Shrinkage Estimation in Risk Parity Portfolios

Nabil Alkafri* Christoph Frey†

September 24, 2021

Abstract

We investigate the impact of shrinkage estimation techniques for the moments of asset returns on risk-parity portfolios. In contrast to mean-variance portfolios, the risk contributions of individual assets in risk-parity portfolios are fixed a priori. This additional restriction is commonly found to stabilize empirical portfolio weights in time. We show that the marginal risk-budget for each portfolio asset indeed serves as a natural shrinkage target and hence provide a new perspective on risk-parity portfolios. In an extensive empirical application, we compare and combine the various shrinkage strategies to popular risk-based approaches from the literature. We find that while using shrinkage estimators in risk-parity portfolios enhances out-of-sample performance based on various criteria, traditional covariance shrinkage estimators dominate all other strategies in high-dimensional settings.

Keywords: Estimation risk, regularization, asset allocation, portfolio optimization, variance-covariance shrinkage

JEL Classifications: C13, C52, C58, C61, G11

*WHU – Otto Beisheim School of Management, Vallendar, Germany. Email: nabil.alkafri@whu.edu

†Erasmus University Rotterdam. Email: christoph.frey@gmail.com

1 Introduction

Risk parity (RP) is a cross-sectional portfolio allocation technique in which each asset contributes the same amount to the total portfolio volatility. While it says nothing about expected returns, an input most difficult to estimate precisely, it limits the risk contribution of individual assets. This appealing property makes such strategies especially popular among practitioners. Recently, RP has also gained some interest from academia due to its well documented performance in empirical analysis (Chaves et al., 2011). Some studies such as Clarke et al. (2013) and Lee (2014), among others, find that RP is superior to traditional strategies such as equal- and value-weighting as well as other weighting techniques that incorporate other moments of asset returns. Further, Maillard et al. (2010) provide theoretical arguments about RP and show for example, that a negative relationship exists between RP asset weights and RP asset betas, which coincides with the findings of Frazzini and Pedersen (2014) that low beta assets outperform high beta assets.

RP is often compared to the Markowitz (1952) minimum variance (MV) portfolio. The main input parameter to both the RP and MV strategies is the variance-covariance matrix of the asset returns: a property difficult to estimate when (i) the correlations among assets is high, i.e. in times of market crashes, or when (ii) the number of assets is high relative to the number of observations. A popular method that has been shown to improve the risk-return properties of mean-variance portfolios that contain estimation errors (Fabozzi et al., 2007) is called *shrinkage*. This technique tries to mitigate estimation error by simply averaging over various models. For example, Ledoit and Wolf (2004a) demonstrate how linearly shrinking the asset covariance matrix enhances the performance of such portfolios. More recently, Ledoit and Wolf (2017) also propose a non-linear shrinkage technique and demonstrate how this further increases Sharpe ratios of empirical portfolios. The latter shrinkage technique especially is an important breakthrough in risk-budgeting techniques such as RP for numerous reasons. Maillard et al. (2010) show that RP can be expressed as a minimum variance portfolio optimization with an additional , non-linaer, constraint for equalizing the asset risk budgets. Also, the authors show that the volatility of RP lies between the volatility of the minimum variance portfolio and that of the equally-weighted portfolio. This implies that RP results from a shrinkage between these two portfolio.

The equal weighting (EW) scheme is an heuristic approach to overcome estimation error by disregarding plug-in estimates for the asset moments altogether. However, disregarding the asset characteristics makes certain properties of EW portfolios, such as the level of diversification, highly sensitive to the underlying asset universe. For example, if the risk of the underlying assets vary significantly, a high risk concentration and a limited degree of diversification are attained, as the weights of the riskiest and the least risky assets in EW are identical. While EW strategies are employed on a broad scale by mutual and pension funds (Benartzi and Thaler, 2001; Windcliff and Boyle, 2004), it is viewed as an arbitrarily manner of deconcentration by equally spreading the initial wealth across all assets. As a matter of fact, the notion of deconcentration in EW has led to the development of RP portfolios, which aim to deconcentrate the portfolio from a risk perspective, where each asset contributes to the same amount of risk to the overall portfolio risk. Amenc et al. (2012) argue that such techniques are classified as *adhoc* allocation schemes, due to their reliance on the notion of deconcentration and the absence of a theoretical framework.

