Journal of International Money and Finance 139 (2023) 102949

Contents lists available at ScienceDirect

Journal of International Money and Finance

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

Sparse and stable international portfolio optimization and currency risk management [☆]

Raphael Burkhardt [a], Urban Ulrych [a,b,c,*]

[a] Department of Banking and Finance, University of Zurich, Switzerland
[b] College of Management of Technology, École Polytechnique Fédérale de Lausanne, Switzerland
[c] Swiss Finance Institute, Switzerland

ARTICLE INFO

JEL classification:
C61
F31
G11
G15

Keywords:
International asset allocation
Currency risk management
Currency overlay
Shrinkage
Sparsity
Stability
Regularization
Mean-variance optimization

ABSTRACT

This paper introduces a sparse and stable optimization approach for multi-currency asset allocation, aiming to improve portfolio performance and currency risk management. We demonstrate that the widespread industry practice of employing currency overlay strategies is suboptimal. In contrast, our proposed regularized joint optimization approach, which integrates assets and currencies, consistently outperforms currency overlay strategies as well as equally-weighted and non-regularized global portfolio benchmarks net of transaction costs. On average, the joint optimization approaches achieve 23.3% higher out-of-sample Sharpe ratios compared to their currency overlay counterparts. By addressing parameter uncertainty and inducing sparsity and stability, our method enhances the mean-variance framework, resulting in improved out-of-sample portfolio performance. These findings challenge the prevailing practice of employing currency overlay strategies and highlight the potential for additional gains in risk-adjusted returns through the joint optimization of assets and currencies.

1. Introduction

International diversification is a common practice among asset managers aimed at improving the portfolio risk-return profile. This improvement can be achieved by the outperformance of foreign markets or by mitigating the (otherwise domestic) systematic market risk. However, global portfolios are prone to a new aspect of risk stemming from fluctuating currency exchange rates. Consequently, managing currency exposure lies at the core of the risk management practice of international investors. In practice, a standard approach for hedging currency risk is performed via a currency overlay, see Kim and Chance (2018). In a currency overlay, an agent first optimizes the asset weights and then subsequently determines the currency exposures by taking positions in currency forward contracts. In this paper, we investigate the magnitude of the sub-optimality associated with currency overlay strategies and propose a novel framework for multi-currency asset allocation that surpasses its limitations. Specifically, we examine the benefits

☆ The authors are grateful to Ming Deng, Jérôme Detemple, Walter Farkas, Damir Filipović, Iftekhar Hasan, Sean Hofland, Valentyn Khmarskyi, Isabella Kooij, Markus Leippold, Patrick Lucescu, Ana Mão de Ferro, Nikola Vasiljević, Rudi Zagst, and the participants of the SFI Research Days 2021, the 7th International Young Finance Scholars Conference (2021), the 14th Edition of the Annual Meeting of The Risk, Banking and Finance Society - International Risk Management Conference 2021, the 10th International Hybrid Conference on Mathematical and Statistical Methods for Actuarial Sciences and Finance (2022), and the Belgian Financial Research Forum 2023 for helpful comments and suggestions.

* Corresponding author at: College of Management of Technology, École Polytechnique Fédérale de Lausanne, Switzerland.
E-mail addresses: raphael.burkhardt@hotmail.de (R. Burkhardt), urban.ulrych@epfl.ch (U. Ulrych).

https://doi.org/10.1016/j.jimonfin.2023.102949

Available online 7 September 2023

of a joint optimization approach to international asset allocation where asset weights and currency exposures are computed in a single (i.e., joint) optimization. While the potential of a joint optimization approach has been explored before, this paper proposes a novel framework that integrates optimization techniques inducing sparsity and stability of the asset and currency weights. Such an approach provides a generalization to the standard Markowitz setting and aims to improve the out-of-sample portfolio performance.

In the context of decentralized investment management, as highlighted in Van Binsbergen et al. (2008), asset allocation decisions often occur in multiple stages. For instance, a Chief Investment Officer (CIO) may first allocate capital to different asset classes, each managed by a specialized asset manager, before determining how to hedge the currency risk of the overall portfolio. This two-stage process can potentially introduce misalignments of incentives and lead to utility costs for the CIO. Similarly, in currency risk management, capital allocation decisions to various assets, such as equities and bonds, are frequently made independently from the subsequent decision of determining the optimal currency hedge for the entire portfolio. Consequently, unexploited expected profits might exist due to the lack of a joint optimization approach. By incorporating a joint optimization framework, which considers both asset allocation and currency risk management simultaneously, this paper seeks to address such potential inefficiencies and enhance the overall risk-adjusted performance of internationally diversified portfolios. The joint optimization approach has the potential to improve portfolio performance in comparison to the traditional separate optimization routine (i.e., a currency overlay) as it takes into account the dependencies between single assets and currencies when determining the optimal portfolio, see Jorion (1994).

This paper sheds light on the inefficiency of currency overlay strategies and highlights the potential for substantial improvement in portfolio performance through a joint optimization approach. The improvement of portfolio performance crucially depends on the estimation of the optimization parameters. To deal with the problems related to parameter uncertainty in the mean-variance framework, we enhance the multi-currency asset allocation problem with regularization methods arising from the area of statistical learning. The main objective of this paper is to study a general approach that produces stable and robust multi-currency portfolios. By considering both asset allocation decisions and currency risk management within a unified framework, we contribute to the existing literature by presenting a novel regularization-based approach for international portfolio optimization and testing its out-of-sample performance net of transaction costs. Our analysis shows that the widespread industry practice of utilizing currency overlay strategies is suboptimal and can be improved by the proposed regularized joint optimization over assets and currencies. We show that, on average, the joint optimization rules achieve 23.3% higher out-of-sample Sharpe ratios compared to their separate optimization counterparts, indicating the robustness and effectiveness of the joint optimization approach in enhancing risk-adjusted returns.

In the methodological part of the paper, we introduce a model that describes a return of an international portfolio that includes domestic and foreign assets together with the currency forward contracts employed to hedge the portfolio currency risk. Building on this formulation, we characterize the separate and joint optimization problems that investors solve when investing internationally. Then, we introduce different portfolio rules that represent extensions of the classical mean-variance framework directed at potentially improving the out-of-sample portfolio performance, see DeMiguel et al. (2009a). These optimization procedures aim at reducing the inherent parameter uncertainty problem by either constraining the norms of the asset- and currency-weight vectors or by shrinking the optimization parameters, hence inducing sparsity and stability (i.e., robust portfolios with potentially only a few active positions). The computational solution to the sparse and stable (SAS) international asset and currency allocation problem is provided in terms of quadratic programming and cross-validation. The introduction of the SAS multi-currency asset allocation framework presents the methodological contribution of this paper.

The empirical analysis compares the joint and separate versions of the international portfolio optimization problem. Leveraging a data sample comprising stocks, spot, and forward currency exchange rates for ten developed economies from January 1999 to December 2019, we thoroughly investigate the out-of-sample performance of the proposed trading strategies. Our key empirical finding highlights the substantial superiority of the SAS joint optimization approach over equivalent strategies based on separate optimization, equally-weighted portfolios, and non-regularized benchmarks, net of transaction costs. Notably, the jointly optimized SAS portfolio attains a Sharpe ratio that is double that of the equally-weighted portfolio, showcasing the significant potential for improved risk-adjusted returns with the joint optimization framework. This empirical contribution unveils the suboptimal nature of the widespread approach that involves constructing an asset portfolio first and then employing currency overlay managers for subsequent currency risk management. In contrast, our results underscore the efficacy of the joint optimization of assets and currencies, driven by the proposed regularization approaches. These regularization techniques address the problem of parameter uncertainty by inducing a bias-variance trade-off, leading to enhanced out-of-sample portfolio performance and superior risk-adjusted returns.

The proposed joint optimization approach not only addresses potential inefficiencies in the current two-stage process of capital allocation and currency risk management but also has broader implications for the behavior of returns and portfolio performance. By simultaneously considering asset allocation decisions and currency hedging within a unified framework, the joint optimization approach aims to capitalize on the unexploited expected profits that may exist under the traditional separate optimization strategies. This integration of asset allocation and currency risk management can lead to more coherent portfolio decisions, reducing misaligned incentives and utility costs. As a result, the SAS joint optimization creates a compelling case for further exploration and adoption of this approach by financial practitioners seeking to enhance the risk-adjusted performance of their portfolios. Moreover, the widespread adoption of the joint optimization strategy could have broader implications for the overall market dynamics. Given the potential for more efficient risk management and enhanced returns by adopting the joint optimization strategy, market participants may witness a shift in portfolio allocation strategies, leading to reduced currency market inefficiencies. All in all, our findings highlight the significant potential benefits of adopting a SAS joint optimization framework for managing internationally diversified portfolios and underscore the importance of considering currency risk in the broader context of asset allocation decisions.

The rest of the paper is organized as follows. Section 2 reviews the existing literature. Section 3 introduces a methodological framework for multi-currency asset allocation and extends this framework by regularization techniques inducing sparsity and stability. Section 4 investigates the empirical performance of the proposed model. Section 5 concludes. Supplementary results are presented in the Appendix, while further derivations and additional figures can be found in the Online Appendix.

2. Literature review

International diversification is a standard practice of institutional investors such as mutual funds, pension funds, insurers, banks, corporate entities, and other financial intermediaries. While investing internationally offers an opportunity of improving the risk-return profile of a portfolio, it also exposes an investor to additional risk stemming from the currency exchange rates. The asset allocation and currency hedging decisions are often considered separately. There are numerous approaches in the literature exploring the benefits of hedging the implicit currency exposure of global portfolios.

The initial studies focus on investigating the potential benefits of fully hedging the implicit currency exposure. Solnik (1974) shows that when assets and currencies are uncorrelated and there is no speculative demand for currencies, it is optimal for an investor to fully hedge the implicit currency exposure. Perold and Schulman (1988) further investigate the benefits of a fully hedged portfolio empirically and find that US investors can reduce the risk of an internationally diversified portfolio of bonds and equities by fully hedging the implicit foreign currency exposure.

Moving away from the full hedging rule, another strand of the literature is positioned around investigating the so-called currency overlay strategies. The overlay strategies aim to reduce the implicit currency risk of a global portfolio by creating a currency portfolio that is laid over the asset portfolio, thereby improving the risk-return profile of the portfolio. A relevant characteristic of a currency overlay is that the asset allocation is separated from the currency allocation, which is determined subsequently after the asset weights have already been fixed. Glen and Jorion (1993) add forward contracts to an internationally diversified portfolio and show that their inclusion results in a statistically significant increase in the portfolio performance. Campbell et al. (2010) find that some currencies move against world equity markets and thus have, despite their low average returns, attractive properties for risk-minimizing international investors. Their model takes into account the correlations between currencies and fully hedged asset returns.

Dynamic currency hedging strategies have been shown to outperform the static approaches by providing significant improvements in portfolio performance.[1] On a more recent note, Boudoukh et al. (2019) derive a decomposition of a currency overlay to a risk-minimizing hedge part and an alpha-generating part. Ulrych and Vasiljević (2020) extend this approach by deriving closed-form expressions for optimal currency allocation under risk and ambiguity aversion. The authors show that ambiguity can be one of the drivers of the insufficient currency diversification puzzle known as the home currency bias. Opie and Riddiough (2020) present a method that exploits the time-series predictability of currency returns by exploiting a forecastable component in global factor returns. Their currency overlay strategy outperforms alternative dynamic hedging benchmarks. Polak and Ulrych (2021) propose a non-Gaussian approach that relies on the extended filtered historical simulation and numerical optimization of a currency overlay with respect to an arbitrary risk measure.

All of the aforementioned papers make substantial contributions to the literature on managing currency risks in international portfolios. However, they share the assumption that the portfolio is first constructed without considering foreign currency exposures, and the optimal hedging decision is subsequently made. Consequently, these overlay strategies are intrinsically suboptimal as they disregard the interactions between individual assets and exchange rates. In the context of decentralized investment settings, Van Binsbergen et al. (2008) analyze the implications of employing multiple asset managers to implement investment strategies in separate asset classes and shed light on the potential for significant utility costs for the chief investment officer. This concept directly connects to the suboptimality of currency overlay strategies, which is a key focus of our investigation in this paper.

