Contents lists available at ScienceDirect

Heliyon

www.elsevier.com/locate/heliyon

Research article Algorithmic portfolio tilting to harvest higher moment gains

ABSTRACT

Kris Boudt^{a,b,c,*}, Dries Cornilly^d, Frederiek Van Holle^e, Joeri Willems^e

^a Universiteit Gent, Sint-Pietersplein 5, 9000 Gent, Belgium

^b Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium

^c Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 Amsterdam, Netherlands

^d SYZ Asset Management, Rue du Commerce 3, 1204 Genève, Switzerland

^e Degroof Petercam Asset Management, Guimardstraat 18, 1040 Brussel, Belgium

ARTICLE INFO

Keywords: Mean-variance-skewness-kurtosis Non-normality Portfolio allocation Tilting Statistics Finance Banking Econometrics Operations management Business Economics Information science Industry

0. Introduction

Many of the popular risk-based portfolios are optimized according to an optimality criterion that does not take into account higher order portfolio return moments (see e.g. Lee (2011)). Typical examples from the literature on risk-based asset allocation include the equally-weighted and equal-risk-contribution portfolios (Maillard et al. (2010)). A further example is the maximum diversification portfolio of Choueifaty and Coignard (2008), which maximizes the ratio between the weighted average standard deviation of the portfolio components and the portfolio standard deviation. These portfolios may still be efficient for investors with mean-variance preferences, as discussed in Ardia and Boudt (2015). They are, however, likely to be inefficient for investors whose preferences depend on the higher order moments of the portfolio return distribution. Several solutions have been proposed in terms of portfolio allocation methods that take the higher order moments into account (see e.g. Jondeau et al. (2007) and Boudt et al. (2012)). In this paper, we study the problem of improving portfolios that are inefficient from a mean-variance-skewness-kurtosis (MVSK) perspective, where skewness and kurtosis are defined as the non-standardized third and fourth central portfolio moments. We do so by tilting such inefficient portfolios to an allocation that is more MVSK efficient, while keeping the good properties of the initial portfolio weights.¹ The proposed methodology is inspired by the shortage minimization framework of Briec et al. (2007).

Many financial portfolios are not mean-variance-skewness-kurtosis efficient. We recommend tilting these

portfolios in a direction that increases their estimated mean and third central moment and decreases their

variance and fourth central moment. The advantages of tilting come at the cost of deviation from the initial

optimality criterion. In this paper, we show the usefulness of portfolio tilting applied to the equally-weighted,

equal-risk-contribution and maximum diversification portfolios in a UCITS-compliant asset allocation setting.

We expect MVSK portfolio tilting to be especially useful in asset allocation. We therefore illustrate the methodology by tilting three common risk-based asset allocation portfolios in a UCITS-compliant setting, where the investment universe consists of 12 broad asset class indices. We find that the tilted risk-based portfolios achieve superior out-of-

E-mail address: kris.boudt@ugent.be (K. Boudt).

https://doi.org/10.1016/j.heliyon.2020.e03516

Received 2 December 2019; Received in revised form 28 January 2020; Accepted 27 February 2020

2405-8440/© 2020 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).







^{*} Corresponding author at: Universiteit Gent, Sint-Pietersplein 5, 9000 Gent, Belgium.

¹ This approach of optimizing non-standardized moments is consistent with the results of Scott and Horvath (1980) which support preferences for third and fourth central moments instead of the standardized skewness and excess kurtosis. We use tilting in the sense of over- or underweighting assets compared to benchmark portfolios (see e.g. Darsinos and Satchell (2004)). We refer the reader to Stutzer (2000, 2003), Haley and McGee (2006), Haley and Whiteman (2008) and Haley and McGee (2011) for an alternative definition of tilting in a portfolio context, namely modifying the probability return distribution.

Table 1	
Impact of MVSK tilting on portfolio criteria.	

Deterioration of reference objective	$h_{\mathrm{ref}}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) \leq h_{\mathrm{ref}}(\mathbf{w}_0 \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) + \kappa$
Increase in mean return	$\mathbf{w}'\boldsymbol{\mu} \ge \mathbf{w}_0'\boldsymbol{\mu} + \boldsymbol{\delta}_{\boldsymbol{\mu}}$
Decrease in variance	$\mathbf{w}' \mathbf{\Sigma} \mathbf{w} \le \mathbf{w}_0' \mathbf{\Sigma} \mathbf{w}_0 - \delta_{\Sigma}$
Increase in skewness	$\mathbf{w}' \mathbf{\Phi}(\mathbf{w} \otimes \mathbf{w}) \geq \mathbf{w}_0' \mathbf{\Phi}(\mathbf{w}_0 \otimes \mathbf{w}_0) + \delta_{\Phi}$
Decrease in kurtosis	$\mathbf{w}'\Psi(\mathbf{w}\otimes\mathbf{w}\otimes\mathbf{w}) \leq \mathbf{w}_0'\Psi(\mathbf{w}_0\otimes\mathbf{w}_0\otimes\mathbf{w}_0) - \delta_{\Psi}$

