\mathbf{Q})^{-1}\mathbf{A}'\mathbf{P}\mathbf{b}. \tag{40}
$$

Using the notation introduced in Section 3.2.1, and defining $\mathbb{1}$ and $\mathbf{u}$ as the vector of ones and the average return vector of the columns in $\mathbf{X}$, respectively, we consider the following specification:

$$\begin{cases} \mathbf{A} := \mathbf{X}, \\ \mathbf{P} := \frac{\lambda}{T}\mathbf{I}, \\ \mathbf{b} := \frac{1}{T}\mathbb{1}, \\ \mathbf{Q} := \mathbf{Z} - \lambda \mathbf{u}\mathbf{u}', \\ \mathbf{x} := \mathbf{\Lambda}_t. \end{cases} \tag{41}$$

By combining Eqs. (40) and (41), we obtain the expression for the ambiguity-adjusted in-sample market price of currency risk:

$$\begin{aligned} \widehat{\mathbf{\Lambda}}^*_{t,\mathbb{H}} &= \Big(\frac{\lambda}{T}\mathbf{X}'\mathbf{X} - \lambda\mathbf{u}\mathbf{u}' + \mathbf{Z}\Big)^{-1}\frac{1}{T}\mathbf{X}'\mathbb{1} \\ &= \Big(\lambda\widehat{\mathrm{Var}}_{\mathbb{H}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,] + \theta\widehat{\mathrm{Var}}_{\mu}[\,\widehat{\mathrm{E}}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]\Big)^{-1}\widehat{\mathrm{E}}_{\mathbb{H}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]. \end{aligned} \tag{42}$$

Therefore, this component of the optimal in-sample currency exposure can be computed by regressing the vector of ones on the excess currency returns while the shrinkage is performed via a generalized L^2-penalty $\|\cdot\|^2_{\mathbf{Q}}$, where $\mathbf{Q}$ is the regularization matrix defined in Eq. (41). The positive definiteness of this matrix depends on the interaction between the ambiguity and pure alpha given by $\mathbf{u}$. In the absence of ambiguity (i.e., $\theta = 0$), the regularization matrix is $\mathbf{Q} = -\lambda\mathbf{u}\mathbf{u}'$, i.e., it is negative definite and its rank is equal to one. For a positive and sufficiently large coefficient of ambiguity aversion θ, the regularization matrix $\mathbf{Q}$ becomes positive definite.

The regression formulation of the ambiguity-adjusted market price of currency risk represents the optimal trade-off between alpha, risk, and ambiguity. This currency sub-portfolio is then overlaid on top of the hedging demand derived in Eq. (35). ■

Appendix B Additional Empirical Results

B.1 In-Sample Analysis: Supplementary Materials

This section presents some supplementary material that enhances the in-sample analysis from Section 4.2. Figure 5 provides an additional illustration of the results presented in Fig. 3 by plotting the entire bootstrapped distribution of the optimal exposures to CHF and EUR for a USD-based investor whose portfolio comprises Swiss, Eurozone, and US equities and bonds. The underlying assets are mixed in a ratio of three (equities) to one (bonds) and are equally weighted between the three analyzed countries within each asset class. The distribution of the estimator for $\theta = 0$ (represented by the blue bars) corresponds to the risk-only case. The distribution is markedly wide, demonstrating that the optimal currency exposure estimator exhibits large parameter uncertainty. For larger values of θ, the