The investigation into the benefits of the joint optimization approach compared to employing specialized currency overlay managers starts with Jorion (1994). Jorion finds that while overlay management can add value to global equity and bond portfolios, it cannot improve portfolio performance as much as a joint approach can. Following his work, Suzuki et al. (1997) propose an efficient algorithm that obtains the globally optimal solution by employing both a mean-absolute deviation and a mean-variance model. Using historical data, they show that this approach is more efficient than a traditional two-step (i.e., separate) optimization. In a subsequent paper, Konno and Li (2000) improve the mean-absolute deviation model in international markets by using a beta pricing method for the expected asset returns and show that this method improves the previous results and leads to an improved and more stable performance. An integrated approach of multi-currency asset allocation, where simulation is performed via the principal component analysis and a subsequent optimization is executed with respect to the expected shortfall risk measure, is presented in Topaloglou et al. (2002). On a similar note, Mhiri and Prigent (2010) extend the mean-variance optimization paradigm in international portfolio optimization to higher moments by adding skewness and kurtosis to the utility function. More recently, Chatsanga and Parkes (2017) include the cost of carry, transaction costs, and margin requirements of forward contracts used for hedging directly in the portfolio return calculation. They show that the inclusion of such costs significantly changes the optimal decisions and improves portfolio performance. Huang and Wang (2021) extend the risk-based (i.e., mean-variance) approach of international portfolio construction to uncertainty. They propose an ambiguity-adjusted mean-chance model for international portfolio selection and derive closed-form solutions.

This paper builds on the existing work on optimal asset and currency allocation in international markets by combining the idea of the joint optimization approach with more recent advances in the portfolio optimization literature. The main objective of this work

[1] See, among others, De Roon et al. (2003), Brown et al. (2012), Caporin et al. (2014), and Cho et al. (2020).

is to develop an approach that produces stable and robust portfolios consisting of multi-currency assets. While the existing literature already shows that the joint optimization approach has the potential to produce portfolios with improved risk-return properties, these results have been principally established in-sample as the same period is used for the optimization and the measurement of the portfolio performance. In-sample optimality implies that the joint approach can only improve portfolio performance. However, the investors are, in practice, concerned with the out-of-sample portfolio performance that can be worsened if in-sample over-fitting takes place.

The classical mean-variance optimization, as introduced by Markowitz (1952), is guaranteed to find the best portfolio in terms of the risk-return trade-off when the true asset means and covariances are known. However, this is practically impossible. Mean-variance portfolios often produce very unstable weights that substantially change for small adjustments in the input parameters. Moreover, the optimal portfolio weights are prone to attaining large positive/negative values. Portfolios constructed via the use of the sample mean and the sample covariance matrix generally perform poorly out-of-sample, see e.g., Michaud (1989) and Jorion (1986).

There exist multiple approaches that extend the Markowitz framework by tackling the inherent parameter uncertainty of the optimization problem. One approach widely used in the literature is the shrinkage of the input parameters. Jorion (1986), for example, uses a Bayes-Stein estimator for the mean to reduce parameter uncertainty. Ledoit and Wolf (2003, 2004) propose shrinking the sample covariance matrix towards a shrinkage target with a lower variance. To deal with the large negative weights of the optimal mean-variance portfolio, as well as potential regulatory requirements, no short-sale constraints are often imposed on the portfolio weights. Jagannathan and Ma (2003) show that imposing lower and upper constraints on the portfolio weights can also be interpreted as shrinkage of the covariance matrix. These approaches have been shown to improve the out-of-sample performance of the optimal mean-variance portfolios.

DeMiguel et al. (2009b) show that many portfolio strategies, including those mentioned above, fail to consistently outperform the naive diversification of the $1/N$ portfolio that assigns equal weights to all assets in the portfolio. In a subsequent paper, DeMiguel et al. (2009a) propose a method that does not focus on the input parameters of the optimization problem, but directly on the portfolio weights. The authors add $\mathcal{L}^1$ or $\mathcal{L}^2$ penalty terms to a minimum variance objective function, thereby penalizing large parameter values. These approaches, also known as the lasso and ridge regression in the statistical learning literature, are commonly utilized not only to deal with parameter uncertainty but also to prevent over-fitting, see e.g., Friedman et al. (2001). DeMiguel et al. (2009a) find that using these regularization approaches in the portfolio optimization context allows for building portfolios that produce improved risk-return performance compared to the naive $1/N$ portfolio. Furthermore, Brodie et al. (2009) also use a $\mathcal{L}^1$ penalty term in the objective function and show that the mean-variance portfolio constructed using this rule consistently outperforms the $1/N$ portfolio. The authors provide additional interpretations to the penalization term, interpreting it as a proxy for trading costs and showing that for a particular value of the penalization parameter, the lasso constraint is equivalent to a no short-sale constraint. Finally, Li (2015) employs an elastic net regularization (i.e., both a lasso and a ridge penalty term) for the mean-variance optimization problem and shows that the resulting portfolios manage to outperform other well-known portfolio optimization rules previously proposed in the literature, as well as the $1/N$ portfolio.

All of the shrinkage and penalization approaches mentioned above extend the basic mean-variance portfolio optimization framework but only deal with assets denominated in a single currency. Our paper contributes to the literature on international portfolio optimization by incorporating the regularization techniques aimed at improving the out-of-sample portfolio performance to the framework of multi-currency asset allocation. Moreover, we, in particular, study the essential treatment of currency risk management and compare the currency overlay approach to the regularized joint optimization of assets and currencies. In addition to a theoretical interest in studying the problem of efficient multi-currency asset allocation, this is a practically oriented issue that is extensively debated in the asset and wealth management industry.

3. Model

In this section, a theoretical framework is presented. We consider a multi-currency-based asset allocation framework as derived in Ulrych and Vasiljević (2020). Building on this model, we formulate separate and joint optimization problems for the optimization of asset and currency weights, introduce the regularized versions of the international asset allocation problems, and provide numerical solutions in terms of quadratic programming algorithms.

3.1. Hedged portfolio return

Let $P_{i,t}$ represent a price of an asset i at time t, expressed in a local currency, with $R_{i,t+1}$ denoting the simple return of this asset over the period t to $t+1$.[2] Given an arbitrary domestic currency, $S_{c_i,t}$ denotes a spot currency exchange rate in domestic currency per unit of foreign currency c_i at time t, where c_i denotes a local currency asset i is denoted in. Let $e_{c_i,t+1}$ denote a currency exchange return from t to $t+1$ of a currency c_i with respect to the domestic currency. An international investor is concerned with returns expressed in her domestic currency. Thereby, an unhedged return $\tilde{R}^u_{i,t+1}$ of an asset i is given by

[2] Local currency is simply the currency in which a certain asset is denominated. Given an investor's base currency, a local currency can refer to either a domestic or a foreign currency.

$$\tilde{R}^u_{i,t+1} = \frac{P_{i,t+1} S_{c_i,t+1}}{P_{i,t} S_{c_i,t}} - 1 = R_{i,t+1} + e_{c_i,t+1} + R_{i,t+1} e_{c_i,t+1}, \tag{1}$$

where, in line with the existing literature, we assume that no rebalancing takes place between t and $t+1$. Observe that the unhedged asset return in domestic currency is expressed as a sum of the asset return in the local currency (i.e., $R_{i,t+1}$), the return of the currency exchange rate (i.e., $e_{c_i,t+1}$), and the (second-order) cross-term between the two (i.e., $R_{i,t+1} e_{c_i,t+1}$).

Now, consider an investor with a portfolio $\mathcal{P}$ consisting of $N \in \mathbb{N}$ assets. These assets can be denominated in different currencies $c = 1, 2, \ldots, K+1$, where $c = 1$ denotes the domestic currency. This implies that by investing in these assets, the investor can, next to the domestic currency, acquire an implicit exposure to up to $K \leq N$ foreign currencies. The investor's portfolio return in domestic currency is given by the weighted average of the single asset returns. Denote with $x_{i,t}$ the relative weight of the wealth invested in asset $i = 1, 2, \ldots, N$ and held from t to $t+1$. Then, the unhedged portfolio return $\tilde{R}^u_{\mathcal{P},t+1}$ is expressed by

$$\tilde{R}^u_{\mathcal{P},t+1} = \sum_{i=1}^{N} x_{i,t} \tilde{R}^u_{i,t+1},$$

where, since the asset weights are given as a fraction of the total portfolio value, $\sum_{i=1}^{N} x_{i,t} = 1$ holds for every t.

Next, we assume that for every time t and every currency c there exists a currency forward contract with a forward exchange rate $F_{c,t}$ expressed in domestic currency per unit of foreign currency c and delivery at $t+1$. Furthermore, we assume that $F_{c,t}$ is set in a way that the price of the corresponding forward contract equals zero at inception. The forward premium (i.e., a rate of return corresponding to the forward exchange rate) is defined by $f_{c,t} := (F_{c,t} - S_{c,t})/S_{c,t}$. Using this expression, the return of a short position in a currency forward contract for currency c is given by $(F_{c,t} - S_{c,t+1})/S_{c,t} = f_{c,t} - e_{c,t+1}$. Because an investor wishes to hedge her implicit foreign currency exposures, she sells these currencies forward. We denote with $\phi_{c,t}$ the relative notional value of a forward exchange rate contract for currency c at time t, expressed as a fraction of total portfolio value. A short position in the forward contract is indicated by $\phi_{c,t} > 0$.

An investor now has a possibility to not only hedge her implicit currency exposure to the currencies $c = 2, \ldots, K+1$, but also to expand her portfolio with forward contracts in currencies $c = K+2, \ldots, M+1$, to which she does not yet have an implicit exposure. The return of a hedged portfolio is then expressed by

$$\tilde{R}^h_{\mathcal{P},t+1} = \tilde{R}^u_{\mathcal{P},t+1} + \sum_{c=2}^{M+1} \phi_{c,t}(f_{c,t} - e_{c,t+1}). \tag{2}$$

For the domestic currency $c = 1$, $S_{1,t} = F_{1,t} = 1$ and $e_{1,t} = f_{1,t} = 0$ trivially hold for all t. This implies that the choice of $\phi_{1,t}$ is arbitrary.

The goal of the inclusion of currency forward contracts into the portfolio is to reduce the currency risk stemming from the investment in foreign assets. In other words, an investor is interested in optimally managing the exposure to foreign currencies. We define the (net) currency exposure to currency c at time t by $\psi_{c,t} := w_{c,t} - \phi_{c,t}$, where $w_{c,t}$ denotes the implicit exposure to currency c. This implicit currency exposure can be expressed as $w_{c,t} = \sum_{i \in \mathcal{A}_c} x_{i,t}$, with $\mathcal{A}_c$ representing the set of all assets denoted in currency c. The case of zero exposure to foreign currencies, $\psi_{c,t} = 0$ for all $c = 2, \ldots, M+1$, corresponds to the fully hedged portfolio, which is equivalent to $\phi_{c,t} = w_{c,t}$. Moreover, we assume that $\sum_{c=1}^{M+1} \phi_{c,t} = 1$, which implies

$$\phi_{1,t} = 1 - \sum_{c=2}^{M+1} \phi_{c,t}, \tag{3}$$

showing that the domestic weight $\phi_{1,t}$ is trivially determined after the foreign currency weights have been set. We impose the summation to one such that, using Eq. (3) and the definition of $\psi_{c,t}$, one obtains $\psi_{1,t} = -\sum_{c=2}^{M+1} \psi_{c,t}$. Therefore, the sum of all net currency exposures is equal to zero, which implies that the currency portfolio is a zero investment portfolio (i.e., a hedging portfolio).