We investigate the impact of shrinkage estimation techniques for the moments of asset returns in risk parity portfolios. Given the number of assets in the asset universe, the risk contributions of individual assets in RP portfolios are fixed a priori. This additional information is commonly found to stabilize empirical portfolio weights over time. Accordingly, we assess whether an asset's risk budget in the RP optimization serves as a natural shrinkage target, and whether the parity element of RP is beneficial in improving portfolio performance relative to conventional linear and non-linear regularization techniques. We therefore bridge the gap between mean-variance and risk parity portfolios in a unique way. We test our results in asset allocation and high-dimensional portfolios and find that the marginal risk-budget for each portfolio asset indeed serves as a natural shrinkage target, in which RP weights are a results of shrinking MV weights towards EW weights. Hence, we provide a new perspective on risk parity portfolios. In higher dimensions however, we find that risk parity strategies can not compete with shrinkage solutions to MV portfolios; a clear indication that parity shrinkage in RP has its limits.

Using shrinkage techniques in RP portfolios is not new. This study is related to Ardia et al. (2017), who assess the impact of estimation errors in the asset variance-covariance matrix in various risk-based portfolios; among them the RP strategy. Using Monte Carlo simulation methods, they find that equal-risk-contribution and inverse-volatility weighted portfolio weights are relatively robust to covariance misspecification while the MV portfolio weights are highly sensitive to errors in both the estimated variances and correlations. While their study only operates in small dimensions with 5 to 30 assets, we investigate the impact of shrinkage estimation techniques also in high dimensions, where RP portfolios are typically not applied. Accordingly, it is important not only to study where the regularization benefits come from, but also to understand where the limitations arise.

The remainder of this study is organized as follows. Section 2 motivates our study by showing the impact of estimation error on RP and MV portfolios in a simulation study. Section 3 shows that the parity element in RP acts as a natural shrinkage device. Section 4 describes the empirical setup and briefly defines other conventional shrinkage and heuristic weighting techniques in the literature. RP, MV, and the strategies in section 4 are assessed in a multi asset allocation setting as well as in high-dimensional equity portfolios in sections 5 and 6. Section 7 concludes.

2 A motivational example

Academia has often focused on mean-variance portfolios (Markowitz, 1952), in which the portfolio weighting is governed by the risk preferences of the investor and the diversification benefits of individual assets. Despite the theoretical elegance and intuitiveness of the mean-variance framework, it did not fare well in practice and resulted in highly concentrated, underperforming, and simply "wrong" portfolios (Michaud, 1989). The observed underperformance of mean-variance optimized portfolios is due to the fact that the deterministic input parameters, namely the asset return vector and the asset covariance matrix, can only be estimated with severe estimation errors.[1] The only applied exception of mean-variance optimized portfolios is MV, which only attempts to minimize portfolio risk whilst neglecting asset returns. Numerous studies such as Chopra and Ziemba (1993), Clarke et al. (2006), and Chow et al. (2011) document the outperformance of MV in equity markets. However, MV has its shortcomings and a main issue is its high concentration in low volatility assets (Chan et al., 1999; Clarke et al., 2011; DeMiguel et al., 2009).

To provide some more empirical intuition for the reader, we follow Jobson and Korkie (1980) to illustrate the difference between RP and MV in terms of its sensitivity to plug-in estimates. We randomly select $N = 10, 50, 100$ assets from stocks that are listed in the NYSE, AMEX, and NASDAQ exchanges. In each scenario, we estimate the asset moments for the mean, $\hat{\mu}_t$, and the variance-covariance-matrix,

[1] See for example Best and Grauer (1991); Chopra and Ziemba (1993); Jobson and Korkie (1980); Michaud (1989)).