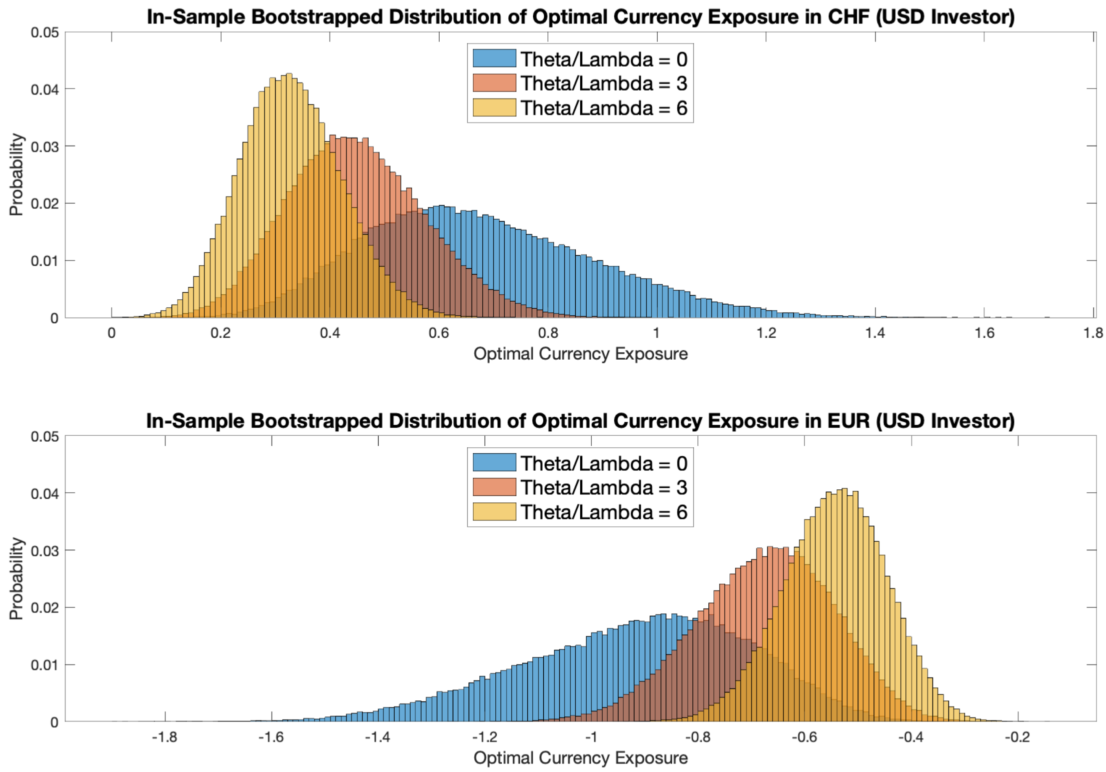

Figure 5 – Bootstrapped distribution of optimal currency exposure in CHF and EUR for a USD-based investor, as a function of the risk and ambiguity aversion parameters λ and θ. We assume that (a) the uncovered interest rate parity condition is satisfied, and (b) the prediction models are independent, and it holds that $\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] = \frac{1}{T^2}\mathbf{I}$. The sample spans the period from January 1999 until December 2019.

distribution shifts toward the ridge target—equal to zero in our example—and becomes significantly narrower. Note that in the limiting case when $\theta \to \infty$, the distribution of the estimated optimal currency exposure converges to the Dirac delta function centered at the origin—implying the limiting optimality of full hedging.

In Fig. 6 we present an efficient surface—a generalization of the standard efficient frontier to a three-dimensional setting. To generate the efficient surface, we maximize the expected portfolio return on a pre-defined two-dimensional grid of the portfolio variance arising from risk and the portfolio variance arising from ambiguity. Mathematically, this optimization problem can be cast as a linear program with quadratic constraints. We consider a USD-based investor who holds an equally-weighted portfolio consisting of broad equity and bond indices representing the seven developed markets described in Section 4. The efficient surface results from an optimization over currencies, while the equity and bond portfolio weights are kept fixed. For practical reasons, we impose an upper bound on the leverage, which is equal to 100% of the currency position implied by the portfolio design, without any initial currency overlay. Note that without any constraint on currency positions, the problem becomes unbounded as infinite leverage is possible.

Observe that, for a given level of ambiguity-induced portfolio volatility, one obtains a standard, two-