Building on the definition of $\mathcal{A}_c$, we define a matrix $\boldsymbol{C} \in \mathbb{R}^{N \times M}$, whose rows indicate a foreign currency denomination of an asset i by setting $\boldsymbol{C}_{i,c} = 1$ if asset i is traded in currency c and zero otherwise. This can then be written as

$$\boldsymbol{C}_{i,c} = \begin{cases} 1, & \text{if } i \in \mathcal{A}_{c+1}, \\ 0, & \text{otherwise}, \end{cases} \tag{4}$$

where the notation $c+1$ is used so that the first column indicates all the assets belonging to the first foreign currency and the domestic currency is skipped. This notation proves convenient in the equations below and arises from the above statement that the choice of the domestic currency weight $\phi_{1,t}$ (and equivalently $\psi_{1,t}$) is implied by the choice of the other M currency weights. This notation allows us to express a vector of implied foreign currency exposures $\boldsymbol{w}_t = (w_{2,t}, \ldots, w_{M+1,t})^T \in \mathbb{R}^{M \times 1}$ at time t as

$$\boldsymbol{w}_t = \boldsymbol{C}^T \boldsymbol{x}_t, \tag{5}$$

where $\boldsymbol{x}_t \in \mathbb{R}^{N \times 1}$ represents the vector of asset weights at time t and $\boldsymbol{C}^T$ is the transpose of the matrix $\boldsymbol{C}$. The domestic implicit exposure is trivially computed as $w_{1,t} = 1 - \boldsymbol{1}_M^T \boldsymbol{w}_t$, where $\boldsymbol{1}_M \in \mathbb{R}^{M \times 1}$ denotes a vector of ones. Note that the vector $\boldsymbol{w}_t$ is of dimension $(M \times 1)$, even though the assets in the portfolio are implicitly exposed to K foreign currencies. This implies that the elements $w_{K+2,t} = w_{K+3,t} = \ldots = w_{M+1,t} = 0$ are all equal to zero and the same holds for the corresponding rows of matrix $\boldsymbol{C}$. Matrix $\boldsymbol{C}$ allows us to express the unhedged return in domestic currency from Eq. (1) in a vectorized form as

$$\tilde{\boldsymbol{R}}^{u}_{t+1} = \boldsymbol{R}_{t+1} + \boldsymbol{C}\boldsymbol{e}_{t+1} + (\boldsymbol{R}_{t+1} \odot \boldsymbol{C}\boldsymbol{e}_{t+1}), \tag{6}$$

where $\odot$ denotes the Hadamard product (i.e., the element-wise product) and $\tilde{\boldsymbol{R}}^{u}_{t+1} \in \mathbb{R}^{N\times 1}$, $\boldsymbol{R}_{t+1} \in \mathbb{R}^{N\times 1}$, and $\boldsymbol{e}_{t+1} = (e_{2,t+1}, ..., e_{M+1,t+1})^T \in \mathbb{R}^{M\times 1}$ are the unhedged return vector in domestic currency, the vector of simple returns in local currencies, and the vector of simple currency returns, respectively.[3] Taking $\boldsymbol{\phi}_t = (\phi_{2,t}, ..., \phi_{M+1,t}) \in \mathbb{R}^{M\times 1}$ and $\boldsymbol{f}_t = (f_{2,t}, ..., f_{M+1,t})^T \in \mathbb{R}^{M\times 1}$ as the vectors of currency forward weights and forward premiums, respectively, Eq. (2) can be expressed in a vectorized form as

$$\tilde{R}^{h}_{\mathcal{P},t+1} = \boldsymbol{x}_t^T \tilde{\boldsymbol{R}}^{u}_{t+1} + \boldsymbol{\phi}_t^T (\boldsymbol{f}_t - \boldsymbol{e}_{t+1}). \tag{7}$$

As mentioned above, the fully hedged portfolio is obtained by $\boldsymbol{\phi}_t = \boldsymbol{w}_t$. Using Eq. (5), this can be expressed as $\boldsymbol{\phi}_t = \boldsymbol{C}^T \boldsymbol{x}_t$. Plugging this into Eq. (7) yields the expression for the return of the fully hedged portfolio $\tilde{R}^{fh}_{\mathcal{P},t+1}$ given by

$$\tilde{R}^{fh}_{\mathcal{P},t+1} = \boldsymbol{x}_t^T \tilde{\boldsymbol{R}}^{u}_{t+1} + \boldsymbol{x}_t^T \boldsymbol{C}(\boldsymbol{f}_t - \boldsymbol{e}_{t+1}) = \boldsymbol{x}_t^T (\boldsymbol{R}_{t+1} + \boldsymbol{C}\boldsymbol{f}_t + (\boldsymbol{R}_{t+1} \odot \boldsymbol{C}\boldsymbol{e}_{t+1})) = \boldsymbol{x}_t^T \tilde{\boldsymbol{R}}^{fh}_{t+1},$$

where $\tilde{\boldsymbol{R}}^{fh}_{t+1} = \boldsymbol{R}_{t+1} + \boldsymbol{C}\boldsymbol{f}_t + (\boldsymbol{R}_{t+1} \odot \boldsymbol{C}\boldsymbol{e}_{t+1})$ denotes the vector of fully hedged asset returns. This is a return that can be achieved by matching a long position in asset i with a short position in a forward contract for a currency c_i at time t, and vice versa. Compared to the unhedged returns from Eq. (6), the term $\boldsymbol{C}\boldsymbol{e}_{t+1}$ is replaced by $\boldsymbol{C}\boldsymbol{f}_t$. This shows that instead of the random currency return $\boldsymbol{e}_{t+1}$, the fully hedged investor secures the safe forward premium $\boldsymbol{f}_t$, thereby reducing the uncertainty arising from fluctuating currency exchange rates. The investor is, however, not able to eliminate all currency risk in the portfolio, as the cross-product term $\boldsymbol{R}_{t+1} \odot \boldsymbol{C}\boldsymbol{e}_{t+1}$ remains. This is also the reason why a perfect currency hedge is impossible to achieve. It would require an investor to know the value of her foreign assets at the maturity of the hedge $t+1$ already when choosing the forward weights at the inception of the currency forward contract at time t.

3.2. International portfolio optimization problem

The international asset allocation problem involves a combined choice of investing in assets and currencies. The corresponding weights $\boldsymbol{x}_t$ and $\boldsymbol{\phi}_t$, as denoted in this paper, are determined such that they maximize a given objective function. We employ a classical mean-variance objective function and formulate two possible optimization problems that arise in the international portfolio optimization setting: i) the currency overlay, where first the asset positions are determined, and subsequently, given the chosen asset positions, the currency weights are optimized, and ii) the joint optimization setting where assets and currencies are determined in a single optimization.

We start with the specification of a separate portfolio optimization problem. By $\boldsymbol{\mu}^a = \mathbb{E}(\tilde{\boldsymbol{R}}^{fh}_{t+1})$ and $\boldsymbol{\mu}^c = \mathbb{E}(\boldsymbol{f}_t - \boldsymbol{e}_{t+1})$ denote the expectations of the fully hedged asset returns and the currency forward excess returns (i.e., arising from the payoffs of the entered currency forward contracts), respectively. The terms $\boldsymbol{\Sigma}^a = \mathrm{Var}(\tilde{\boldsymbol{R}}^{fh}_{t+1})$ and $\boldsymbol{\Sigma}^c = \mathrm{Var}(\boldsymbol{f}_t - \boldsymbol{e}_{t+1})$ denote the $(N \times N)$ covariance matrix of the fully hedged asset returns and the $(M \times M)$ covariance matrix of the currency forward excess returns, respectively. Furthermore, we define $\boldsymbol{\Sigma}^{ac} = \mathrm{Cov}(\tilde{\boldsymbol{R}}^{u}_{t+1}, \boldsymbol{f}_t - \boldsymbol{e}_{t+1})$ as the $(N \times M)$ cross-covariance matrix between asset and currency returns. The term $\gamma > 0$ represents the coefficient of relative risk aversion. The separate international portfolio optimization problem is then given by

$$\begin{aligned}
\boldsymbol{x}_t^* &= \arg\max_{\boldsymbol{x}_t} \left\{ \ \mathbb{E}(\tilde{R}^{fh}_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^{fh}_{\mathcal{P},t+1}) \ \right\} \\
&= \arg\max_{\boldsymbol{x}_t} \left\{ \ \boldsymbol{x}_t^T \boldsymbol{\mu}^a - \frac{\gamma}{2} \boldsymbol{x}_t^T \boldsymbol{\Sigma}^a \boldsymbol{x}_t \ \right\} \\
&\qquad \text{subject to} \qquad \boldsymbol{x}_t^T \mathbf{1}_N = 1, \\
\boldsymbol{\phi}_t^* \mid \boldsymbol{x}_t^* &= \arg\max_{\boldsymbol{\phi}_t} \left\{ \ \mathbb{E}(\tilde{R}^{h}_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^{h}_{\mathcal{P},t+1}) \mid \boldsymbol{x}_t^* \ \right\} \\
&= \arg\max_{\boldsymbol{\phi}_t} \left\{ \ \boldsymbol{\phi}_t^T \boldsymbol{\mu}^c - \frac{\gamma}{2} \boldsymbol{\phi}_t^T \boldsymbol{\Sigma}^c \boldsymbol{\phi}_t - \gamma \boldsymbol{x}_t^{*T} \boldsymbol{\Sigma}^{ac} \boldsymbol{\phi}_t \ \mid \boldsymbol{x}_t^* \right\}.
\end{aligned} \tag{8}$$

The first optimization step is a standard mean-variance optimization of the asset portfolio. Here, the currencies are not considered and the goal is to optimize the mean-variance trade-off for the fully hedged portfolio.[4] In the second optimization step, the currency overlay step, an investor maximizes the mean-variance objective conditionally on the already determined asset weights (i.e., as obtained in the first optimization step), with the hedged return as the object of interest, focusing on the optimization of the currency exposure. Note that only the asset weights need to add up to one. For the currency weights, the weight $\phi_{1,t}$, which is not a part of the optimization, can simply be chosen such that the currency weights sum up to one, see Eq. (3).

Given the conditioning on $\boldsymbol{x}_t^*$, terms that do not depend on $\boldsymbol{\phi}_t$ can be dropped out of the formulation of the second optimization step. Notice that the choice of the optimal currency weights is only influenced by the previously determined asset weights through

[3] The Hadamard product, also known as the element-wise product, is an operation that takes two matrices of the same dimensions and produces another matrix, where each element (i, j) is the product of elements (i, j) of the original two matrices.

[4] In the separate optimization approach, an investor optimizes the asset allocation without the consideration of the underlying implicit currency exposure. This is the reason for the fully hedged, as opposed to the unhedged, formulation of the first optimization step. The currency allocation is then performed in the second (i.e., currency overlay) optimization step.

the term $\boldsymbol{x}_t^{*T}\boldsymbol{\Sigma}^{ac} = \mathrm{Cov}(\tilde{R}^u_{\mathcal{P},t+1}, \boldsymbol{f}_t - \boldsymbol{e}_{t+1})$, which is a $(1 \times M)$ vector representing the covariances between the unhedged portfolio return and the currency forward returns. Clearly, the correlations between the single assets and currencies are not considered. Thus, the separate optimization approach can be understood as a form of dimensionality reduction. Moreover, one can think of specialized equity (bond) asset managers being employed in the first optimization step, see Van Binsbergen et al. (2008). Thereby, asset positions are determined, and a specialized equity (bond) asset manager does not speculate on currencies when deciding on the asset allocations – hence, the fully hedged return specification in the first optimization step. Subsequently, a specialized currency overlay manager determines the optimal currency allocations, given the already determined asset allocation. As discussed in Kim and Chance (2018), the described separate optimization is a standard approach of optimizing international portfolios in practice.

Next, we formulate a joint international asset allocation optimization problem. Consider the vector $\boldsymbol{\theta}_t = (x_{1,t}, ..., x_{N,t}, \phi_{2,t}, ..., \phi_{M+1,t})^T \in \mathbb{R}^{(N+M)\times 1}$ which is a combined vector of the asset weights $\boldsymbol{x}_t$ and the currency weights $\boldsymbol{\phi}_t$. We denote by $\boldsymbol{r}_{t+1} = (\tilde{R}^u_{1,t+1}, ..., \tilde{R}^u_{N,t+1}, (f_{2,t} - e_{2,t+1}), ..., (f_{M+1,t} - e_{M+1,t+1}))$ the combined vector of asset returns expressed in a domestic currency and currency forward excess returns. Furthermore, the term $\boldsymbol{\mu} = \mathbb{E}(\boldsymbol{r}_{t+1})$ denotes the $((N+M)\times 1)$ vector of its expected returns and $\boldsymbol{\Sigma} = \mathrm{Var}(\boldsymbol{r}_{t+1})$ the corresponding $((N+M)\times(N+M))$ covariance matrix. The joint international portfolio optimization problem is given by

$$\begin{aligned} \boldsymbol{\theta}_t^* &= \underset{\boldsymbol{\theta}_t}{\arg\max} \left\{ \ \mathbb{E}(\tilde{R}^h_{\mathcal{P},t+1}) - \frac{\gamma}{2}\mathrm{Var}(\tilde{R}^h_{\mathcal{P},t+1}) \ \right\} \\ &= \underset{\boldsymbol{\theta}_t}{\arg\max} \left\{ \ \boldsymbol{\theta}_t^T\boldsymbol{\mu} - \frac{\gamma}{2}\boldsymbol{\theta}_t^T\boldsymbol{\Sigma}\boldsymbol{\theta}_t \ \right\} \\ &\qquad \text{subject to} \qquad \boldsymbol{\theta}_t^T\mathbf{q}_{N,M} = 1, \end{aligned} \tag{9}$$

where $\mathbf{q}_{N,M}$ denotes a $((N+M)\times 1)$ vector with the first N elements equal to one and the remaining M elements equal to zero. Therefore, the constraint is equivalent to $\boldsymbol{x}^T\mathbf{1}_N = 1$ in the separate optimization problem.

This formulation of the optimization problem determines the mean-variance efficient vector $\boldsymbol{\theta}_t$ by jointly optimizing over the asset weights and currency exposures. Here, also the correlations between the single assets and currencies are considered. The optimization problem is, correspondingly, of higher dimension compared to the separate optimization. As such, it is more prone to parameter uncertainty and over-fitting, which can potentially lead to a worsened out-of-sample portfolio performance. To tackle this problem, we introduce regularization to the international asset allocation problem in the next section.