Note: The parameters δ_{μ} , δ_{Σ} , δ_{Φ} , δ_{Ψ} are positive when the MVSK tilted portfolio has preferential moments compared to the benchmark, at the cost of an increase of κ in h_{ref} .

sample moments while still showing the characteristics of the initial allocation criterion.

The remainder of the article is organized as follows. Section 1 introduces our notation. Section 2 presents the proposed methodology for portfolio tilting. Section 3 documents the usefulness of the method in a UCITS-compliant asset allocation framework. We end the paper with a conclusion. The supplementary appendix provides an example using the R package **mvskPortfolios**.

1. Notation

We consider the decision problem of optimizing a portfolio invested in N assets with weight vector $\mathbf{w} \equiv (w_1, \dots, w_N)'$. The optimality of the portfolio weights is evaluated with respect to a reference portfolio \mathbf{w}_0 and the first four portfolio return moments. For those portfolio return moments, accurate estimators for the corresponding comoment matrices have been proposed by Martellini and Ziemann (2010) and Boudt et al. (2020a,b), among others. The interpretation and reliable estimation of comoment matrices greater than four is a subject for further research. It is also an open research question whether the integration of portfolio moments greater than four in portfolio optimization improves the out-of-sample welfare of the investor (see e.g. Jondeau and Rockinger (2006)). Investors interested in portfolio optimization using all moments can directly optimize the portfolio return distribution, as in Haley and McGee (2006). This requires to re-estimate the portfolio return distribution for every candidate portfolio vector w. The computational advantage of our MVSK tilting portfolio is that the optimization does not require such re-estimation.

1.1. Asset returns, their comoments and the portfolio moments

The asset return over the investment horizon is denoted by $\mathbf{r} \equiv (r_1, \dots, r_N)'$. Its mean, covariance, coskewness and cokurtosis matrices are

$$\mu \equiv \mathbb{E}[\mathbf{r}],$$

$$\Sigma \equiv \mathbb{E}\left[(\mathbf{r} - \mu)(\mathbf{r} - \mu)'\right],$$

$$\Phi \equiv \mathbb{E}\left[(\mathbf{r} - \mu)(\mathbf{r} - \mu)' \otimes (\mathbf{r} - \mu)'\right],$$

$$\Psi \equiv \mathbb{E}\left[(\mathbf{r} - \mu)(\mathbf{r} - \mu)' \otimes (\mathbf{r} - \mu)' \otimes (\mathbf{r} - \mu)'\right],$$
(1)

where \otimes stands for the Kronecker product. These comoments determine the portfolio higher order moments. Define $m_q \equiv \mathbb{E} \left[(\mathbf{w'r} - \mathbf{w'}\mu)^q \right]$ as the *q*-th central portfolio moment, implicitly depending on the weights **w**. We have that $m_2 \equiv \mathbf{w'} \Sigma \mathbf{w}$, $m_3 \equiv \mathbf{w'} \Phi(\mathbf{w} \otimes \mathbf{w})$ and $m_4 \equiv \mathbf{w'} \Psi(\mathbf{w} \otimes \mathbf{w} \otimes \mathbf{w})$. We denote the portfolio tracking error volatility by the volatility of the difference in portfolio composition:

$$TE_{vol}(\mathbf{w}) = \sqrt{(\mathbf{w} - \mathbf{w}_0)^{\prime} \Sigma(\mathbf{w} - \mathbf{w}_0)}.$$
(2)

1.2. Optimality of the reference portfolio

We are agnostic about the construction of the reference portfolio \mathbf{w}_0 . It can be expert-based, algorithmic or a combination of both. The proposed MVSK tilting requires that the investor elicit an optimality criterion $h_{\text{ref}}(\cdot)$ for which \mathbf{w}_0 is the optimal solution. This criterion is a

function of the portfolio weights w, the asset return moments μ , Σ , Φ and Ψ and possibly other parameters stacked into the vector ζ :

$$\mathbf{w}_0 \equiv \operatorname{argmin}_{\mathbf{w}\in\mathcal{C}} h_{\operatorname{ref}}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}), \tag{3}$$

where *C* is the feasible set. It is always possible to find a criterion for which this reference portfolio is optimal. The simplest criterion is the tracking error volatility in (2), which is obviously minimized when $\mathbf{w} = \mathbf{w}_0$.