$\hat{\Sigma}_t$, and accordingly determine the "true" RP and MV portfolios and compare them with 10,000 estimated portfolios, which are computed based on 250 hypothetical returns that are simulated using a multivariate normal distribution with the "true" return vector ($\hat{\mu}_t$) and "true" covariance matrix ($\hat{\Sigma}_t$). The purpose of the simulation is to evaluate how close these simulated portfolios are compared to the true portfolio when increasing the number of assets held in the portfolio. The results of the simulation experiment are illustrated in Figure B.1.

In fact, our results are in line with Jobson and Korkie (1980) for MV. Simulated MV portfolios either yield a higher return with a higher variance or are suboptimal allocation compared with the true MV portfolio. Furthermore, the difference in MV between the true portfolio and estimations thereof increases with the number of assets. Meanwhile, the simulated RP portfolios are heavily concentrated around the true RP portfolio, which lies roughly in the center of all simulated portfolios. This observation holds true regardless of the number of assets considered. The results suggest that RP can be better estimated using plug-in estimates, as the simulated portfolios lie closely around the true portfolio. Also, simulated RP portfolios can not only be suboptimal compared to the true portfolio, but can also dominate the true RP portfolio in the mean-variance space. Accordingly, the results imply that the RP technique, by strictly defining the total risk contribution of each asset, induces a structure on the portfolio weights and hence can be deemed as a form of portfolio regularization.

3 Methodology

RP portfolios constitute a middle-ground between MV and EW. They maintain the notion of deconcentration as in EW whilst considering single and joint asset total risk contributions so that all assets contribute equally to the portfolio risk. Although MV equalizes asset risk contributions, it does so merely on a marginal basis. This implies that a minimal change in the weight of any asset in MV should result in the same ex-ante change in overall portfolio risk. However, the total risk contributions are generally not equal, which causes portfolio risk to be mainly concentrated in a few assets, foregoing the benefits of diversification. Maillard et al. (2010) show that risk parity can be denoted as the minimum variance portfolio with an additional risk budget constraint, formally it can be expressed as

$$y^* = \arg\min \sqrt{y'\Sigma y}, \quad s.t. \quad \begin{cases} \sum_{i=1}^{n} \ln y_i \geqslant c \\ y \geqslant 0 \end{cases}, \tag{1}$$

where the optimal portfolio weights $\omega^* = \{\omega_1^*, \ldots, \omega_N^*\}$, which sum up to 100%, are determined as $\omega_i^* = y_i^* / \sum_{j=1}^{N} y_j^* \quad \forall i = 1, \ldots, N$. γ is an arbitrarily chosen constant and Σ describes the asset variance-covariance matrix. Indeed, Maillard et al. (2010) show that, based on the risk budget determined by γ, the volatility of RP lies between those of MV and EW, i.e. $\sigma_{MV} < \sigma_{RP} < \sigma_{EW}$.[2]

An alternative to overcoming estimation error is to adjust the plug-in estimates objectively or subjectively, which can be based on the expectations of fund managers, asset pricing theories, or a combination of both (Black and Litterman, 1992). However, most prominently applied are various statistical methods that are less prone to sampling errors or fat-tail effects in the data. Among them are shrinkage techniques, which are a routine of averaging various estimators.[3] The notion underlying shrinkage is that portfolios that yield a stable performance regardless of market phases, namely portfolios that are more robust to estimation errors and model risk, are better suited for investors, especially risk-averse investors. This property is paramount in finance, given that financial returns are heavily entailed with extreme returns or fat-tails that typically have a significant impact on the portfolio optimization process and performance. Hence, shrinkage techniques reduce the impact of fat-tails and estimation errors, yielding more conservative portfolios that perform well regardless of market conditions (Fabozzi et al., 2007).