3.3. Sparse and stable international portfolios

The solutions to both the separate and the joint optimization problem depend on the true expected returns and variances that are generally unknown. One commonly used approach is based on using the sample counterparts $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ in place of the true parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. This plug-in approach, however, tends to yield poor out-of-sample performance, see DeMiguel et al. (2009b). This is the reason for the introduction of alternative methods that deal with the problem of parameter uncertainty in international portfolio optimization. This section introduces some extensions to the standard mean-variance framework and translates them into the context of the multi-currency portfolio optimization problem. We start by focusing on the sparsity of asset and currency weights, introduce stable international portfolios, and subsequently combine the two into a SAS international asset allocation approach.

3.3.1. Sparse multi-currency portfolio

A sparse portfolio is a portfolio with some, or possibly many, of the portfolio weights set to zero. This is a desirable property, especially when the asset universe is large. Furthermore, setting some weights to zero has the potential benefit of lowering transaction costs. The idea of sparse portfolios has been studied in the finance literature, Britten-Jones (1999), for example, introduce an F-test to test whether multiple asset weights are statistically different from zero. Another approach is presented by Garlappi et al. (2007), who propose setting all portfolio weights to zero (i.e., only investing in risk-free assets) when the squared Sharpe ratio is not statistically different from zero.

Sparsity is introduced into the portfolio optimization problem by constraining the $\mathcal{L}^1$ vector norm of the portfolio weights. The $\mathcal{L}^1$-norm is defined by $\|\boldsymbol{x}\|_1 = \sum_{i=1}^N |x_i|$. Sparsity arises from the fact that the optimal portfolio is usually located at one of the corners of the ($\mathcal{L}^1$-induced) constraint region, where one or more weights are equal to zero, e.g., see Figure 3.12 in Hastie et al. (2009). Next to sparsity, there are multiple potential benefits of introducing a constraint on the $\mathcal{L}^1$-norm of the portfolio weights. Such a constraint can induce stability into a portfolio by reducing the sensitivity of the optimization to potential collinearities between the assets, see Daubechies et al. (2004). Through the introduction of the constraint, the estimation becomes biased; however, it has the potential to improve the out-of-sample portfolio performance through a bias-variance trade-off, see Tibshirani (1996). Furthermore, the constraint also limits the amount of short-sale permitted in the portfolio, see DeMiguel et al. (2009a), which is a desirable property from a practical point of view. Another advantage is that the constraint can be seen as a proxy for transaction costs, see Brodie et al. (2009).[5]

[5] Transaction costs normally consist of two components: a fixed "overhead" cost for every transaction, which is usually comparably small for large investments, and a proportional cost that results from the bid-ask spread of the market. As the $\mathcal{L}^1$ constraint limits the absolute size of the portfolio weights, this can also be seen as imposing a limit on the proportional transaction costs.

Next, we specify a sparse international portfolio optimization problem. Consider a separate optimization problem as defined in Eq. (8). Consistent with the literature on regularization, we cast the optimization problem in terms of minimization and introduce the $\mathcal{L}^1$-normed penalty directly in the objective function. The sparse separate optimization is given by

$$\begin{aligned} \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{x}_t} \left\{ \frac{\gamma}{2}\boldsymbol{x}_t^T\hat{\boldsymbol{\Sigma}}^a\boldsymbol{x}_t - \boldsymbol{x}_t^T\hat{\boldsymbol{\mu}}^a + \lambda_1^a\|\boldsymbol{x}_t\|_1 \right\} \\ &\qquad \text{subject to} \qquad \boldsymbol{x}_t^T\mathbf{1}_N = 1, \\ \boldsymbol{\phi}_t^* \mid \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{\phi}_t} \left\{ \frac{\gamma}{2}\boldsymbol{\phi}_t^T\hat{\boldsymbol{\Sigma}}^c\boldsymbol{\phi}_t - \boldsymbol{\phi}_t^T\hat{\boldsymbol{\mu}}^c + \gamma\boldsymbol{x}_t^{*T}\hat{\boldsymbol{\Sigma}}^{ac}\boldsymbol{\phi}_t + \lambda_1^c\|\boldsymbol{\phi}_t\|_1 \mid \boldsymbol{x}_t^* \right\}, \end{aligned} \tag{10}$$

where the symbol ˆ denotes a sample estimate, and the parameters λ_1^a and λ_1^c govern the penalization of the assets and currencies, respectively.[6] Employing two separate parameters for assets and currencies allows for a more precise calibration of the model to the international portfolio optimization problem. A single parameter for both would indicate that both should be treated in the same way, i.e., currencies are just other assets. However, as explained above, only the assets are constrained to add up to one. Moreover, the transaction costs for currencies are generally lower than those of equities. The interpretation of the lasso constraint as a proxy for transaction costs also provides motivation for a separate penalization structure. The sparsity property induced by the lasso constraint could also imply that possibly many of the currency weights can become equal to zero if there is a single penalization parameter for assets and currencies. However, this is not a desirable property from the currency risk management perspective.

Equivalently to the separate optimization case, the sparse joint optimization problem is given by

$$\begin{aligned} \boldsymbol{\theta}_t^* &= \arg\min_{\boldsymbol{\theta}_t} \left\{ \frac{\gamma}{2}\boldsymbol{\theta}_t^T\hat{\boldsymbol{\Sigma}}\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^T\hat{\boldsymbol{\mu}} + \|\boldsymbol{S}\boldsymbol{\theta}_t\|_1 \right\} \\ &\qquad \text{subject to} \qquad \boldsymbol{\theta}_t^T\mathbf{q}_{N,M} = 1, \end{aligned} \tag{11}$$

where $\boldsymbol{S}$ is an $((N+M)\times(N+M))$ matrix with all off-diagonal elements equal to zero, the first N diagonal elements equal to λ_1^a and the last M diagonal elements equal to λ_1^c. Such specification again allows for different optimal levels of penalization for assets and currencies.

The $\mathcal{L}^1$-norm proposed above is non-differentiable with respect to the weight vector, which is why the lasso problem does not yield an analytical solution. There are different techniques to solve the lasso problem numerically, such as the conjugate-gradient method or the LARS algorithm, see Efron et al. (2004). In this paper, we introduce various types of constraints on the asset and currency weights and consequently require a general algorithm that is able to handle such constraints. We follow the approach described by Roncalli (2013), where auxiliary variables are introduced and the problem is reformulated as a quadratic programming problem in a higher dimension. For details on the exact specification of this quadratic program, see the Online Appendix.

3.3.2. Stable multi-currency portfolio

The stability of portfolio weights is another property that is desirable for the investor. The standard mean-variance efficient portfolio is sensitive to changes or errors in the inputs and is prone to producing extreme weights. This is problematic in practice since unstable weights have the potential to cause high transaction costs over time that can significantly diminish realized portfolio performance. Ledoit and Wolf (2003) show that instability is related to the estimation errors of the sample covariance matrix and its inverse needed for computing the mean-variance efficient portfolio. The inverse becomes especially problematic in the presence of multicollinearity. These instabilities motivate several approaches in the literature that aim at constraining the portfolio in a way that induces stability and thereby reduces the problem of parameter uncertainty.

A well-known approach of shrinking the covariance matrix is proposed in Ledoit and Wolf (2004). The authors propose a shrinkage estimator for a covariance matrix given by

$$\hat{\boldsymbol{\Sigma}}_s = \upsilon\hat{\boldsymbol{\Sigma}}_g + (1-\upsilon)\hat{\boldsymbol{\Sigma}}, \tag{12}$$

where $\hat{\boldsymbol{\Sigma}}$ is a sample covariance matrix, $\hat{\boldsymbol{\Sigma}}_g$ is a shrinkage target with a lower variance, see Ledoit and Wolf (2004) for details, and υ is a shrinkage intensity. Furthermore, the authors demonstrate that optimized portfolios resulting from the covariance matrix shrinkage improve out-of-sample performance. Moreover, Li (2015) shows that using a shrinkage estimator of the form of Eq. (12) is equal to a linear transformation of the portfolio weights that depends on both the shrinkage target $\hat{\boldsymbol{\Sigma}}_g$ and the shrinkage intensity υ. Another interpretation offered by the author is that the shrinkage estimator can be seen as a constraint on the $\mathcal{L}^2$-norm, weighted by the target covariance matrix $\hat{\boldsymbol{\Sigma}}_g$.

While the implementation of this estimator is straightforward for the joint international portfolio optimization approach, where the matrix $\hat{\boldsymbol{\Sigma}}$ can simply be replaced by its shrunk version $\hat{\boldsymbol{\Sigma}}_s$, this does not apply for the separate optimization problem. The reason for this is that the cross-covariance matrix between the assets and currencies $\hat{\boldsymbol{\Sigma}}^{ac}$ should theoretically be shrunk as well. This can be achieved by slightly modifying the currency overlay step in the separate optimization problem. We do so by augmenting the currency overlay problem by one dimension and introducing the vector $\boldsymbol{\varphi}_t = (1, \phi_{2,t}, .., \phi_{M+1,t})$. Here, the first weight corresponds to the asset portfolio. Its weight is already predetermined from the asset optimization step since the asset weights sum up to one.

[6] For $\lambda_1^a = \lambda_1^c = 0$, the penalization is absent, and the optimization problem corresponds to the one from Eq. (8). Moreover, larger choices of λ_1^a and λ_1^c imply larger penalization of the absolute portfolio weights, which can be interpreted as shrinking the weights towards zero.

Accordingly, we introduce parameters $\hat{\boldsymbol{\Sigma}}^{\tilde{c}} \in \mathbb{R}^{(M+1)\times(M+1)}$ and $\hat{\boldsymbol{\mu}}^{\tilde{c}} \in \mathbb{R}^{(M+1)\times 1}$, which represent the augmented covariance matrix and the augmented mean vector of the joint asset portfolio and currencies vector, as defined below. The stable separate international portfolio optimization with shrinkage of the covariance matrix is given by

$$\begin{aligned} \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{x}_t} \left\{ \frac{\gamma}{2} \boldsymbol{x}_t^T \hat{\boldsymbol{\Sigma}}_s^a \boldsymbol{x}_t - \boldsymbol{x}_t^T \hat{\boldsymbol{\mu}}^a \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{x}_t^T \mathbf{1} = 1, \\ \boldsymbol{\varphi}_t^* \mid \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{\varphi}_t} \left\{ \frac{\gamma}{2} \boldsymbol{\varphi}_t^T \hat{\boldsymbol{\Sigma}}_s^{\tilde{c}} \boldsymbol{\varphi}_t - \boldsymbol{\varphi}_t^T \hat{\boldsymbol{\mu}}^{\tilde{c}} \mid \boldsymbol{x}_t^* \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{A}\boldsymbol{\varphi}_t = 1, \end{aligned} \tag{13}$$

with

$$\boldsymbol{A} = \begin{pmatrix} 1 & \mathbf{0} \end{pmatrix}, \qquad \hat{\boldsymbol{\Sigma}}^{\tilde{c}} = \begin{pmatrix} \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^a \boldsymbol{x}_t^* & \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^{ac} \\ (\boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^{ac})^T & \hat{\boldsymbol{\Sigma}}^c \end{pmatrix}, \text{ and } \qquad \hat{\boldsymbol{\mu}}^{\tilde{c}} = \begin{pmatrix} \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\mu}}^a \\ \hat{\boldsymbol{\mu}}^c \end{pmatrix}.$$

Here, the linear equality constraint in the second optimization step simply enforces that the weight assigned to the asset portfolio is equal to one, and the optimization is performed only with respect to the currency weights. The matrices $\hat{\boldsymbol{\Sigma}}_s^a$ and $\hat{\boldsymbol{\Sigma}}_s^{\tilde{c}}$ are calculated via Eq. (12), whereby the shrinkage target and the shrinkage intensity are computed as proposed in Ledoit and Wolf (2004).[7]

The stable joint international portfolio optimization with shrinkage of the covariance matrix is then accordingly given by

$$\begin{aligned} \boldsymbol{\theta}_t^* &= \arg\min_{\boldsymbol{\theta}_t} \left\{ \frac{\gamma}{2} \boldsymbol{\theta}_t^T \hat{\boldsymbol{\Sigma}}_s \boldsymbol{\theta}_t - \boldsymbol{\theta}_t^T \hat{\boldsymbol{\mu}} \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{\theta}_t^T \mathbf{q}_{N,M} = 1. \end{aligned} \tag{14}$$

Another approach to induce stability in the portfolio is achieved by constraining the $\mathcal{L}^2$-norm of the portfolio weights. This approach is well-connected to the shrinkage of the covariance matrix described above. While $\mathcal{L}^1$ penalization induces extreme shrinkage (i.e., sparsity), the $\mathcal{L}^2$ penalization induces smooth shrinkage (i.e., stability). As Li (2015) shows, imposing a ridge type of constraint is equivalent to shrinking the covariance matrix towards the identity matrix. Therefore, constraining the $\mathcal{L}^2$-norm of the portfolio weights or employing a covariance matrix shrinkage estimator can be interpreted as two analogous approaches.