In the application, we consider three examples of h_{ref} -functions. The first equals the Herfindahl index of the portfolio weights:

$$h_{\rm EW}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) = \sum_{i=1}^{N} w_i^2, \tag{4}$$

which is minimized by the equally-weighted portfolio. DeMiguel et al. (2009) refer to this portfolio as the 1/N allocation. The portfolio uses a naive approach to diversification as it ignores the heterogeneity in the component return distributions when optimizing the weights. The second objective function is the sum of squared deviations of the *N* percentage risk contributions with respect to 1/N:

$$h_{\text{ERC}}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) = \sum_{i=1}^{N} \left(\% R C_i - \frac{1}{N} \right)^2,$$
(5)

with $\Re RC_i \equiv \frac{w_i [\Sigma w]_i}{w' \Sigma w}$ the percentage volatility contribution of asset *i*. The solution to this optimization is the equal-risk contribution portfolio (Maillard et al. (2010)). The third objective that we consider is the negative of the diversification ratio (Choueifaty and Coignard (2008)):

$$h_{\rm DR}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) = -\frac{\mathbf{w}' \sqrt{\rm diag(\boldsymbol{\Sigma})}}{\sqrt{\mathbf{w}' \boldsymbol{\Sigma} \mathbf{w}}},\tag{6}$$

where diag(Σ) = ($\Sigma_{11}, ..., \Sigma_{pp}$)' is the $N \times 1$ vector containing the diagonal elements of Σ .

Note that the framework has a large scope and also includes the mean-variance efficient portfolios of Markowitz (1952) and the minimum modified Value-at-Risk and Expected Shortfall portfolios of Boudt et al. (2008, 2012), among others. Moreover, the function $h_{\rm ref}$ is not necessarily the one used to determine w_0 . For example, it is possible to construct the maximum diversification portfolio with weights w_0 by minimizing $h_{\rm DR}$, but use $h_{\rm ref} = {\rm TE}_{\rm vol}$ when tilting the portfolio.

2. MVSK portfolio tilting

We assume that the investor is willing to sacrifice a margin κ of the value of the reference objective function in order to improve the portfolio performance in terms of an increase in the mean and skewness and a decrease in the variance and the kurtosis. We impose these improvements through inequality constraints, as can be seen in Table 1, where the minimum levels of improvements in the mean, variance, skewness and kurtosis are δ_{μ} , δ_{Σ} , δ_{Φ} , and δ_{Ψ} , respectively.

Our approach to simultaneously look for improvements in expected return, variance, skewness and kurtosis is compatible with the accepted direction of preferences for the first four moments (see e.g. Scott and Horvath (1980)) and is closely related to the utility-based framework (Jondeau and Rockinger (2006), Harvey et al. (2010)). In the utilitybased approach to non-Gaussian preferences, the expected utility is often approximated by a Taylor expansion, or more generally, a function

$$\mathcal{U}_{\alpha}(\boldsymbol{w}) = \alpha_1 \boldsymbol{w}' \boldsymbol{\mu} - \alpha_2 \boldsymbol{w}' \boldsymbol{\Sigma} \boldsymbol{w} + \alpha_3 \boldsymbol{w}' \boldsymbol{\Phi}(\boldsymbol{w} \otimes \boldsymbol{w}) - \alpha_4 \boldsymbol{w}' \boldsymbol{\Psi}(\boldsymbol{w} \otimes \boldsymbol{w} \otimes \boldsymbol{w}).$$
(7)

The vector α contains positive values and governs the relative importance of the higher moments with respect to the expected utility of the investor. Such an objective function is generally non-concave and hence it is impossible to guarantee a global maximum. In this regard, Briec et al. (2007) prove a primal-dual relationship between the primal, shortage-based approach followed in this paper and the dual approach of directly maximizing a non-concave expected utility approximation.

2.1. Choice of δ parameters

n

S

In the current formulation we have at least three unknowns: the tilted portfolio weights **w**, the minimum improvement levels $\delta_{(\cdot)}$ and the sacrifice parameter κ . In principle, we can assume that the investor will elicit these. In practice, however, this is not realistic and a data-driven calibration is required. To achieve this, we simplify the problem by assuming that the value of κ is fixed and that the $\delta_{(\cdot)}$'s are functions of a single δ .