Shrinkage techniques date back to James and Stein (1961) and typically consist of: (i) An arbitrary estimator such as historical plug-in estimates, (ii) a structured shrinkage target, and (iii) a parameter that determines the shrinkage intensity between (i) and (ii). James and Stein (1961) show that the benefit of shrinkage techniques is a lower mean-squared error than the plug-in estimate for the cost of a typically biased shrinkage target. This is true for any shrinkage intensity parameter larger than zero.

[2]For more information on the derivation of these properties of RP as well as the constant γ, refer to Maillard et al. (2010), Appendix B, pp. 68-69.

[3]Other methods to mitigate the effect of estimation errors include Bayesian techniques (Barry, 1974), the Bayesian approach in combination with asset pricing models (Pástor, 2000), robust optimization models (Garlappi et al., 2007), Bayesian robust optimization techniques (Wang, 2005), robust estimation methods (DeMiguel and Nogales, 2009), and employing portfolio constraints (Jagannathan and Ma, 2003).

Linear shrinkage in risk parity

Provided our previous arguments, the RP strategy can be defined as a simple form of portfolio regularization. While shrinkage and factor approaches reduce dimensionality by reducing the number of estimable parameters, risk parity induces a structure on the portfolio weights by restricting the risk contribution of individual assets to the total portfolio variance.

Provided N assets, let $r = \{r_1, \ldots, r_N\}$ be a collection of the N asset returns and $\omega = \{\omega_1, \ldots, \omega_N\}$ the corresponding portfolio weight vector. The portfolio variance is defined by $\sigma_p^2(\omega) = \omega'\Sigma\omega = \sum_{i=1}^N \omega_i^2\sigma_i^2 + \sum_{i=1}^N \sum_{i\neq j} \omega_i\omega_j\sigma_{ij}$, where Σ is the $N \times N$ variance-covariance matrix of r with variances σ_i^2 and covariances σ_{ij}. The marginal risk contribution of a single asset i towards the total portfolio is $\partial\sigma_p(\omega)\big/\partial\omega_i$. Given that the total portfolio volatility is the sum of the individual risk contributions of all assets, $\sigma_p(\omega) = \sum_{i=1}^N \left(\omega_i \times \partial\sigma_p(\omega)\big/\partial\omega_i\right)$, the risk budgeting portfolio is defined by

$$\omega_{rp} = \left\{\omega \in [0,1]^N : \sum_{i=1}^N \omega_i = 1, \omega_i \geq 0, \omega_i \times (\Sigma\omega)_i = b_i \times (\omega'\Sigma\omega)\right\}, \qquad (2)$$

where ω_i is the i^{th} element of the weight vector, $(\Sigma\omega)_i$ denotes the i^{th} element of the corresponding vector and $b_i \geq 0$ is the weight risk budget for asset i with $\sum_{i=1}^N b_i = 1$. For the equally-weighted risk contribution (RP) portfolio case in which $b_i = b_j \; \forall i, j$, Maillard et al. (2010) show that $\sigma_p(\omega_{mv}) \leq \sigma_p(\omega_{rp}) \leq \sigma_p(\omega_{EW})$.

The question now is how shrinkage approaches and the risk budgeting restriction are related in a portfolio context. Let us start with the most basic setup of a bivariate RP portfolio. With $\omega = (\omega_1, 1 - \omega_1)$ and σ_1^2 and σ_2^2 denoting the variances of asset 1 and 2, it can be shown that for a long-only case ($0 \leq \omega_1 \leq 1$), the unique RP portfolio solution to Equation 1 is

$$\omega_{RP} = \left(\frac{\sigma_1^{-1}}{\sigma_1^{-1} + \sigma_2^{-1}}, \frac{\sigma_2^{-1}}{\sigma_1^{-1} + \sigma_2^{-1}}\right) = \left(\frac{\sigma_2}{\sigma_1 + \sigma_2}, \frac{\sigma_1}{\sigma_1 + \sigma_2}\right). \qquad (3)$$