3.3.3. Sparse and stable multi-currency portfolio

Last, we combine the methods introduced above and construct international portfolios that are both sparse and stable. We start with the separate optimization and combine sparsity as induced by the $\mathcal{L}^1$ penalization, such as in Eq. (10) with shrinkage of the covariance matrix as presented in Eq. (13). Thereby, the SAS separate international portfolio optimization with shrinkage of the covariance matrix is given by

$$\begin{aligned} \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{x}_t} \left\{ \frac{\gamma}{2} \boldsymbol{x}_t^T \hat{\boldsymbol{\Sigma}}_s^a \boldsymbol{x}_t - \boldsymbol{x}_t^T \hat{\boldsymbol{\mu}}^a + \lambda_1^a \|\boldsymbol{x}_t\|_1 \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{x}_t^T \mathbf{1} = 1, \\ \boldsymbol{\varphi}_t^* \mid \boldsymbol{x}_t^* &= \arg\min_{\boldsymbol{\varphi}_t} \left\{ \frac{\gamma}{2} \boldsymbol{\varphi}_t^T \hat{\boldsymbol{\Sigma}}_s^{\tilde{c}} \boldsymbol{\varphi}_t - \boldsymbol{\varphi}_t^T \hat{\boldsymbol{\mu}}^{\tilde{c}} + \lambda_1^c \|\boldsymbol{\varphi}_t\|_1 \mid \boldsymbol{x}_t^* \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{A}\boldsymbol{\varphi}_t = 1, \end{aligned} \tag{15}$$

with

$$\boldsymbol{A} = \begin{pmatrix} 1 & \mathbf{0} \end{pmatrix}, \qquad \hat{\boldsymbol{\Sigma}}^{\tilde{c}} = \begin{pmatrix} \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^a \boldsymbol{x}_t^* & \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^{ac} \\ (\boldsymbol{x}_t^{*T} \hat{\boldsymbol{\Sigma}}^{ac})^T & \hat{\boldsymbol{\Sigma}}^c \end{pmatrix}, \quad \text{and} \quad \hat{\boldsymbol{\mu}}^{\tilde{c}} = \begin{pmatrix} \boldsymbol{x}_t^{*T} \hat{\boldsymbol{\mu}}^a \\ \hat{\boldsymbol{\mu}}^c \end{pmatrix}.$$

Here, $\hat{\boldsymbol{\Sigma}}_s^a$ and $\hat{\boldsymbol{\Sigma}}_s^c$ are the shrinkage estimators of the covariance matrix as specified in Eq. (12), with the shrinkage target and the shrinkage intensity as proposed in Ledoit and Wolf (2004).

Equivalently, the SAS joint international portfolio optimization with shrinkage of the covariance matrix is given by

$$\begin{aligned} \boldsymbol{\theta}_t^* &= \arg\min_{\boldsymbol{\theta}_t} \left\{ \frac{\gamma}{2} \boldsymbol{\theta}_t^T \hat{\boldsymbol{\Sigma}}_s \boldsymbol{\theta}_t - \boldsymbol{\theta}_t^T \hat{\boldsymbol{\mu}} + \|\boldsymbol{S}\boldsymbol{\theta}_t\|_1 \right\}, \\ &\quad \text{subject to} \quad \boldsymbol{\theta}_t^T \mathbf{q}_{N,M} = 1, \end{aligned} \tag{16}$$

with parameters as already defined before. The optimization problems from Eqs. (15) and (16) are solved analogously to the augmented quadratic problem as specified in the Online Appendix, with the appropriately shrunk covariance matrix.

Instead of employing the covariance shrinkage to construct SAS portfolios, it is also possible to induce stability into the sparse portfolio by constraining also the $\mathcal{L}^2$-norm of the portfolio weights. Such a specification that includes both the lasso and the ridge

[7] The shrinkage target is constructed from the common constant correlation (i.e., the average of all sample correlations) and the vector of sample variances. The shrinkage intensity is optimized with respect to a quadratic measure of distance between the true and the estimated covariance matrices based on the Frobenius norm. For more information, consult the original paper.

Table 1
Summary statistics: 1999 to 2019.

	Unhedged	Fully Hedged	Unhedged	Fully Hedged
	Annual Average Return (%)		Annual Volatility (%)	
Australia	10.91	7.49	20.67	12.20
Canada	10.20	8.91	18.87	12.87
Switzerland	9.70	9.28	18.99	17.18
Eurozone	6.66	7.08	19.24	14.94
United Kingdom	5.73	6.05	16.44	13.51
Japan	5.51	7.13	16.26	17.23
Norway	12.01	11.18	24.73	18.63
New Zealand	12.79	8.40	19.10	10.93
Singapore	10.35	10.02	20.98	18.17
United States	7.75		14.16	

This table presents summary statistics for a dataset comprising 21 equity broad market indices from 1999 to 2019. All returns are converted to US dollars. A fully hedged asset position consists of a long position in an asset and an equal-sized short position in the currency of the asset's denomination. Average returns and return volatilities are computed using daily data and reported as annualized percentages.

type of constraint, often referred to as the elastic net, can be implemented by adding both the $\mathcal{L}^1$ and $\mathcal{L}^2$ penalty terms to the mean-variance objective function. For details on the exact specification of the elastic net SAS international portfolio optimization, see the Online Appendix.

The portfolio optimization rules introduced above do not address the practical concern of the allowed minimal/maximal currency exposure in an international portfolio. In practice, many institutional investors, including pension funds, face restrictions on the permissible level of foreign currency exposure in their portfolios. Such restrictions are usually driven by the regulators. Moreover, constraining currency weights can, from another angle, also be seen as a form of shrinkage. The detailed formulation of incorporating the net currency exposure constraints into the proposed model is provided in the Appendix.

4. Empirical analysis

The goal of this section is to empirically examine the out-of-sample performance of the regularized international portfolio optimization methods. We compare the proposed joint portfolio optimization approach with its currency overlay counterpart and the $1/N$ benchmark.

4.1. Data and summary statistics

The analysis is based on data ranging from January 1999 to December 2019. The data was obtained from Refinitiv Datastream and consists of 21 equity broad market indices. We consider ten major developed-market currencies, the US dollar, euro, Japanese yen, Swiss franc, British pound, Canadian dollar, New Zealand dollar, Singapore dollar, Norwegian krone, and Australian dollar. All of the assets are denominated in one of these currencies and there is at least one asset denominated in every currency throughout the analysis. The data series are available at a daily frequency, and the sample period ranges from January 1999, when the euro was introduced to the global financial markets, to December 2019. Table 1 shows the average annual returns and annual volatilities for the ten currency zones. The returns are converted to the US dollar and are benchmarked against their fully hedged counterparts.

The average annual returns range from 5% to 13%, with Japan exhibiting the lowest average return of 5.51%, while New Zealand demonstrates the highest average return of 12.79% over the sample period. The table further shows that fully hedging the assets yields a lower average return for all markets except for the Eurozone, the United Kingdom, and Japan. Furthermore, the volatilities of the fully hedged returns are lower than those of their unhedged counterparts for all markets except for Japan. This shows, as already previously observed in the international portfolio optimization literature, that full hedging has the potential to improve the risk-return spectrum by lowering the portfolio volatility.

Table 2 reports the full-sample correlations of currency returns (Panel A) and the cross-correlations between currencies and stocks (Panel B). Panel A shows that all currency returns are positively correlated. Moreover, markets that are geographically closely related, such as Australia and New Zealand, or Switzerland and the Eurozone, demonstrate particularly high correlations in currency returns. The cross-correlations between currencies and equities in Panel B are generally positive, with only the Japanese yen, Swiss franc, and the US dollar being negatively correlated to any of the stock markets. This is similar to the findings of Campbell et al. (2010), who found that the Swiss franc and the Japanese yen move against global equity markets and are consequently attractive for risk-minimizing international investors based in the US. Panel B further shows that these three currencies exhibit very low correlations with international stock markets. This further increases their value to international investors by acting as a potential source of additional diversification. Consequently, these currencies are known as the so-called safe-haven currencies.

Note that the currency correlations, as shown in Panel A of Table 2, are considered in both separate and joint optimization problems, while the cross-correlations between currencies and assets are only considered in the joint optimization. The reason for this is that the currency overlay approach only takes into account the correlation between the return of the asset portfolio and the currency returns. Hence, the currency overlay approach can be regarded as a form of dimensionality reduction. From Panel B, it is

Table 2
Cross-country return correlations: 1999 to 2019.

	AUD	CAD	CHF	EUR	GBP	JPY	NZD	USD	NOK	SGD
	Panel A: Currencies									
AU	1.00									
CA	0.56	1.00								
CH	0.33	0.26	1.00							
EU	0.47	0.43	0.69	1.00						
UK	0.34	0.49	0.41	0.55	1.00					
JP	0.12	0.27	0.43	0.31	0.26	1.00				
NZ	0.68	0.41	0.37	0.45	0.32	0.09	1.00			
US	0.17	0.51	0.37	0.38	0.53	0.57	0.15	1.00		
NO	0.48	0.49	0.48	0.65	0.48	0.19	0.34	0.27	1.00	
SG	0.49	0.60	0.56	0.58	0.57	0.60	0.42	0.75	0.48	1.00
	Panel B: Stocks									
AU	0.76	0.39	0.05	0.21	0.12	-0.19	0.55	-0.10	0.34	0.20
CA	0.60	0.67	-0.02	0.17	0.22	-0.10	0.40	0.03	0.36	0.24
CH	0.47	0.34	0.31	0.28	0.23	0.00	0.42	0.09	0.33	0.27
EU	0.52	0.37	0.18	0.45	0.25	-0.17	0.45	-0.04	0.44	0.22
UK	0.55	0.50	0.11	0.31	0.44	-0.09	0.43	0.10	0.44	0.32
JP	0.43	0.41	0.17	0.21	0.30	0.27	0.31	0.28	0.23	0.39
NZ	0.60	0.34	0.12	0.25	0.17	-0.10	0.76	-0.02	0.28	0.23
US	0.51	0.56	0.08	0.25	0.31	0.05	0.41	0.33	0.32	0.45
NO	0.52	0.37	0.04	0.23	0.19	-0.21	0.31	-0.13	0.56	0.11
SG	0.59	0.41	0.03	0.13	0.18	-0.07	0.46	0.02	0.25	0.35

This table presents the cross-country correlations between currency and stock returns. In Panel A, each cell reports the correlation of currency returns between the row and column currencies. Panel B presents the same for stock returns, where each cell represents the correlation between an equity market (row) and a currency (column). The correlations are calculated as averages across all possible base currencies.

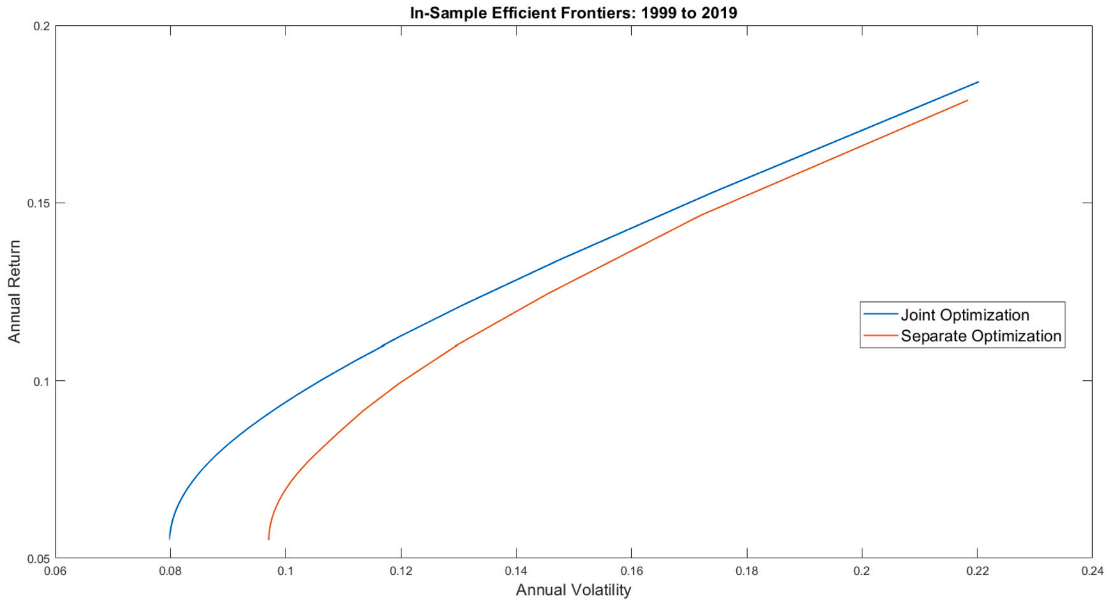

Fig. 1. This figure illustrates the in-sample efficient frontiers for the dataset of 21 equity indices and 10 currencies spanning the period from 1999 to 2019. The curves represent the results obtained by varying the risk aversion parameter γ in both steps of the separate optimization as well as in the joint optimization.

clear that for the period covered by the dataset, the cross-correlations are different from zero and take on values of up to 0.76 for correlations between local currencies and local markets. Therefore, the joint optimization approach proposed in this paper has the potential to improve portfolio performance by taking into account this additional information.