The investor wishes to maximize this δ , thereby achieving more preferable portfolio moments at a maximum cost of κ in the reference objective function. The general MVSK portfolio tilting problem is then:

$$\begin{aligned} \underset{\delta \in \mathbb{R}, \mathbf{w} \in C}{\operatorname{tabulk}} \quad \delta \\ \underset{\delta \in \mathbb{R}, \mathbf{w} \in C}{\operatorname{tabulk}} \quad & h_{\operatorname{ref}}(\mathbf{w} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) \leq h_{\operatorname{ref}}(\mathbf{w}_0 \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\Phi}, \boldsymbol{\Psi}, \boldsymbol{\zeta}) + \kappa, \\ \mathbf{w}' \boldsymbol{\mu} \geq \mathbf{w}'_0 \boldsymbol{\mu} + \delta_{\boldsymbol{\mu}}, \\ \mathbf{w}' \boldsymbol{\Sigma} \mathbf{w} \leq \mathbf{w}'_0 \boldsymbol{\Sigma} \mathbf{w}_0 - \delta_{\boldsymbol{\Sigma}}, \\ \mathbf{w}' \boldsymbol{\Phi}(\mathbf{w} \otimes \mathbf{w}) \geq \mathbf{w}'_0 \boldsymbol{\Phi}(\mathbf{w}_0 \otimes \mathbf{w}_0) + \delta_{\boldsymbol{\Phi}}, \\ \mathbf{w}' \boldsymbol{\Psi}(\mathbf{w} \otimes \mathbf{w} \otimes \mathbf{w}) \leq \mathbf{w}'_0 \boldsymbol{\Psi}(\mathbf{w}_0 \otimes \mathbf{w}_0) - \delta_{\boldsymbol{\Psi}}, \end{aligned}$$
(8)

where δ_{μ} , δ_{Σ} , δ_{Φ} , and δ_{Ψ} are all increasing functions of δ . We denote this as $(\delta_{\mu}, \delta_{\Sigma}, \delta_{\Phi}, \delta_{\Psi})' = g(\delta)$. The threshold κ determines the maximum deterioration in the reference function. In the extreme case of setting $\kappa = \infty$, we have the MVSK optimization setup of Briec et al. (2007).

2.2. Choice of κ parameter

. .

The optimization in (8) yields the optimal value of δ for a given value of κ , and increasing κ always increases the value of δ . To gauge this trade-off between low values of κ being more in line with the initial portfolio objective and higher values of δ yielding preferable moments, we recommend visualizing δ on a grid of κ values that ranges from zero to a pre-set upper bound $\overline{\kappa}$. Based on this, an investor or portfolio manager can select a value of κ that limits the decrease in the reference objective for a significant gain in moments. This choice will depend on the level, slope and curvature of the curve.

In the context of portfolio tilting, it is also natural to set κ at the highest possible value that respects an upper bound on the portfolio tracking error volatility TE_{vol}. This means that we look for the highest improvement in the MVSK objectives (as measured by δ) for which the tilted portfolio is sufficiently close to the initial portfolio. The MVSK portfolio tiling problem is then:

$$\begin{split} \max_{\delta \in \mathbb{R}, w \in \mathcal{C}} & \delta \\ \text{subject to} & \text{TE}_{\text{vol}}(\mathbf{w}) \leq \tau, \\ & \mathbf{w}' \, \mu \geq \mathbf{w}'_0 \mu + \delta_\mu, \\ & \mathbf{w}' \, \Sigma \mathbf{w} \leq \mathbf{w}'_0 \Sigma \mathbf{w}_0 - \delta_\Sigma, \\ & \mathbf{w}' \, \Phi(\mathbf{w} \otimes \mathbf{w}) \geq \mathbf{w}'_0 \Phi(\mathbf{w}_0 \otimes \mathbf{w}_0) + \delta_\Phi, \\ & \mathbf{w}' \, \Psi(\mathbf{w} \otimes \mathbf{w} \otimes \mathbf{w}) \leq \mathbf{w}'_0 \Psi(\mathbf{w}_0 \otimes \mathbf{w}_0 \otimes \mathbf{w}_0) - \delta_\Psi, \end{split}$$
(9)

where $(\delta_{\mu}, \delta_{\Sigma}, \delta_{\Phi}, \delta_{\Psi})' = g(\delta)$ and τ determines the maximum tracking error volatility of the tilted portfolio as compared to the reference portfolio.