This reveals the enriching property of the RP portfolio: The higher the variance in asset 2, the higher the weight of asset 1 and vice versa. Further, assuming that assets 1 and 2 are uncorrelated ($\rho_{12} = 0$), then the RP portfolio coincides with the minimum-variance portfolio ($\omega_{RP} = \omega_{MV}$).[4]

[4]It holds that $\omega_{MV} = \left(\frac{\sigma_2 - \sigma_{12}}{\sigma_1 + \sigma_2 - 2\sigma_{12}}, \frac{\sigma_1 - \sigma_{12}}{\sigma_1 + \sigma_2 - 2\sigma_{12}}\right)$ where $\sigma_{12} = \rho_{12}\sigma_1\sigma_2$ and ρ_{12} is the correlation

The empirical counterpart with estimated variance-covariance matrix is

$$\hat{\omega}_{RP} = \left(\frac{\hat{\sigma}_1^{-1}}{\hat{\sigma}_1^{-1} + \hat{\sigma}_2^{-1}}, \frac{\hat{\sigma}_2^{-1}}{\hat{\sigma}_1^{-1} + \hat{\sigma}_2^{-1}} \right) = \left(\frac{\hat{\sigma}_2}{\hat{\sigma}_1 + \hat{\sigma}_2}, \frac{\hat{\sigma}_1}{\hat{\sigma}_1 + \hat{\sigma}_2} \right), \tag{4}$$

where

$$\hat{\Sigma} = \begin{pmatrix} \hat{\sigma}_1^2 & \hat{\sigma}_{12} \\ \hat{\sigma}_{12} & \hat{\sigma}_2^2 \end{pmatrix} = \frac{1}{T-1} \sum_{t=1}^{T} \left(r_t - \bar{r} \right) \left(r_t - \bar{r} \right)', \text{ with} \tag{5}$$

$$\bar{r} = \left(\frac{1}{T} \sum_{s=1}^{T} r_{s,1}, \ldots, \frac{1}{T} \sum_{s=1}^{T} r_{s,N} \right). \tag{6}$$

While for low N inverting the estimated variance-covariance matrix ($\hat{\Sigma}$) to obtain the RP portfolio weights is not very costly in terms of estimation errors, it becomes a problem for larger asset universes or more correlated markets. Typically, dimension reduction techniques such as factor modeling or shrinkage towards a shrinkage target are valid attempts to mitigate estimator errors. A traditional shrinkage estimator for the variance-covariance matrix by Ledoit and Wolf (2004a) shrinks the sample variance-covariance matrix estimate towards the identity matrix (I). Given a shrinkage intensity $\kappa \in [0, 1]$, the approach yields

$$\begin{aligned} \hat{\Sigma}^* &= \kappa \, \mathrm{I} + (1-\kappa)\hat{\Sigma} \\ &= \begin{pmatrix} \kappa & 0 \\ 0 & \kappa \end{pmatrix} + (1-\kappa) \begin{pmatrix} \hat{\sigma}_1^2 & \hat{\sigma}_{12} \\ \hat{\sigma}_{12} & \hat{\sigma}_2^2 \end{pmatrix} \\ &= \begin{pmatrix} \kappa + (1-\kappa)\hat{\sigma}_1^2 & (1-\kappa)\hat{\sigma}_{12} \\ (1-\kappa)\hat{\sigma}_{12} & \kappa + (1-\kappa)\hat{\sigma}_2^2 \end{pmatrix} \end{aligned} \tag{7}$$

Since we assumed $\sigma_{12} = 0$ making the off-diagonals in $\hat{\Sigma}^*$ equal to zero, the RP portfolio for this case and replacing the traditional variance-covariance matrix with a shrinkage equivalent yields

$$\hat{\omega}_{RP}^* = \left(\frac{\frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_1^2}}}{\frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_1^2}} + \frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_2^2}}}, \frac{\frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_2^2}}}{\frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_1^2}} + \frac{1}{\sqrt{\kappa + (1-\kappa)\hat{\sigma}_2^2}}} \right).$$

coefficient between assets 1 and 2.