The potential benefits of the joint optimization are further depicted in Fig. 1, which depicts the in-sample efficient frontiers for both the separate and the joint international portfolio optimization problems. The frontiers were derived by solving the joint and the separate optimization problems from Section 3 for a range of risk-aversion parameters.

As this analysis is carried out in-sample (i.e., the same data is used for the estimation and the performance evaluation), it is clear that the joint optimization approach, by definition, has to provide improved results compared to the separate optimization approach. The separate approach is clearly suboptimal as it does not take into consideration the dependencies between the individual assets and the currencies but only between the whole asset portfolio and the currencies. The two frontiers in Fig. 1 would be equal in the case that assets and currencies are uncorrelated. In this case, the joint and the separate optimization would take into consideration the same amount of information.

4.2. Out-of-sample analysis

In this section, the out-of-sample performance of the regularized international portfolio optimization approaches is investigated. These approaches encompass the sparse and/or stable methods described in Section 3, as well as the standard mean-variance and $1/N$ portfolio benchmarks. In addition to the (currency unconstrained) SAS portfolios, we include the same optimization methods with added constraints on the allowed net currency exposures. Such constraints are unique to the international portfolio optimization problem and are driven by a practical consideration arising from potential regulatory restrictions. For these portfolio rules, the maximal absolute net currency exposure is bounded by 30% of the total portfolio value. The list of out-of-sample analyzed regularized international portfolio optimization rules is:

- **MV**: The classical mean-variance portfolio (without regularization).
- **NC1**: The sparse portfolio constructed by imposing an $\mathcal{L}^1$-norm constraint.
- **NC2**: The stable portfolio constructed by imposing an $\mathcal{L}^2$-norm constraint.
- **NC1-CC**: The sparse and stable portfolio constructed by imposing an $\mathcal{L}^1$-norm constraint and shrinkage of the covariance matrix towards a given constant correlation matrix.
- **NC1-NC2**: The sparse and stable portfolio constructed by constraining both the $\mathcal{L}^1$- and the $\mathcal{L}^2$-norm.
- **RULE-CB**: A given optimization rule (RULE) as specified above with a constraint that the absolute value of the net currency exposure to any currency does not exceed 30%, i.e., a currency bound (CB).
- **1/N**: The fully hedged $1/N$ portfolio (i.e., the equally-weighted portfolio).

For each portfolio optimization rule, a separate optimization (i.e., the currency overlay) and a corresponding joint optimization over assets and currencies in a single optimization step are carried out. Our objective is to assess the out-of-sample performance of the regularized joint optimization approach compared to the regularized separate and non-regularized optimization methods. While the joint optimization achieves in-sample optimality, it is important to examine whether this advantage carries over to out-of-sample scenarios. The joint optimization approach requires estimating dependencies between single assets and currencies. Such granular estimation is alleviated in the separate optimization that can consequently be seen as an approach that inherently reduces dimensionality. To gain insights into the optimal approach for multi-currency portfolios, it is crucial to analyze the out-of-sample performance of both optimization approaches. Additionally, we present the fully hedged $1/N$ portfolio as a benchmark, which serves as a widely explored model-free reference in the international asset allocation literature.

Following the existing literature, we adopt the US dollar as the base currency throughout the analysis and accordingly report all the results in USD. We assume that an investor rebalances her portfolio with a monthly frequency, while daily data is used for estimation. We further assume that an investor only hedges her currency exposure one period ahead, with forward contracts entered at time t expiring at time $t+1$. At expiry, the payoffs from the currency hedging are reinvested in the assets. The optimization employs a rolling estimation window of two years (i.e., 500 daily observations). The covariance matrix is then computed according to the given respective portfolio optimization rule outlined above, and the Bayes-Stein estimator is utilized for the shrinkage of the mean, see Jorion (1986).

The risk-aversion parameter is set to $\gamma = 3$, which is a choice suitable for most portfolio allocation decisions, see for example Fabozzi et al. (2007). The other model parameters, λ_1^a, λ_1^c, λ_2^a, and λ_2^c, corresponding to the $\mathcal{L}^1$ and $\mathcal{L}^2$ penalization parameters for assets and currencies, are determined via cross-validation. Here, we follow the approach of DeMiguel et al. (2009a), where leave-one-out cross-validation is carried out on the data sample of each estimation window. Using adaptive grid search, the combination of values that achieves the lowest cross-validation portfolio variance is then chosen as the model parameter set. This set is used to estimate the optimal asset and currency weights on the given rolling window.

Table 3 details the results of the out-of-sample analysis. The table provides a comprehensive analysis of the portfolio rules, considering different optimization approaches and hedging strategies. All results are reported net of transaction costs, with the transaction costs being computed as a proportion of the traded assets and currencies based on the bid-ask spread. We assume transaction costs of twenty basis points for the assets and two basis points relative to the notional of each entered currency forward contract. The Online Appendix provides the derivation of the multi-currency portfolio returns adjusted for transaction costs. For each portfolio, we report the annualized Sharpe Ratio, denoted by SR, the annualized Sortino Ratio, denoted by SOR, the annualized certainty equivalent return, denoted by CEV, the annualized volatility, the maximum drawdown, denoted by MD, and the turnover of the assets and currencies, denoted by $Turnover^a$ and $Turnover^c$, respectively. The asset turnover reported in Table 3 is an average monthly turnover, as widely used in the existing literature, see, for example, Li (2015) or DeMiguel et al. (2009a). We also provide the turnover associated with hedging the currency exposure. The foreign currency exposure is hedged via the forward contracts with delivery at their maturity. Thereby, an investor at every time t makes a new decision on the amount of entered currency forward contracts (i.e., possibly long or short). Thus, the turnover for the currencies is calculated as the average over all absolute currency weights $|\phi_{c,t}|$, for $c = 2, \ldots, M+1$.

In Table 3, we present the results in two panels: Panel A and Panel B. In Panel A, we explore the case where optimization is performed over both assets and currencies. This panel includes separate optimization and joint optimization approaches, allowing us to compare the performance of these strategies. The results highlight the impact of considering currency exposures in the portfolio

Table 3
Out-of-sample portfolio performance.

Panel A:	SR		SOR		CEV		Volatility		MD		$Turnover^a$		$Turnover^c$	
	Sep.	Joint	Sep.	Joint	Sep.	Joint	Sep.	Joint	Sep.	Joint	Sep.	Joint	Sep.	Joint
MV	0.3481	0.4625	0.3939	0.5314	0.0217	0.0322	0.2167	0.2701	0.6359	0.6808	1.8663	2.1036	3.1820	5.1781
MV-CB	0.2488	0.4119	0.2779	0.4919	0.0091	0.0390	0.1924	0.2005	0.5884	0.5778	1.8425	2.0160	2.3346	3.5834
NC1	0.4586	0.5035	0.5474	0.5702	0.0489	0.0580	0.1091	0.1425	0.5166	0.5110	0.5151	0.4708	1.3771	0.8574
NC1-CB	0.4496	0.5048	0.5193	0.6020	0.0473	0.0552	0.1046	0.1170	0.5281	0.4476	0.4873	0.4893	1.4731	1.2008
NC2	0.5107	0.5600	0.6142	0.6620	0.0582	0.0641	0.1343	0.1296	0.5942	0.4785	0.4291	0.3548	0.4499	1.2842
NC2-CB	0.5198	0.5171	0.6319	0.6246	0.0583	0.0573	0.1255	0.1212	0.5770	0.4673	0.4301	0.3847	0.8074	1.6956
NC1-CC	0.3699	0.4542	0.4307	0.5092	0.0394	0.0510	0.1150	0.1443	0.5437	0.5307	0.6007	0.4589	1.1480	0.7989
NC1-CC-CB	0.3880	0.5168	0.4397	0.6125	0.0413	0.0569	0.1110	0.1182	0.5628	0.4541	0.5506	0.4526	1.2373	1.1435
NC1-NC2	0.4962	0.6172	0.6305	0.7646	0.0488	0.0659	0.0883	0.1082	0.4504	0.4389	0.3472	0.4542	1.6709	1.0144
NC1-NC2-CB	0.5532	0.6810	0.6882	0.8314	0.0538	0.0715	0.0881	0.1046	0.4612	0.4392	0.3619	0.4058	1.6223	1.4011
1/N	0.3413		0.4071		0.0357		0.1302		0.5504		0.0369		0.9482	
Panel B:	SR		SOR		CEV		Volatility		MD		$Turnover^a$		$Turnover^c$	
	Full	Zero	Full	Zero	Full	Zero	Full	Zero	Full	Zero	Full	Zero	Full	Zero
MV	0.2768	0.4461	0.3189	0.5465	0.0238	0.0480	0.1537	0.1839	0.5630	0.6556	1.8248	1.8235	2.4480	0
NC1	0.4845	0.5756	0.5519	0.7036	0.0509	0.0681	0.1042	0.1411	0.5306	0.6095	0.5137	0.5077	0.9804	0
NC2	0.3058	0.4469	0.3492	0.5400	0.0323	0.0494	0.0971	0.1294	0.5235	0.5935	0.4348	0.4282	1.0022	0
NC1-CC	0.3515	0.4769	0.3966	0.5488	0.0373	0.0531	0.1175	0.1271	0.5828	0.5816	0.6023	0.5987	0.8897	0
NC1-NC2	0.4813	0.5671	0.5472	0.6872	0.0494	0.0658	0.0975	0.1341	0.4962	0.5714	0.3459	0.3437	0.9539	0

This table presents the performance measures for the international portfolio optimization rules introduced in this paper. The portfolio rules are abbreviated as defined in Section 4.2. The table is divided into two panels. Panel A showcases the results when optimization is performed over both assets and currencies, either for separate or joint optimization. Panel B focuses on the scenario where optimization is conducted solely over the assets (as the first step of the currency overlay), followed by either full hedging or zero hedging, which remains constant across all times and currencies. The dataset spans from January 1999 to December 2019, with the first portfolio formed in January 2001. It includes 21 equity indices from ten different countries, along with the corresponding spot and forward currency rates. All measures are computed out-of-sample and net of transaction costs.

construction process. Panel B, on the other hand, focuses on the scenario where optimization is conducted solely over the assets, representing the first step of the currency overlay. In this panel, we examine the performance of portfolios with two distinct hedging strategies: full hedging and zero hedging. These strategies maintain a constant hedging level across all times and currencies. By presenting the results in two panels, we provide a comprehensive view of the performance measures for each portfolio rule, considering different optimization and hedging approaches. The insights gained from this analysis contribute to our understanding of the out-of-sample effectiveness of the proposed portfolio optimization methods.

The results in Table 3 demonstrate that the portfolios derived with the joint optimization consistently achieve higher Sharpe ratios compared to their separate optimization counterparts (Panel A) and the constant hedging benchmarks (Panel B). On average, the joint optimization rules achieve 23.3% higher Sharpe ratios compared to their separate optimization counterparts. The comparison between portfolio rules with and without currency bounds shows that SAS portfolio rules that include bounds on the absolute currency exposures (i.e., NC1-CC-CB and NC1-NC2-CB) outperform the rules without such constraints (i.e., NC1-CC and NC1-NC2) in terms of Sharpe ratios. The other portfolio rules generally exhibit slightly inferior performance upon the introduction of currency bounds, which remains consistent for both separate and joint optimization approaches. Additionally, the regularized portfolio rules, along with the constant hedging approaches, outperform the fully hedged $1/N$ and MV portfolio benchmarks. For instance, the jointly optimized NC1-NC2-CB portfolio attains a Sharpe ratio that is double that of the $1/N$ portfolio, highlighting the substantial improvement achieved by the joint optimization combined with the additional currency exposure constraints.