2.3. Shrinkage estimation of comoments

In order to implement the proposed framework, reliable estimates of the comoments μ , Σ , Φ and Ψ are required. One possibility is to use parametric conditions linking the mean estimated return to the estimated comoments, as in Ardia and Boudt (2015). Another possibility is to combine the sample estimator and a parametric structured matrix using the approach of shrinkage estimation, as recommended by Ledoit and Wolf (2003), Martellini and Ziemann (2010) and Boudt et al. (2020a). In general, the shrinkage estimator for the central moments θ is given by

$$\hat{\theta} = (1 - \hat{\alpha}^{\star})\hat{\theta}_{\text{sample}} + \hat{\alpha}^{\star}\hat{\theta}_{\text{structured}},\tag{10}$$

with $\hat{\theta}_{sample}$ the sample moments and $\hat{\theta}_{structured}$ the moments under a structured estimation approach. Several choices of targets are possible. In this paper, we use the structured moments under the assumption that the marginals are independent, which is a popular choice in practice.

The value of \hat{a}^* is determined in a data-driven way by minimizing the mean squared estimation error:

$$\alpha^{\star} = \arg\min_{\alpha \in [0,1]} \mathbb{E}\left[\left\| (1-\alpha)\widehat{\theta}_{\text{sample}} + \alpha\widehat{\theta}_{\text{structured}} - \theta \right\|^2 \right], \tag{11}$$

for which a consistent estimator is derived in Ledoit and Wolf (2003) for the covariance matrix and in Boudt et al. (2020a) for the higher moments.

3. Illustration in asset allocation

In this application, we focus on variance-skewness-kurtosis (VSK) efficient portfolios, aiming to decrease the variance and fourth central moment while increasing the third central moment. We do not incorporate the estimated expected returns since their estimation precision is typically low compared to the precision in estimating the portfolio variance, skewness and kurtosis. We further use the following parameters for tilting the portfolio in order to improve the variance, skewness and kurtosis:

VSK :
$$(\delta_{\mu}, \delta_{\Sigma}, \delta_{\Phi}, \delta_{\Psi})$$

= $\delta \times (-\infty, \mathbf{w}_0' \Sigma \mathbf{w}_0, |\mathbf{w}_0' \Phi(\mathbf{w}_0 \otimes \mathbf{w}_0)|, \mathbf{w}_0' \Psi(\mathbf{w}_0 \otimes \mathbf{w}_0 \otimes \mathbf{w}_0)).$ (12)

The VSK tilting above is equivalent to MVSK tilting in case of homogeneous expected asset returns (i.e., $\mu = \mu v_N$, with v_N the *N*-dimensional vector of ones and μ the common expected return). Note that in (12) we let the function $g(\delta)$ depend in a specific way on the benchmark portfolio \mathbf{w}_0 . Alternative specifications may be more desirable depending on the application. In particular, when the minimum variance, maximum skewness or minimum kurtosis portfolio is used as benchmark portfolio, the calibration needs to be modified such that a deterioration of the second, third or fourth moment portfolio performance is allowed. This is possible by setting δ_{Σ} , δ_{Φ} , or δ_{Ψ} to $-\infty$, while setting $h_{\text{ref}}(\cdot)$ to the objective function corresponding to the minimum variance, maximum skewness or minimum kurtosis portfolio.

3.1. Data

The data consist of weekly returns of 12 broad scale investment indices over the period January 26, 2001 until March 15, 2019. All series are expressed in EUR. Table 2 provides an overview of the indices along with their mean, volatility, skewness and kurtosis over the full period. The annual volatilities lie between 3% for iBoxx EUR Corporate and 21% for MSCI Emerging Markets. The standardized skewness is generally

Table 2

Summary statistics of the weekly returns of the asset class indices.

		geom. mean	volatility	skewness	ex. kurt.
1	MSCI Europe	2.82	0.19	-0.71	7.46
2	MSCI USA	4.55	0.18	-0.34	3.05
3	MSCI Japan	1.63	0.19	-0.11	1.31
4	MSCI Emerging Markets	7.35	0.21	-0.15	4.97
5	EPRA Eurozone	5.16	0.19	-1.29	9.01
6	JPM EMU	4.64	0.04	-0.14	2.16
7	iBoxx EUR Corporate	4.50	0.03	-0.94	5.41
8	Bloomberg Barclays Global Aggregate Treasuries	3.17	0.06	0.57	4.01
9	JPM EMBI Global Diversified Composite	7.29	0.11	-0.12	4.87
10	Bloomberg Barclays Global High Yield	6.96	0.10	-0.21	4.70
11	Bloomberg Barclays US Corporates	4.37	0.10	0.29	1.46
12	Bloomberg Barclays World Govt Inflation Linked	4.34	0.07	0.11	1.39

Note: We report the annualized geometric mean (geom. mean, %), the annualized standard deviation (volatility), the standardized skewness and excess kurtosis (ex. kurt.) over the period January 26, 2001 until March 15, 2019.

negative, except for Barclays Global Aggregate Treasuries, US Corporates and World Govt Inflation Linked, for which it is slightly positive. All assets show heavy-tailed behavior and have a positive excess kurtosis.