It directly follows that

$$\lim_{\kappa \to 0} \hat{\omega}^*_{RP} = \left(\frac{\hat{\sigma}_2}{\hat{\sigma}_1 + \hat{\sigma}_2}, \frac{\hat{\sigma}_1}{\hat{\sigma}_1 + \hat{\sigma}_2} \right) \tag{8}$$

$$\lim_{\kappa \to 1} \hat{\omega}^*_{RP} = \left(\frac{1}{2}, \frac{1}{2} \right). \tag{9}$$

Hence, we see that the shrinkage approach forces the RP portfolio towards the naive EW portfolio.

4 Empirical setup

Strategies

Given the results of Section 3, the question arises as to how RP fares compared to other shrinkage-based allocation schemes. The following section briefly describes the allocation techniques employed in the empirical analysis, which consist of heuristic weighting schemes as well as the most prominent shrinkage techniques in the literature. Table A.1 provides an overview of all strategies investigated in our study.

Inverse volatility

The inverse volatility strategy($VOLA$) is a heuristic allocation scheme that assigns weights according to asset volatilities, formally it is defined as

$$\hat{\omega}_i = \frac{1/\hat{\sigma}_i}{\sum_{i=1}^{N} 1/\hat{\sigma}_i}, \tag{10}$$

where $\hat{\sigma_i}$ is the estimated volatility of asset i.[5] The weight of an asset i equals its inverse volatility divided by the sum of inverse volatilities for all assets. As a results, $VOLA$ overweights low volatility stocks and underweights high volatility stocks. Hence, it is related to the low-volatility effect (Baker et al., 2011), which has been shown to outperform other heuristic weighting schemes such as equal- and value-weighting (Anderson et al., 2012). However, the total portfolio volatility is not minimized with $VOLA$ since asset covariances are not considered in the weighting process.

[5]This strategy is also referred to as naive risk parity, as it equalizes the risk contributions of assets when the correlations among assets are (almost) identical. See Maillard et al. (2010).

Linear shrinkage

Further, we employ the Ledoit and Wolf (2004a,b) linear shrinkage techniques, which can be described as an optimally weighted averages of the sample variance-covariance matrix of the asset universe ($\hat{\Sigma}$) and a shrinkage target ($\hat{F}$). The shrinkage approach can be expressed as follows

$$\hat{\Sigma}_{LWLS} = \kappa\hat{F} + (1-\kappa)\hat{\Sigma},$$

where α denotes the shrinkage intensity, $\hat{\Sigma}$ represents the sample variance-covariance matrix, and $\hat{F}$ describes the shrinkage target. We employ two different shrinkage targets, which are the constant correlation model ($MVCCM$; $RPCCM$) (Ledoit and Wolf, 2004a) and an identity matrix with dimension N ($MVID$;$RPID$) (Ledoit and Wolf, 2004b). Further, the shrinkage intensity κ is chosen accordingly to the references, minimizing the quadratic loss between the true and the estimated covariance matrices.[6]

Non-linear shrinkage

Recently, a new class of shrinkage methods has emerged, which are referred to as non-linear shrinkage techniques. Unlike linear shrinkage, in which all variance and covariance parameters shrink towards the corresponding target with the same amount, nonlinear shrinkage approaches shrink the eigenvalues of the estimated variance-covariance matrix. Thus, depending on how far these parameters lie from the shrinkage target, translates into regularized variance parameters and ultimately portfolio weights in a nonlinear fashion. In this study, we will focus on the analytical non-linear shrinkage technique proposed by Ledoit and Wolf (2017) as this is superior in terms of computational power and feasibility as compared with numerical approaches such as Abadir et al. (2014), and Lam (2016), among others. The analytical non-linear shrinkage estimator $\hat{\Sigma}_{LWNLS}$ is determined by