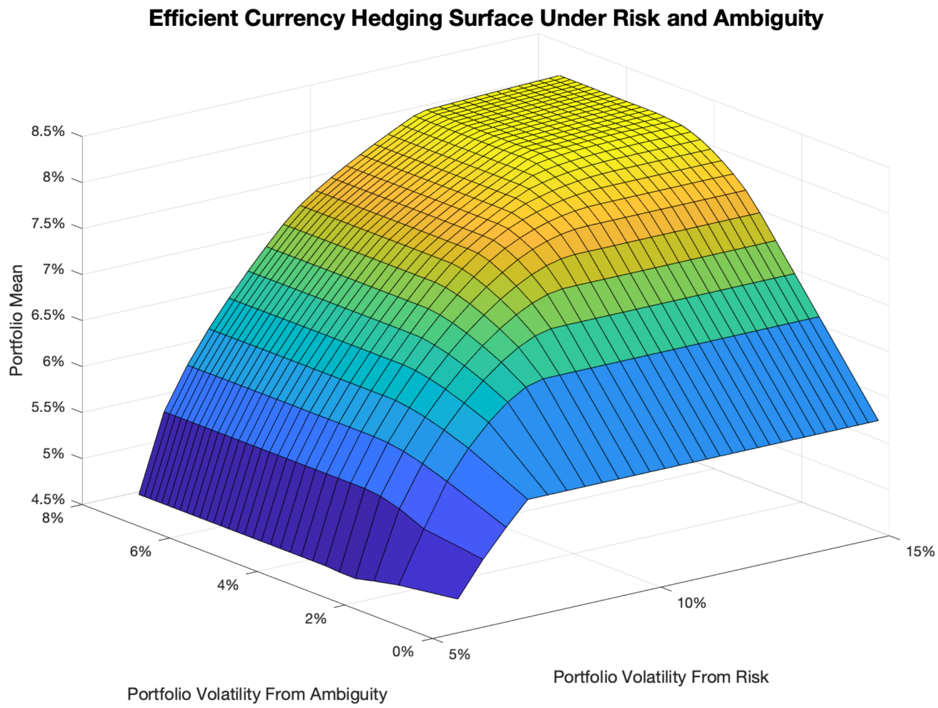

Figure 6 – The efficient currency hedging surface for a USD-based investor who holds a global equally-weighted portfolio that consists of developed market equities and bonds. The efficient surface illustrates the optimal in-sample trade-off between the portfolio return and the two components of portfolio volatility driven by risk and ambiguity, respectively. We assume that (a) the uncovered interest rate parity condition is satisfied, and (b) the prediction models are independent, and it holds that $\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] = \frac{1}{T^2}\mathbf{I}$. The sample spans the period from January 1999 until December 2019.

dimensional efficient frontier (i.e., the expected return versus the volatility arising from risk). Moreover, the point on the efficient surface which corresponds to the minimum variance and the maximum ambiguity converges to the global minimum variance case in a standard mean–variance framework, which does not include ambiguity. The convergence is achieved only if the leverage constraint is sufficiently relaxed or removed, otherwise, the standard minimum variance portfolio might not be feasible.

Table 6, in analogy to Panel A of Table III in Campbell, Serfaty-De Medeiros, and Viceira (2010), reports optimal currency exposures for an investor that is fully invested in a single-country equity portfolio and considers a single foreign currency for hedging. A multivariate version of this problem (where investment in all available currencies is possible) is presented in Table 3. We assume a quarterly hedging horizon and employ a non-parametric bootstrap to estimate the standard deviations of the optimal exposures for both risk-only (RO) and ambiguity-adjusted (AA) investors.

The results show that single-country stock investors considering a single foreign currency almost always choose positive exposures to the Swiss franc, Japanese yen, and US dollar (i.e., these currencies are perceived as safe-haven currencies), mixed exposures to the euro, and predominantly negative exposures to the Australian dollar, Canadian dollar, and British pound. These results are consistent between both RO and AA investors. The difference between both investors again arises when one looks at the standard deviations of the optimal currency exposures. These standard deviations are lower in the case of the AA investor, which is an effect of the ridge regression. However, this decrease is in a