3.2. Illustration of portfolio tilting

Figs. 1 and 2 illustrate the concept of VSK portfolio tilting for the last day of our sample, namely March 15, 2019. The reference portfolio is the maximum diversification portfolio (DR), and the moments are estimated as outlined inbelow.

The dots in Fig. 1 represent different randomly generated feasible portfolios and provide an idea of the set in which we are optimizing. We highlight the DR portfolio and associated tilted portfolios with $\kappa = -(0.01, 0.025, 0.05)h_{\text{DR}}(\mathbf{w}_0)$. When the margin κ increases, the tilted portfolios attain a higher skewness while decreasing in variance and kurtosis.

Fig. 2a shows the joint improvements in the moments measured by the objective value δ in (8). It is remarkable how much improvement is possible while only marginally decreasing the diversification ratio. A κ of 1% results in a relative improvement of at least 15%, whereas increasing κ to 5% yields a relative improvement of around 23%. This is also seen in the bottom row of Fig. 1, where the largest difference between subsequent portfolios occurs when $\kappa = 1\%$.

For the same portfolios, the effect on the weights is illustrated in Fig. 2b. The tilted portfolios lower the equity allocation (assets 1, 2, 3, and 4) and increase the bond component through Global High Yield (asset 10). As mentioned by a referee, bootstrapping can be used to evaluate the estimation uncertainty of the optimized weights or to robustify the weights with respect to estimation uncertainty. We refer the interested reader to Scherer (2002) for a review.

3.3. Out-of-sample results

We now turn to analyzing the out-of-sample effects of portfolio tilting on portfolio performance when the reference portfolio is obtained by minimizing $h_{\rm EW}$, $h_{\rm ERC}$ and $h_{\rm DR}$. For $h_{\rm EW}$ and $h_{\rm ERC}$, we consider deviations of $\kappa = 0.01, 0.05, 0.10$. For $h_{\rm DR}$, we take κ proportional to the diversification ratio of the maximum diversification portfolio; $\kappa = -(0.01, 0.025, 0.05)h_{\rm DR}(\mathbf{w}_0)$. In addition to these fixed κ , we automatically select the tilted portfolios with maximum annual tracking error volatilities of 0.5% and 1%. We adopt the industry-practice of rolling five-year estimation windows and estimate the moments using the shrinkage approach toward the target that assumes independent marginals. Before estimation, we winsorize the observations following the procedure in Boudt et al. (2008). The portfolios are long-only and fully invested with a maximum weight of 20% in a single asset, complying with UCITS regulations in a fund-of-funds setting. In practice, the indices can be replaced by funds covering different asset classes.

Fig. 3a illustrates that the covariance and cokurtosis shrinkage coefficients are stable over time and put most of the weight on the sample covariance and cokurtosis estimators. The coskewness shrinkage coefficient is more volatile and occasionally puts all weight on the target. This is in line with the findings of Martellini and Ziemann (2010) and Boudt et al. (2020a). The course of the diversification ratio over time is shown in Fig. 3b. Note the large decrease of the diversification ratio in the period 2014 to 2018 due to an increase in the average correlation between the indices from around 24% to 50%.

Table 3 reports summary statistics for the benchmark and tilted portfolios. In all three settings, we observe a decrease in volatility and fourth central moment (ψ) and an increase in third central moment (ϕ), which is exactly what VSK portfolio tilting aims to achieve. The effect is more pronounced when κ increases, but it comes at the cost of a slightly lower average return. As a measure for the trade-off, we follow the recommendation of Zakamouline and Koekebakker (2009) and use the skew-adjusted Sharpe ratio,

$$ASSR = SR\sqrt{1 + \frac{skewness}{3}SR},$$
(13)

where the Sharpe ratio (SR) and skewness are computed on the annual out-of-sample returns. The last column of Table 3 shows that tilting improves the ASSR in all cases.

It is further important to notice in Table 3 that portfolio tilting also improves the downside risk, since the maximum drawdown decreases compared to the benchmark portfolio. In the case of the DR portfolio, the maximum drawdown even goes from -15.6% to -11.1% when allowing a 5% lower diversification ratio. The effect on turnover is highly dependent on the benchmark portfolio. Due to the VSK portfolio being slightly more concentrated, the turnover generally increases with κ . However, turnover can be dealt with in practice by imposing an additional constraint in the portfolio optimization step.