$$\hat{\Sigma}_{LWNLS} = U_T\hat{\delta}^o_T U_T \text{ with } \hat{\delta}^o_T := \text{ Diag}\left(\hat{\phi}^o_T(\lambda_{T,1}),\ldots,\lambda_{T,N})\right),$$

where $U_T := [u_{T,1},\ldots,u_{T,N}]$ is a (full) orthogonal matrix whose columns are the sample eigenvectors $u_{T,i}$, $\lambda_{T,i}$ denotes the sample eigenvalues, and $\hat{\phi}^o_T(x)$ represents the optimal shrinkage function.[7] We employ this technique on the risk parity ($RPNLS$) and minimum variance ($MVNLS$) portfolios.

---

[6]For more information the reader is refered to Ledoit and Wolf (2004a,b).

[7]For more information on the derivation of $\hat{\phi}^o_T(x)$ see Ledoit and Péché (2011), Theorem 3.

Empirical setup

Data were collected from Datastream and consist of daily total return price indices, which are price indices that are adjusted for dividends and stock splits, from 01.01.1994 to 30.09.2019. We employ an out-of-sample rolling-window approach as in DeMiguel et al. (2009). Each data set consists of $M = 6,783$ business days, where the asset moments are estimated based on the previous $T = 250$ days. Portfolios are held for $\theta = 20$ business days and are then rebalanced using the previous T business days. We test the portfolios in Table A.1 in different settings. These are: (1) In a multi asset allocation where $T \gg N$ and the risk-return properties of different asset classes vary to greater extent than when only considering assets within a certain class such as equity or bonds; and (2) high dimensional equity portfolios where plug-in estimates are quite unreliable (Jobson and Korkie, 1980).[8]

Let $P_{i,t}$ be the price of asset i at time t and $D_{i,t}$ be the dividend payment of asset i in period t, then the total return of asset i in period $t+1$ ($r_{i,t+1}$) is defined as

$$r_{i,t+1} = \frac{P_{i,t+1} + D_{i,t}}{P_{i,t}} - 1. \tag{11}$$

Accordingly, the total return of a strategy j with $j \in \{1, \ldots, 10\}$ (we compare 10 different strategies in total) and N assets at time $t+1$ is $r_{j,t+1} = \sum_{i=1}^{N}(r_{i,t+1} \times \hat{\omega}_{i,j,t})$, where $\hat{\omega}_{i,j,t}$ represents the weight of asset i for strategy j at time t. At rebalancing, a trade of $|\hat{\omega}_{i,j,t+1} - \hat{\omega}_{i,j,t^+}|$ occurs for each asset i held in strategy j, where $\hat{\omega}_{i,j,t^+}$ denotes the weight of asset i for strategy j immediately before rebalancing at time $t+1$, which is determined as $\hat{\omega}_{i,j,t^+} = \hat{\omega}_{i,j,t} \times (1 + r_{i,t+1})$. Following DeMiguel et al. (2009), let ξ represent a proportional transaction cost paid for each asset when rebalancing the portfolio, then the total cost incurred from trading all assets for portfolio j is obtained by $\xi \times \sum_{i=1}^{N}|\hat{\omega}_{i,j,t+1} - \hat{\omega}_{i,j,t^+}|$. As a result, the wealth from investing in strategy j at time $t+1$ is determined as

$$W_{j,t+1} = W_{j,t}(1 + r_{j,t+1})(1 - \xi \sum_{i=1}^{N}|\hat{\omega}_{i,j,t+1} - \hat{\omega}_{i,j,t^+}|), \tag{12}$$

where the return net of transaction costs for strategy j is given by $(W_{j,t+1}/W_{j,t}) - 1$ with $W_{j,t+1}$ and $W_{j,t}$ describing wealth at time $t+1$ and t, respectively. Moreover, as in DeMiguel et al. (2009), the turnover of strategy j is defined as the average sum over time of the absolute value of the trades incurred for all assets or formally