Optimal Currency Exposure:
Single-Country and Single-Currency Portfolios

Stock Market	Model	AUD	CAD	CHF	EUR	GBP	JPY	USD
Australia	RO		0.53	0.70	0.62	0.38	0.55	0.57
			(0.31)	(0.20)	(0.21)	(0.21)	(0.09)	(0.17)
	AA		0.44	0.63	0.55	0.34	0.53	0.53
			(0.27)	(0.18)	(0.19)	(0.19)	(0.09)	(0.17)
Canada	RO	-0.75		0.65	0.55	0.32	0.64	1.06
		(0.35)		(0.20)	(0.17)	(0.18)	(0.21)	(0.33)
	AA	-0.63		0.57	0.48	0.27	0.60	0.89
		(0.30)		(0.17)	(0.14)	(0.15)	(0.20)	(0.30)
Switzerland	RO	-0.78	-0.62		-0.72	-0.70	0.24	-0.10
		(0.15)	(0.18)		(0.38)	(0.24)	(0.22)	(0.23)
	AA	-0.70	-0.54		-0.49	-0.58	0.22	-0.08
		(0.14)	(0.15)		(0.23)	(0.20)	(0.20)	(0.20)
Eurozone	RO	-0.60	0.16	1.12		-0.02	0.82	1.09
		(0.33)	(0.24)	(0.62)		(0.31)	(0.22)	(0.27)
	AA	-0.53	0.14	0.80		-0.01	0.77	0.97
		(0.29)	(0.20)	(0.36)		(0.25)	(0.21)	(0.24)
UK	RO	-0.70	-0.44	0.34	0.24		0.37	0.41
		(0.17)	(0.16)	(0.25)	(0.24)		(0.20)	(0.31)
	AA	-0.62	-0.37	0.29	0.19		0.35	0.35
		(0.16)	(0.14)	(0.21)	(0.19)		(0.19)	(0.27)
Japan	RO	-0.74	-0.94	-0.46	-0.43	-0.78		-0.89
		(0.17)	(0.21)	(0.29)	(0.27)	(0.21)		(0.28)
	AA	-0.70	-0.86	-0.42	-0.40	-0.72		-0.79
		(0.16)	(0.19)	(0.27)	(0.25)	(0.20)		(0.25)
USA	RO	-0.72	-1.05	-0.09	-0.24	-0.64	0.35	
		(0.17)	(0.25)	(0.24)	(0.25)	(0.31)	(0.24)	
	AA	-0.66	-0.88	-0.08	-0.22	-0.54	0.31	
		(0.16)	(0.23)	(0.21)	(0.22)	(0.27)	(0.21)	

Table 6 – This table presents optimal currency exposures and their standard deviations for a domestic equity investor who chooses positions from a single available foreign currency to optimize her portfolio. The rows of the table indicate the base country (which corresponds also to the base currency) and the columns represent the optimal positions in (single) currencies for both risk-only (RO) and ambiguity-adjusted (AA) investors. The currency exposures are computed by Table 3, where $\theta = 0$ by construction (i.e., ordinary least squares regression) for the RO investor and $\theta/\lambda = 2$ is assumed for the AA investor (i.e., ridge regression). We assume a quarterly hedging horizon and employ a non-parametric bootstrap to estimate the standard deviations of the optimal exposures for both RO and AA investors.

single foreign currency case smaller compared to the multivariate case presented in Table 3. This shows that the ridge regression provides a larger influence on the distribution of optimal currency exposures in a multivariate compared to a univariate setting. The largest decrease in the standard deviations from the RO to the AA investor is again attained for the Swiss and Eurozone investors due to the high correlation between their markets and currencies.

B.2 Out-of-Sample Backtest with Speculative Demand

This section extends the out-of-sample (OOS) backtest from Section 4.3 by including also the speculative currency demand. We study global portfolios of equities and bonds for various choices of base currencies. The underlying assets are mixed in a ratio of three (equities) to one (bonds) and are equally weighted between the three analyzed countries within each asset class. We study three different model-based currency overlay strategies: the risk-only minimum variance (RO), which is equivalent to the risk-only case from Section 4.3, the risk-only mean–variance (RO-Spec), which is a risk-only investor with added speculative currency demand, and the ambiguity-adjusted mean–variance (AA-Spec), which is the ambiguity-adjusted investor with added speculative currency demand. The OOS performance of these portfolios is compared to three constant-hedging benchmarks: zero, half, and full hedging. Additionally, we include the leverage constraints on all models and allow optimal currency positions to lie between -1 and 1 (which prevents extreme leverage and still gives models enough space to possibly take advantage of the estimated speculative demands).

We worked with the so-called naive ambiguity in the main part of the paper. Here, in order to account for the (ambiguity-adjusted) speculative currency demand, we employ several popular exchange rate forecasting models to estimate the expected currency return $\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]$ and the ambiguity covariance matrix $\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]$.[15] We follow the work of Cheung, Chinn, Pascual, and Zhang (2019) who examine the predictive performance of various exchange rate models. More specifically, we consider seven forecasting models in total: (a) Historical average, (b) Uncovered interest rate parity, (c) Relative purchasing power parity, (d) Sticky price monetary model, (e) Behavioral equilibrium exchange rate (BEER) model, (f) Sticky price monetary model augmented by risk and liquidity factors, and (g) Yield curve slope model. Further information about the models and the macroeconomic data used for the exchange rate forecasting is provided in Appendix B.3.