Fig. 4a illustrates the minimum improvement in the variance, skewness and kurtosis as measured by δ in (8) when tilting the portfolios. It is clear that most of the improvement happens for low values of κ , such as $\kappa = 1\%$. Increasing κ is still beneficial, but its effect is dampened. For $\kappa = 1\%$, Fig. 4b shows the difference in drawdown over time between the maximum diversification portfolio and the tilted portfolio. The cause of the difference is found in the weights given in Fig. 4c and 4d. The clearest difference is that the tilted portfolio reduced the overall allocation to equities (assets 1, 2, 3 and 4) while increasing the bond component. Moreover, the tilted portfolio re-allocates part of the bond exposure. The main differences here are the weights of Global High Yield (asset 10) and US Corporates (asset 11). The DR portfolio does not invest in Global High Yield, while the tilted portfolio is invested in the periods 2006-2009 and 2016-2019. Since both are highly correlated and have about the same volatility, this is due to a difference in higher order comoments in relation to the other components of the portfolio.



Fig. 1. Effect of VSK tilting when the reference portfolio is the maximum diversification portfolio.

We verified that the choice of window length, shrinkage target and pre-processing does not alter the conclusion presented in this section. Robustness checks are available from the authors upon request.

4. Conclusion

This paper presents a mean-variance-skewness kurtosis (MVSK) portfolio tilting framework designed to modify portfolio weights in a direction that improves the first four moments of the portfolio return distribution. Under the MVSK portfolio tilting framework, the initial portfolio is assumed to be optimal with respect to an initial objective function. However, it is possible to jointly improve the moments at the cost of only a slight deterioration in this initial objective function. Hence, portfolio tilting is a trade-off between the deterioration in portfolio construction objective on the one hand, and the improvements in the portfolio return moments on the other. The empirical application shows the improvements in out-of-sample performance when tilting the equally-weighted, equal-risk-contribution and maximum diversification portfolios in a UCITS-compliant asset allocation setting. These performance gains are low-hanging-fruit for all investors. A practical policy



Fig. 2. Effect of portfolio tilting on portfolio moments and weights.



Fig. 3. Portfolio statistics over time.

 Table 3

 Portfolio summary: variance-skewness-kurtosis tilting.

		geom. mean	volatility	ϕ	Ψ	MD	Turn.	ASSR
DR		4.506	0.052	-0.222	0.228	-15.605	86.713	0.671
	$\kappa = 0.01$	4.318	0.049	-0.197	0.190	-13.603	85.650	0.694
	$\kappa = 0.025$	4.326	0.048	-0.171	0.166	-12.392	92.869	0.718
	$\kappa = 0.05$	4.473	0.047	-0.147	0.145	-11.053	99.652	0.761
	$\tau = 0.005$	4.462	0.049	-0.200	0.195	-13.834	81.661	0.694
	$\tau = 0.01$	4.547	0.047	-0.157	0.153	-11.627	90.682	0.755
ERC		4.838	0.053	-0.248	0.264	-14.106	42.527	0.691
	$\kappa = 0.01$	4.666	0.050	-0.161	0.176	-11.846	63.846	0.738
	$\kappa = 0.025$	4.611	0.049	-0.130	0.148	-10.795	89.080	0.755
	$\kappa = 0.05$	4.563	0.048	-0.119	0.134	-10.045	125.653	0.762
	$\tau = 0.005$	4.706	0.050	-0.186	0.194	-12.401	82.228	0.711
	$\tau = 0.01$	4.688	0.048	-0.144	0.153	-10.768	119.563	0.749
EW		5.447	0.084	-1.305	2.007	-26.562	55.653	0.509
	$\kappa = 0.01$	5.175	0.066	-0.559	0.713	-19.264	50.015	0.590
	$\kappa = 0.025$	4.824	0.059	-0.333	0.402	-16.088	95.241	0.617
	$\kappa = 0.05$	4.670	0.050	-0.197	0.211	-12.570	73.487	0.705
	$\tau = 0.005$	5.335	0.079	-1.092	1.566	-25.251	59.247	0.520
	$\tau = 0.01$	5.226	0.074	-0.902	1.203	-23.890	68.179	0.532

Note: We report the annualized geometric mean (geom. mean, %), the annualized standard deviation (volatility), the skewness (ϕ , 10⁻⁶) and kurtosis (ψ , 10⁻⁷), maximum drawdown (%), annual turnover (Turn.) and skew-adjusted Sharpe ratio (ASSR) based on annual returns. The DR, ERC and EW portfolios are reported, along with the performance of their tilted versions for $\kappa = 0.01, 0.05, 0.10$ and tilting with a maximum annual tracking error volatility τ of 0.5% and 1%.



Fig. 4. Portfolio statistics over time.

recommendation is therefore that regulators in the field of investor protection should promote knowledge of how higher order moments affect financial portfolio outcomes and how they can be optimized. In order to facilitate this process, we have released an open source implementation in the R package **mvskPortfolios**.