To statistically compare the risk-adjusted performance of the presented portfolio rules, we utilize a robust test for differences in Sharpe ratios based on Ledoit and Wolf (2008). The results reveal that the regularized portfolio rules consistently exhibit lower p-values compared to the non-regularized MV benchmark, providing stronger evidence for rejecting the null hypothesis of equal Sharpe ratios. Notably, the jointly optimized NC1-NC2-CB strategy achieves a notable p-value of 5.75%. Further details can be found in the Appendix.

The results in terms of Sortino ratios and certainty equivalents are similar to the ones based on the Sharpe ratio. Most notably, portfolios constructed with the joint optimization approach consistently outperform their separate optimization counterparts. The norm-constrained portfolios achieve improved out-of-sample performance compared to the $1/N$ and MV benchmarks. This shows the importance of including regularization in the international portfolio optimization process. Moreover, the performance of the SAS portfolios is improved by adding additional bounds on the allowed absolute currency exposures. However, this is not the case for the rest of the portfolio rules.

Analyzing the annualized volatilities in Panel A of Table 3, we observe that the joint optimization generally exhibits higher volatility compared to the separate optimization rules. However, as evidenced by the increased Sharpe and Sortino ratios, this increase in risk is compensated by higher average returns. Moreover, when considering the portfolio rules with additional currency exposure constraints in both joint and separate optimization approaches, we find that the introduction of currency bounds generally leads to a reduction in volatility. This highlights the effectiveness of incorporating additional shrinkage through the introduction of currency exposure bounds in managing portfolio risk. Turning our attention to Panel B, where the optimization is performed only over the assets with either full or zero hedging, we observe that the constant hedging benchmarks demonstrate that full hedging decreases volatility compared to zero hedging. This suggests that actively managing currency exposure can help reduce volatility in the portfolio. Overall, we find that the SAS separate optimization approach achieves the lowest volatility of 8.8%.

On the other hand, the joint optimization tends to decrease the portfolio maximum drawdowns compared to the separate optimization for the norm-constrained portfolio rules (Panel A). This indicates that while the joint optimization may potentially increase portfolio risk as measured by volatility, it tends to mitigate risk in terms of maximum drawdown. Moreover, the imposition of additional currency exposure constraints also contributes to a decrease in maximum drawdowns. In Panel B, we find that full hedging exhibits a lower maximum drawdown compared to zero hedging, indicating the risk-mitigating effect of full hedging. Furthermore, both the separate and joint optimization approaches achieve even smaller maximum drawdowns compared to constant hedging. Notably, the joint SAS portfolio rule emerges as the strategy with the smallest maximum drawdown, showcasing its effectiveness in managing downside risk.

Last, the joint optimization approach tends to reduce asset and currency turnover in norm-constrained portfolios, with only minor exceptions. In contrast, the non-regularized MV benchmark and the NC1-NC2 portfolio rule exhibit the opposite trend. Notably, the model-free $1/N$ portfolio, by its construction, achieves the lowest turnover, with an average monthly asset turnover of 3.7%. It is worth highlighting that the asset turnovers between the RULE and RULE-CB separate optimization strategies (Panel A), as well as between full and zero hedging (Panel B), slightly differ due to the reinvestment of payoffs from the currency hedging back into the asset portfolio. These reinvestments create small discrepancies in asset turnovers over the course of the backtest.

All in all, this analysis shows that regularization and joint optimization improve the out-of-sample performance of internationally diversified portfolios. Such an improvement is achieved in terms of increased risk-adjusted returns (i.e., Sharpe and Sortino ratios) as well as an increased certainty equivalent, and reduced maximum drawdown net of transaction costs. The best-performing model is the SAS portfolio with the additionally imposed currency bounds (i.e., NC1-NC2-CB). Compared to the $1/N$ benchmark, this portfolio rule exhibits doubled Sharpe and Sortino ratios as well as a doubled certainty equivalent, while reducing the portfolio maximum drawdown by 20%.

Fig. 2 illustrates the computed optimal currency exposure in EUR for the SAS portfolio rules, namely NC1-NC2 and NC1-NC2-CB, under both separate and joint optimization. The figure highlights the optimal currency exposure differences between the separate and joint optimization approaches. Despite these differences, both methods effectively enhance the out-of-sample portfolio performance, as demonstrated in Table 3. The optimal exposures are generally stable over time, with the joint optimization approach tending to provide increased stability compared to the separate optimization approach. Notably, the portfolio rule incorporating currency

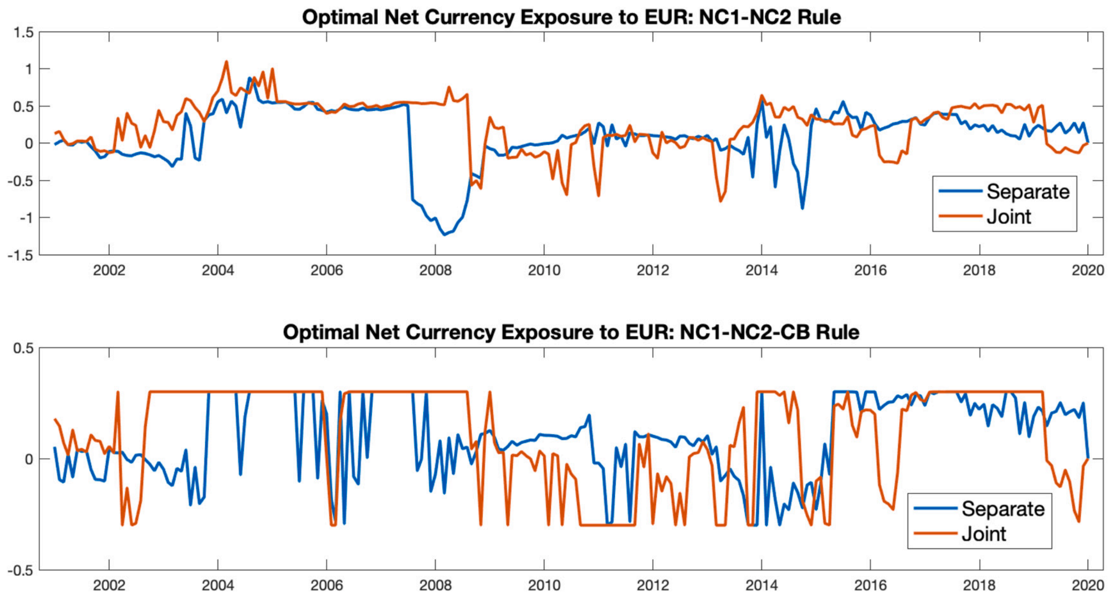

Fig. 2. This figure presents the computed optimal currency exposure in EUR for two sparse and stable portfolio rules: the NC1-NC2 portfolio (top) and the NC1-NC2-CB portfolio (bottom). The results are shown for both the separate and joint optimization approaches.

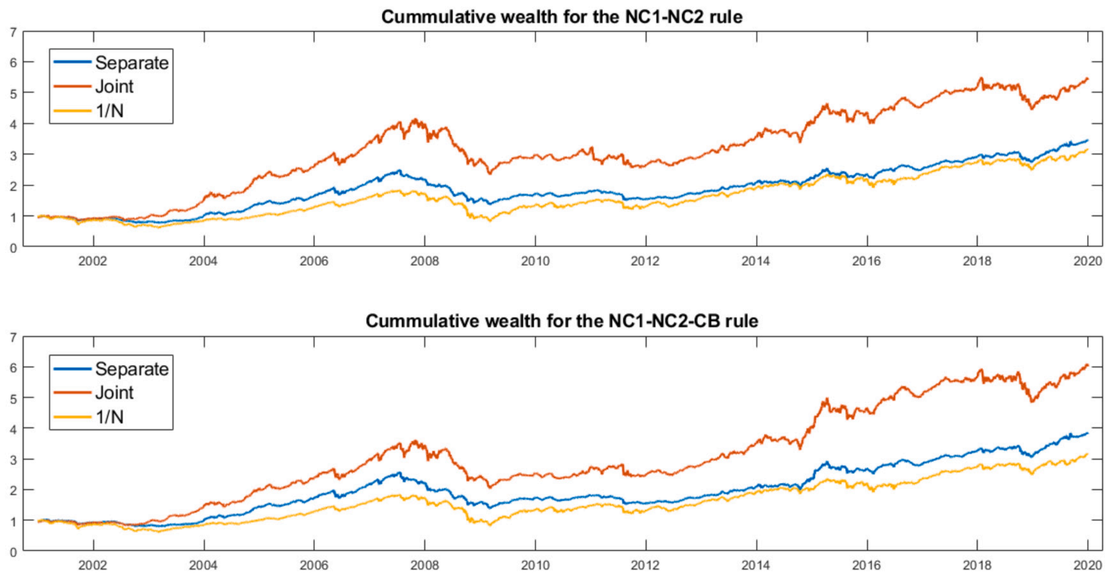

Fig. 3. This figure portrays the cumulative wealth evolution for two sparse and stable portfolio rules: the NC1-NC2 portfolio (top) and the NC1-NC2-CB portfolio (bottom), with an initial wealth of \$1. The reported values are adjusted for transaction costs.

bounds demonstrates enhanced stability over time, as the currency bounds can be seen as an additional form of shrinkage. This improved stability is one of the main drivers behind the out-of-sample outperformance of the NC1-NC2-CB rule. Additionally, the optimal exposures for other currencies are presented in the Online Appendix.

Fig. 3 illustrates the evolution of the cumulative wealth for the two SAS portfolio rules, NC1-NC2 and NC1-NC2-CB. For each rule, the separate and joint optimizations are depicted, and the fully hedged $1/N$ portfolio is shown as the benchmark. One can observe a consistent outperformance of the joint optimization approach compared to its separate counterpart as well as the $1/N$ portfolio. This outperformance is here achieved in terms of cumulative wealth and is observed over the whole backtest period. Furthermore, one can observe that the constraints on the currency exposure improve the out-of-sample performance in terms of cumulative wealth for the joint SAS portfolio rule. The other portfolio rules covered in this paper show similar cumulative wealth performance patterns, with the joint optimization consistently outperforming the separate optimization as well as the $1/N$ portfolio net of transaction costs.

5. Conclusions

In this paper, we develop a novel framework for international asset allocation that incorporates regularization techniques to determine asset and currency weights in a sparse and stable manner. Additionally, we conduct an in-depth analysis investigating the out-of-sample benefits of a joint optimization approach compared to the prevalent practice of employing currency overlay strategies for managing portfolio currency risk. Through this comprehensive analysis, we contribute to the literature on multi-currency asset allocation.

To the best of our knowledge, this is the first paper presenting a methodology for sparse and stable multi-currency asset allocation. Similar approaches, aimed at reducing parameter uncertainty in the mean-variance framework, have been proposed for the allocation of assets denominated in a single currency. We extend these approaches by proposing a framework where asset and currency weights are determined in a sparse and stable manner. Furthermore, we present an extensive empirical analysis of the out-of-sample performance of the international asset allocation problem and the corresponding currency risk management policy. We demonstrate that the proposed sparse and stable joint optimization outperforms the equivalent approaches based on separate (i.e., currency overlay) optimization as well as the equally-weighted and the non-regularized global portfolios net of transaction costs. On average, the joint optimization approaches attain 23.3% higher out-of-sample realized Sharpe ratios compared to their separate optimization counterparts. These findings challenge the prevailing practice of employing currency overlay strategies and highlight the potential for substantial additional gains in risk-adjusted returns through the joint optimization of assets and currencies.

This paper establishes a compelling case for the adoption of the proposed sparse and stable international asset allocation framework, as it provides a superior alternative to currency overlay strategies. By integrating assets and currencies within a unified optimization framework, our approach accounts for the interdependencies between them and offers improved risk management and enhanced portfolio performance. While such an approach has a larger potential to suffer from parameter uncertainty and over-fitting, we employ several regularization approaches that tackle these problems by inducing sparsity and stability of portfolio weights. We show that the out-of-sample outperformance of the joint optimization is driven by regularization that induces a bias-variance trade-off leading to improved portfolio performance. Moreover, constraining the allowed net currency exposures improves the performance of sparse and stable portfolios. Such constraint can also be interpreted as an additional form of shrinkage.

This paper allows for possible extensions of the proposed international asset optimization approach into several research directions. For example, one could generalize the introduced methodology based on the mean-variance objective to the sparse and stable tail-based (e.g., expected shortfall) risk measures. Moreover, time-series models capturing the empirical stylized facts present in financial data (e.g., volatility clustering, asymmetry, and excess kurtosis) could also be employed for modeling the asset and currency return dynamics. The performance of such models can then be benchmarked upon our model. The current framework could also be extended to the multi-period setting. Thereby, the asset and currency rebalancing would be considered within a longer time horizon. Such a model would require the development of a method based on stochastic programming that captures the multi-period asset and currency allocation decision dynamics. On the empirical side, the performance of bond portfolios could be investigated as well. Given the generally low correlations between currency and bond returns, the common practice of institutional investors is to fully hedge bond portfolios. One could explore if this strategy could potentially be improved within a joint optimization format. Furthermore, one could also study the performance of multi-currency portfolios exposed to emerging market currencies in combination with the already explored developed markets.