We label the models by their respective probability measures $\mathbb{Q}_i$, where $i = 1, 2, \ldots, n$. Specifically, in our application $n = 7$. The set of all considered probability measures is denoted by $\mathcal{Q}$. In the first step, our goal is to estimate $\mathrm{E}_{\mathbb{Q}_i}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]$ for each model $\mathbb{Q}_i$. To compute $\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]$ and $\mathrm{Var}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,]$, we need to specify a prior μ, which governs the weighting of the models in $\mathcal{Q}$. A common approach is to use equal weighting with $\mu_i = 1/n$, where μ_i is the weight corresponding to the model $\mathbb{Q}_i$. However, the model performance can be potentially improved by calibrating the model weights μ to the data. Forecast combinations have been advocated to outperform individual forecasting models from both theoretical and empirical perspectives, e.g., see Timmermann (2006). We follow Conflitti, De Mol, and Giannone (2015) and compute the weights μ by minimizing the mean square forecasting error for the ambiguity-adjusted mean–variance model under the constraints of $\sum_{i=1}^{n} \mu_i = 1$ and $\mu_i \geq 0$ (for $i = 1, 2, \ldots, n$).

[15] We note that the assumption $\mathrm{Cov}_\mu[\,\mathrm{E}_{\mathbb{Q}}[\,\tilde{R}^{fh}_{\mathcal{P},t+1}\,], \mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1} - \mathbf{f}_t\,]\,] = \mathbf{0}$ remains intact. The economic interpretation of this condition is that fully hedged positions are treated as purely risky (i.e., non-ambiguous). Furthermore, this choice fixes the shrinkage target at the origin.

The set up of the OOS backtest is the same as in Section 4.3. Namely, currency hedging is implemented using forward contracts which are rolled over quarterly and the asset positions are also rebalanced quarterly. When a currency forward contract expires, the resulting P&L is reinvested in the assets such that the equal portfolio weighting is preserved over the backtest. A new hedge is formed immediately thereafter and the process is repeated over time. The covariance structure is computed on a rolling window of two years of daily historical data, whereas the predictive models that employ macroeconomic data (which is available on a monthly or quarterly frequency) take six years of historical data for calibration. Meaning, the first six years of the data set described in Section 4.1 are used for calibration and the rest of the data set corresponds to the out-of-sample (rolling-window-based) empirical exercise. All results in this section are presented net of hedging (i.e., transaction) costs, which are assumed to be ten basis points relative to the notional of a currency forward contract. Moreover, we assume that the asset positions are rebalanced without transaction costs in order to isolate and study solely the effects of different currency overlay strategies on the portfolio performance. Note that our findings are robust to alternative assumptions on the level of transaction costs, hedge maturity (i.e., rebalancing frequency), and portfolio composition. The corresponding results are available on request.

Table 7 reports the annualized average returns, volatilities, and risk-adjusted performance metrics (i.e., Sharpe and Sortino ratios) of portfolio daily returns, as well as the average hedging turnover. The statistics are presented for all considered base currencies and across different hedging strategies described above. Our findings indicate that the out-of-sample benefit of currency hedging strongly depends on the investor's base currency and the model.

Similarly to Table 5, all currency overlay strategies outperform constant-hedging benchmarks in terms of risk reduction as well as risk-adjusted performance. As expected, the RO model achieves the largest volatility reduction. However, also the AA-Spec strategy manages to reduce the risk of portfolio returns well. The robust hypothesis test for the difference in variances from Ledoit and Wolf (2011) shows that RO and AA-Spec strategies produce a highly statistically significant reduction in variance compared to both zero and full hedging (except for the AUD-based AA-Spec investor). On the other hand, the RO-Spec model does not produce any statistically significant reduction in variance compared to both zero and full hedging.

Next, we investigate the risk-adjusted portfolio performance as measured by the Sharpe ratio. Again, the model-based strategies tend to outperform their constant hedging benchmarks. The largest relative increase in the Sharpe ratio is observed for Swiss and US investors. The robust hypothesis test for the difference in Sharpe ratios from Ledoit and Wolf (2008) shows that these differences are not statistically significant. Note that statistically significantly distinguishing Sharpe ratios is difficult since Sharpe ratios sensitively depend on the realized average portfolio returns that are noisy estimates of the true mean portfolio returns.