Declarations

Author contribution statement

K. Boudt, D. Cornilly, F. Van Holle and J. Willems: Conceived and designed the experiments; Performed the experiments; Analyzed and interpreted the data; Contributed reagents, materials, analysis tools or data; Wrote the paper.

Funding statement

The authors gratefully acknowledge support from the Research Foundation - Flanders (FWO research grant G023815N and PhD fellowship 1114117N).

Competing interest statement

The authors declare no conflict of interest.

Additional information

The analysis can be reproduced using the code available at https://github.com/cdries/mvskPortfolios.

Supplementary content related to this article has been published online at https://doi.org/10.1016/j.heliyon.2020.e03516.

References

- Ardia, D., Boudt, K., 2015. Implied expected returns and the choice of a mean-variance efficient portfolio proxy. J. Portf. Manag. 41 (4), 68–81.
- Boudt, K., Carl, P., Peterson, B.G., 2012. Asset allocation with conditional value-at-risk budgets. J. Risk 15 (3), 39–68.
- Boudt, K., Cornilly, D., Verdonck, T., 2020a. A coskewness shrinkage approach for estimating the skewness of linear combinations of random variables. J. Financ. Econom. 18 (1), 1–23.
- Boudt, K., Cornilly, D., Verdonck, T., 2020b. Nearest comment estimation with unobserved factors. J. Econom. https://doi.org/10.1016/j.jeconom.2019.12.009. forthcoming.
- Boudt, K., Peterson, B.G., Croux, C., 2008. Estimation and decomposition of downside risk for portfolios with non-normal returns. J. Risk 11 (2), 79–103.
- Briec, W., Kerstens, K., Jokung, O., 2007. Mean-variance-skewness portfolio performance gauging: a general shortage function and dual approach. Manag. Sci. 53 (1), 135–149.
- Choueifaty, Y., Coignard, Y., 2008. Toward maximum diversification. J. Portf. Manag. 35 (1), 40–51.
- Darsinos, T., Satchell, S., 2004. Measuring style tilting and decomposing style risk. J. Asset Manag. 5 (1), 64–71.
- DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: how inefficient is the 1/n portfolio strategy? Rev. Financ. Stud. 22 (5), 1915–1953.
- Haley, M.R., McGee, M.K., 2006. Tilting safety first and the Sharpe portfolio. Finance Res. Lett. 3 (3), 173–180.
- Haley, M.R., McGee, M.K., 2011. "KLICing" there and back again: portfolio selection using the empirical likelihood divergence and Hellinger distance. J. Empir. Finance 18 (2), 341–352.
- Haley, M.R., Whiteman, C.H., 2008. Generalized safety first and a new twist on portfolio performance. Econom. Rev. 27 (4–6), 457–483.
- Harvey, C.R., Liechty, J.C., Liechty, M.W., Müller, P., 2010. Portfolio selection with higher moments. Quant. Finance 10 (5), 469–485.

- Jondeau, E., Poon, S.-H., Rockinger, M., 2007. Financial Modeling Under Non-Gaussian Distributions. Springer Science & Business Media.
- Jondeau, E., Rockinger, M., 2006. Optimal portfolio allocation under higher moments. Eur. Financ. Manag. 12 (1), 29–55.
- 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.
- Lee, W., 2011. Risk-based asset allocation: a new answer to an old question? J. Portf. Manag. 37 (4), 11–28.

Maillard, S., Roncalli, T., Teïletche, J., 2010. The properties of equally weighted risk contribution portfolios. J. Portf. Manag. 36 (4), 60–70.

Markowitz, H., 1952. Portfolio selection. J. Finance 7 (1), 77-91.

- Martellini, L., Ziemann, V., 2010. Improved estimates of higher-order comoments and implications for portfolio selection. Rev. Financ. Stud. 23 (4), 1467–1502.
- Scherer, B., 2002. Portfolio resampling: review and critique. Financ. Anal. J. 58 (6), 98–109.
- Scott, R.C., Horvath, P.A., 1980. On the direction of preference for moments of higher order than the variance. J. Finance 35 (4), 915–919.
- Stutzer, M., 2000. A portfolio performance index. Financ. Anal. J. 56 (3), 52-61.
- Stutzer, M., 2003. Portfolio choice with endogenous utility: a large deviations approach. J. Econom. 116, 365–386.
- Zakamouline, V., Koekebakker, S., 2009. Portfolio performance evaluation with generalized Sharpe ratios: beyond the mean and variance. J. Bank. Finance 33 (7), 1242–1254.