Appendix A

A.1. Embedding currency exposure constraints

Consider the notation from Section 3. Take the net currency exposure ψ_c and assume that

$$V_c^l \leq \psi_c = w_c - \phi_c \leq V_c^u \qquad \text{for} \quad c = 2, \dots, M+1, \tag{17}$$

where V_c^l and V_c^u are constants indicating the lower and the upper allowed exposure to currency c, expressed as a fraction of total portfolio value.[8] In what follows, we assume $V^l = V_2^l = \dots = V_{M+1}^l$ and $V^u = V_2^u = \dots = V_{M+1}^u$, indicating that the lower and the upper bound are equal among all foreign currencies. It is straightforward to relax this assumption and allow for constraints that vary among different currencies. Using Eqs. (5) and (17), the upper bound on the net currency exposure can be expressed as

$$\boldsymbol{w}_t - \boldsymbol{\phi}_t \leq V^u \mathbf{1}_M \iff \boldsymbol{C}^T \boldsymbol{x}_t - \boldsymbol{\phi}_t \leq V^u \mathbf{1}_M \iff \begin{pmatrix} \boldsymbol{C}^T & -\boldsymbol{I}_M \end{pmatrix} \begin{pmatrix} \boldsymbol{x}_t \\ \boldsymbol{\phi}_t \end{pmatrix} \leq V^u \mathbf{1}_M. \tag{18}$$

The lower bound can be expressed equivalently.

For the separate optimization problem, the net currency exposure constraints are implemented by imposing an upper and a lower bound on the currency weights in the second optimization (i.e., currency overlay) step. The first optimization step (i.e., asset allocation) remains unchanged. Let $\boldsymbol{\phi}_t^-$ and $\boldsymbol{\phi}_t^+$ denote the lower and upper bound on the currency weight, $\boldsymbol{\phi}_t^- \leq \boldsymbol{\phi}_t \leq \boldsymbol{\phi}_t^+$. Using Eq. (18), the lower and the upper bounds of the currency weights are given by

[8] For a specific choice of $V_c^l = 0$ and $V_c^u = w_c$, currency c is not allowed to be over or under hedged (i.e., hedging between 0% and 100% of the exposure is allowed). For $V_c^l < 0$ and $V_c^u > w_c$, over- and under-hedging, respectively, are permitted.

Table 4
Sharpe Ratio Difference Test: p-values.

	Separate vs. $1/N$	Joint vs. $1/N$
MV	97.29%	66.03%
MV-CB	73.75%	77.01%
NC1	56.37%	46.96%
NC1-B	55.48%	38.69%
NC2	27.58%	27.51%
NC2-CB	24.89%	34.66%
NC1-CC	83.14%	61.44%
NC1-CC-CB	74.90%	37.58%
NC1-NC2	33.40%	11.99%
NC1-NC2-CB	21.04%	5.75%

This table presents the outcomes of the robust hypothesis test examining the difference in Sharpe ratios between two investment strategies as proposed by Ledoit and Wolf (2008). The null hypothesis assumes that the Sharpe ratios of the two strategies are identical. A p-value of $x\%$ implies that the null hypothesis can be rejected at the $x\%$ significance level.

$$\boldsymbol{\phi}_t^+ \mid \boldsymbol{x}_t^* = \boldsymbol{C}^T \boldsymbol{x}_t^* - V^l \mathbf{1}_M \quad \text{and} \quad \boldsymbol{\phi}_t^- \mid \boldsymbol{x}_t^* = \boldsymbol{C}^T \boldsymbol{x}_t^* - V^u \mathbf{1}_M, \tag{19}$$

highlighting that the currency bounds depend on the asset weights chosen in the first optimization step. A general SAS separate international portfolio optimization setting, as described before, can easily be augmented for currency exposure constraints by adding the constraint from Eq. (19) in the second optimization step.

Next, we investigate the joint optimization problem. The representation in Eq. (18) shows that an implicit currency exposure, attained through optimizing over $\boldsymbol{x}_t$, affects the possible values currency weights $\boldsymbol{\phi}_t$ are allowed to take. In the joint optimization problem, these two vectors are determined in a combined fashion and influence each other. Hence, the currency exposure constraints in the joint optimization have to be enforced with a linear inequality constraint specification.

Consider the definition of $\boldsymbol{\theta}_t$ and the constraints imposed by the upper and lower bounds, as obtained in Eq. (18). Then, a joint international portfolio optimization problem can be augmented for currency exposure constraints via an additional inequality constraint given by

$$\boldsymbol{Q}\boldsymbol{\theta}_t \leq \boldsymbol{D},$$

with

$$\boldsymbol{Q} = \begin{pmatrix} \boldsymbol{C}^T & -\boldsymbol{I}_M \\ -\boldsymbol{C}^T & \boldsymbol{I}_M \end{pmatrix} \quad \text{and} \quad \boldsymbol{D} = \begin{pmatrix} V^u \mathbf{1}_M \\ -V^l \mathbf{1}_M \end{pmatrix},$$

where $\boldsymbol{Q} \in \mathbb{R}^{2M\times(N+M)}$ and $\boldsymbol{D} \in \mathbb{R}^{2M\times 1}$.

A.2. Sharpe ratio difference test

To statistically compare the risk-adjusted performance of portfolio rules described in Table 3, we conduct a test for differences in Sharpe ratios between each model-based portfolio rule and the $1/N$ benchmark. We employ a robust performance hypothesis testing approach proposed by Ledoit and Wolf (2008) that assumes the null hypothesis of equal Sharpe ratios for the two strategies. The results presented in Table 4 indicate that many portfolio rules do not exhibit statistically significant differences in their Sharpe ratios compared to the $1/N$ benchmark. This is expected, as the Sharpe ratio estimates are known to be noisy. However, the jointly optimized NC1-NC2-CB strategy achieves a noteworthy p-value of 5.75%. Moreover, the regularized strategies consistently demonstrate lower p-values compared to the non-regularized MV benchmark. These findings highlight the effectiveness of the SAS international portfolios in enhancing risk-adjusted portfolio performance as measured by the Sharpe ratio.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jimonfin.2023.102949.

References

Boudoukh, J., Richardson, M., Thapar, A., Wang, F., 2019. Optimal currency hedging for international equity portfolios. Financ. Anal. J. 75 (4), 65–83.
Britten-Jones, M., 1999. The sampling error in estimates of mean-variance efficient portfolio weights. J. Finance 54 (2), 655–671.
Brodie, J., Daubechies, I., De Mol, C., Giannone, D., Loris, I., 2009. Sparse and stable Markowitz portfolios. Proc. Natl. Acad. Sci. 106 (30), 12267–12272.
Brown, C., Dark, J., Zhang, W., 2012. Dynamic currency hedging for international stock portfolios. Rev. Futures Mark. 20, 419–455.
Campbell, J.Y., Medeiros, K. Serfaty-De, Viceira, L.M., 2010. Global currency hedging. J. Finance 65 (1), 87–121.
Caporin, M., Jimenez-Martin, J.-A., González Serrano, L., 2014. Currency hedging strategies in strategic benchmarks and the global and Euro sovereign financial crises. J. Int. Financ. Mark. Inst. Money 31, 159–177.
Chatsanga, N., Parkes, A.J., 2017. International portfolio optimisation with integrated currency overlay costs and constraints. Expert Syst. Appl. 83, 333–349.
Cho, J.-B., Min, H.-G., McDonald, J.A., 2020. Volatility and dynamic currency hedging. J. Int. Financ. Mark. Inst. Money 64, 101163.

Daubechies, I., Defrise, M., De Mol, C., 2004. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. Commun. Pure Appl. Math. 57 (11), 1413–1457. A Journal Issued by the Courant Institute of Mathematical Sciences.
De Roon, F.A., Nijman, T.E., Werker, B.J., 2003. Currency hedging for international stock portfolios: the usefulness of mean–variance analysis. J. Bank. Finance 27 (2), 327–349.
DeMiguel, V., Garlappi, L., Nogales, F.J., Uppal, R., 2009a. A generalized approach to portfolio optimization: improving performance by constraining portfolio norms. Manag. Sci. 55 (5), 798–812.
DeMiguel, V., Garlappi, L., Uppal, R., 2009b. Optimal versus naive diversification: how inefficient is the 1/N portfolio strategy? Rev. Financ. Stud. 22 (5), 1915–1953.
Efron, B., Hastie, T., Johnstone, I., Least Angle Regression, R. Tibshirani, 2004. Ann. Stat. 32 (2), 407–499.
Fabozzi, F.J., Kolm, P.N., Pachamanova, D.A., Focardi, S.M., 2007. Robust Portfolio Optimization and Management. John Wiley & Sons.
Friedman, J., Hastie, T., Tibshirani, R., 2001. The Elements of Statistical Learning. Springer Series in Statistics, vol. 1. New York.
Garlappi, L., Uppal, R., Wang, T., 2007. Portfolio selection with parameter and model uncertainty: a multi-prior approach. Rev. Financ. Stud. 20 (1), 41–81.
Glen, J., Jorion, P., 1993. Currency hedging for international portfolios. J. Finance 48 (5), 1865–1886.
Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer Science & Business Media.
Huang, X., Wang, X., 2021. International portfolio optimization based on uncertainty theory. J. Math. Prog. Oper. Res. 70 (2), 225–249.
Jagannathan, R., Ma, T., 2003. Risk reduction in large portfolios: why imposing the wrong constraints helps. J. Finance 58 (4), 1651–1683.
Jorion, P., 1986. Bayes-Stein estimation for portfolio analysis. J. Financ. Quant. Anal. 21 (3), 279–292.
Jorion, P., 1994. Mean/variance analysis of currency overlays. Financ. Anal. J. 50 (3), 48–56.
Kim, S.F., Chance, D.M., 2018. An empirical analysis of corporate currency risk management policies and practices. Pac.-Basin Finance J. 47, 109–128.
Konno, H., Li, J., 2000. Applications of the integrated approach to international portfolio optimization. Asia-Pac. Financ. Mark. 7 (2), 121–144.
Ledoit, O., Wolf, M., 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. J. Empir. Finance 10 (5), 603–621.
Ledoit, O., Wolf, M., 2004. Honey, I shrunk the sample covariance matrix. J. Portf. Manag. 30 (4), 110–119.
Ledoit, O., Wolf, M., 2008. Robust performance hypothesis testing with the Sharpe ratio. J. Empir. Finance 15 (5), 850–859.
Li, J., 2015. Sparse and stable portfolio selection with parameter uncertainty. J. Bus. Econ. Stat. 33 (3), 381–392.
Markowitz, H.M., 1952. Portfolio selection. J. Finance 7 (1), 77–91.
Mhiri, M., Prigent, J.-L., 2010. International portfolio optimization with higher moments. Int. J. Econ. Finance 2 (5), 157–169.
Michaud, R.O., 1989. The Markowitz optimization enigma: is 'optimized' optimal? Financ. Anal. J. 45 (1), 31–42.
Opie, W., Riddiough, S.J., 2020. Global currency hedging with common risk factors. J. Financ. Econ. 136 (3), 780–805.
Perold, A.F., Schulman, E.C., 1988. The free lunch in currency hedging: implications for investment policy and performance standards. Financ. Anal. J. 44 (3), 45–50.
Polak, P., Ulrych, U., 2021. Dynamic currency hedging with non-Gaussianity and ambiguity. In: Swiss Finance Institute Research Paper No. 21-60.
Roncalli, T., 2013. Introduction to Risk Parity and Budgeting. CRC Press.
Solnik, B.H., 1974. An equilibrium model of the international capital market. J. Econ. Theory 8 (4), 500–524.
Suzuki, K.-I., Konno, H., Morjiri, M., 1997. An international portfolio optimization model hedged with forward currency contracts. Financ. Eng. Jpn. Markets 4 (3), 275–286.
Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. J. R. Stat. Soc., Ser. B, Methodol. 58 (1), 267–288.
Topaloglou, N., Vladimirou, H., Zenios, S.A., 2002. CVaR models with selective hedging for international asset allocation. J. Bank. Finance 26 (7), 1535–1561.
Ulrych, U., Vasiljević, N., 2020. Global currency hedging with ambiguity. In: Swiss Finance Institute Research Paper No. 20-73.
Van Binsbergen, J.H., Brandt, M.W., Koijen, R.S., 2008. Optimal decentralized investment management. J. Finance 63 (4), 1849–1895.