Finally, we study the turnover of our currency overlay strategies as defined in Eq. (13). The turnover

Out-of-Sample Backtest of Currency Overlay Strategies

	Constant Hedge			**Model Hedge**		
Panel A: AUD	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	6.75	7.38	8.01	8.54	9.97	9.86
Volatility (%)	8.67	7.81	9.37	6.89	8.46	7.86
Sharpe Ratio	0.78	0.94	0.85	1.24	1.18	1.26
Sortino Ratio	1.12	1.32	1.17	1.77	1.77	1.95
Turnover	0	0.43	0.86	1.48	1.58	1.47
Panel B: CAD	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	6.69	6.43	6.17	6.42	7.95	6.95
Volatility (%)	9.42	8.75	9.35	6.90	8.47	7.27
Sharpe Ratio	0.71	0.74	0.66	0.93	0.94	0.96
Sortino Ratio	1.00	1.01	0.90	1.31	1.39	1.37
Turnover	0	0.43	0.86	1.58	1.59	1.51
Panel C: CHF	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	12.87	10.70	9.34	6.89	8.37	7.36
Volatility (%)	11.01	8.68	7.22	5.91	5.78	5.76
Sharpe Ratio	0.41	0.48	0.53	0.70	0.75	0.74
Sortino Ratio	0.56	0.66	0.73	0.98	1.09	1.06
Turnover	0	0.43	0.86	1.64	1.67	1.59
Panel D: EUR	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	7.21	6.39	5.56	5.73	7.57	6.63
Volatility (%)	10.10	9.41	9.35	6.89	8.94	7.31
Sharpe Ratio	0.71	0.68	0.59	0.83	0.85	0.87
Sortino Ratio	0.99	0.93	0.81	1.17	1.30	1.26
Turnover	0	0.43	0.86	1.38	1.45	1.34
Panel E: GBP	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	8.96	7.66	6.36	6.48	8.09	7.27
Volatility (%)	10.39	9.33	9.36	6.89	8.49	7.41
Sharpe Ratio	0.86	0.82	0.68	0.94	0.95	0.98
Sortino Ratio	1.24	1.15	0.93	1.33	1.42	1.43
Turnover	0	0.43	0.86	1.68	1.70	1.63
Panel F: JPY	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	7.18	6.18	5.18	4.71	6.30	5.52
Volatility (%)	16.59	12.61	9.33	6.89	8.47	7.54
Sharpe Ratio	0.43	0.49	0.56	0.68	0.74	0.73
Sortino Ratio	0.59	0.67	0.76	0.95	1.09	1.06
Turnover	0	0.43	0.86	1.71	1.67	1.63
Panel G: USD	Zero	Half	Full	RO	RO-Spec	AA-Spec
Mean (%)	6.75	6.67	6.59	6.41	7.94	7.02
Volatility (%)	12.82	10.78	9.34	6.89	8.47	7.25
Sharpe Ratio	0.53	0.62	0.71	0.93	0.94	0.97
Sortino Ratio	0.74	0.86	0.97	1.31	1.39	1.39
Turnover	0	0.43	0.86	1.66	1.68	1.62

Table 7 – We report the out-of-sample (annualized) average return, volatility, Sharpe ratio, Sortino ratio, and turnover for six currency overlay strategies (their definitions are provided at the beginning of this section) and across seven base currencies. We assume $\lambda = 3$ and $\theta = 6$ for all the strategies, while $\theta 0$ only affects the AA investor. The underlying global portfolio consists of equities and bonds mixed in proportion three (equities) to one (bonds). Within each asset class, we assume an equally-weighted sub-portfolio that consists of the corresponding developed market indices.

of the constant hedging strategies remains unchanged across all base currencies. By construction, the currency overlay strategies are more reactive to changing market conditions than the constant-hedging strategies, hence generating a higher average turnover. RO-Spec achieves a higher turnover compared to the RO model, which is expected since the speculative demand, which is generally less stable over time, is added to the pure hedging demand RO. It is interesting to note that even though the speculative demand is also added to the AA-Spec model, its turnover is not elevated given the shrinkage induced by ambiguity aversion. Again, the shrinkage stabilizes the estimation of optimal currency exposures over time.

B.3 A Compendium of the Implemented Exchange Rate Forecasting Models

In Appendix B.2, we utilize different exchange rate forecasting models to estimate the expected excess currency return $\mathrm{E}_{\bar{\mathbb{Q}}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]$ and the ambiguity covariance matrix $\mathrm{Var}_{\mu}[\,\mathrm{E}_{\mathbb{Q}}[\,\mathbf{e}_{t+1}-\mathbf{f}_t\,]\,]$. The choice of the models is based on the work of Cheung, Chinn, Pascual, and Zhang (2019). Here, we provide the specifications of the implemented models. All macroeconomic variables described in this section, and on top of those listed in Section 4.1, are sourced from Bloomberg.

(1) Historical average: A sample average is calculated for a given set of historical spot and forward exchange rates.

(2) Uncovered interest rate parity (UIRP):

$$S_{c,t+1} = S_{c,t} + \widetilde{i}_t, \tag{43}$$

where $S_{c,t}$ is the exchange rate of currency c with respect to the home/base currency at time t, i_t is the interest rate at time t with maturity $t+1$, and tilde denotes the inter-country difference. The tilde notation applies to all listed models.

(3) Relative purchasing power parity:

$$S_{c,t+1} = \beta_0 + \beta_1 \widetilde{p}_t + \epsilon_{t+1}, \tag{44}$$

where p_t is the logarithm of the consumer price index (CPI) at time t, and ϵ_{t+1} is the error term.

(4) Sticky price monetary model:

$$S_{c,t+1} = \beta_0 + \beta_1 \widetilde{m}_t + \beta_2 \widetilde{y}_t + \beta_3 \widetilde{i}_t + \beta_4 \widetilde{\pi}_t + \varepsilon_{t+1}, \tag{45}$$

where m_t is the logarithm of the money supply, and y_t is the logarithm of the real (i.e., inflation-adjusted) gross domestic product (GDP). The variables i_t and π_t represent the interest and inflation rate, respectively, and ε_{t+1} is the error term.

(5) Behavioral equilibrium exchange rate (BEER) model:

$$S_{c,t+1} = \beta_0 + \beta_1 \widetilde{p}_t + \beta_2 \widetilde{\omega}_t + \beta_3 \widetilde{r}_t + \beta_4 \widetilde{gdebt}_t + \beta_5 \widetilde{tot}_t + \beta_6 \widetilde{nfa}_t + \upsilon_{t+1}, \tag{46}$$

where p_t is the logarithm of the CPI, ω_t is the relative price of non-tradables, r_t is the real interest rate, $gdebt_t$ is the government-debt-to-GDP ratio, tot_t is the logarithm of the terms of trade, nfa_t is the net foreign assets, and υ_{t+1} is the error term.

(6) Sticky price monetary model augmented by risk and liquidity factors:

$$S_{c,t+1} = \beta_0 + \beta_1 \widetilde{m}_t + \beta_2 \widetilde{y}_t + \beta_3 \widetilde{i}_t + \beta_4 \widetilde{\pi}_t + \beta_5 VIX_t + \beta_6 TED_t + \xi_{t+1}, \tag{47}$$

which can be understood as the Sticky price monetary model augmented for the VIX and the three-month Treasury-Libor (TED) spread.

(7) Yield curve slope:

$$S_{c,t+1} - S_{c,t} = \beta_0 + \beta_1 \widetilde{i}_t + \beta_2 slope_t + \eta_{t+1}, \tag{48}$$

where the inter-country difference in the level of the three-month interest rate is combined with the difference in the slope (ten-year minus three-month yields).

The first specification is based on sample estimates. The second model does not require any estimation to generate a prediction. All remaining models are based on OLS regressions.

The financial market variables (i.e., the interest rates, VIX, TED spread, spot and forward exchange rates) are obtained from Refinitiv Datastream and Bloomberg. The macroeconomic variables listed above are sourced from the International Monetary Fund (IMF), the Organization for Economic Co-operation and Development (OECD), Bloomberg, and the websites of the relevant central banks, i.e., the Reserve Bank of Australia (RBA), the Bank of Canada (BoC), the Swiss National Bank (SNB), the European Central Bank (ECB), the Bank of England (BoE), the Bank of Japan (BoJ), and the Federal Reserves (Fed). ♣