[8] Typically, when the number of observations T is not significantly larger than the number of assets N, a portfolio is referred to as a high dimensional portfolio.

$$Turnover_j = \frac{1}{M-T} \sum_{t=1}^{M-T} \sum_{i=1}^{N} (|\hat{\omega}_{i,j,t+1} - \hat{\omega}_{i,j,t^+}|). \quad (13)$$

Turnover measures the average percentage of total wealth that is traded at each rebalancing period and is related to the transaction costs incurred when implementing a strategy j. It is noteworthy to mention that a portfolio strategy which is superior in gross terms, that is in the absence of transaction costs, might not be optimal to implement when transaction costs are taken into account.

5 Applications in multi asset allocation

To demonstrate how RP fares compared with the aforementioned strategies, we analyze the characteristics of each strategy in an asset allocation setting, where all strategies are rebalanced on a monthly basis. The data set comprises of equities and bonds of all developed countries according to the MSCI classification.[9] Equities are represented by the MSCI total return indices and are denoted in USD. Bond markets are denoted by the Datastream 10 year government bond indices, and are hedged to USD according to the interest rate parity model using the Thomson Reuters spot and 1 month forward rates. Commodities are described by the Bloomberg Commodity Index and are denoted in USD. Table A.2 reports the risk-return characteristics of RP, EW, MV, and the shrinkage techniques highlighted in Table A.1.

All strategies result in a positive growth of wealth over time. In gross terms, MV yields the highest terminal wealth (\$340.40), followed by $RPNLS$ (\$299.57), RP (\$290.09), and $MVCCM$ (\$283.96), while other shrinkage and heuristic strategies yield a much lower gross wealth of circa \$240. However, MV is the most concentrated portfolio with an HHI of 53%. Shrinkage techniques decrease the portfolio concentration, with $MVID$ yielding the lowest concentration of 17.45% while $MVCCM$ yields an HHI of 32.85%. EW is the least concentrated portfolio by definition (2.36%), followed by $VOLA$ (3.88%), whereas the simple RP yields an HHI of 10.43%.

Secondly, shrinkage techniques for RP do not necessarily improve portfolio characteristics and generally do not vary significantly from RP. RP yields slightly better results than its shrinkage counterparts in terms of return, volatility, maximum drawdown, and even turnover. Nevertheless, the variation as mentioned earlier remains

[9] We exclude Hong Kong, Israel, and Singapore due to insufficient historical bond data.

marginal and merits no significant improvements. In terms of net Sharpe ratio, only $RPNLS$ (0.29) yields a value on par with RP (0.28), whilst $RPID$ (0.24) and $RPCCM$ (0.23) underperform. This indicates that, unlike for the case of MV, the regularization implied by the risk budget constraint on the weights of RP portfolios is quite stark and no longer enables the portfolio weights of significant further alteration.

Since a higher portfolio concentration typically results in a higher turnover ratio and, thus, higher transaction costs, the results change regarding net terminal wealth. MV performs worst due to having the highest turnover ratio (4.94), which translates into transaction costs of roughly 54% of total gross wealth, resulting in a terminal net wealth of merely \$156.81. This result coincides with the findings of Best and Grauer (1991), Michaud (1989), among others, that mean-variance optimized portfolios are highly sensitive to estimation risk, and typically result in high turnover ratios which ultimately leads to a significant underperformance relative to other weighting techniques. Meanwhile, RP yields the highest net terminal wealth of roughly \$260, followed by $VOLA$ (\$235.11) and EW (\$224.97). Prominent shrinkage techniques pose an improvement compared to MV, albeit remaining suboptimal compared to RP and the heuristic techniques in terms of net terminal wealth.

All portfolio returns are leptokurtic; indicating the existence of fat-tails in the return distribution. This is in line with the longstanding finding that financial returns are heavily entailed with extreme returns or fat-tails. Interestingly, although the traditional view on financial returns is that they are negatively skewed (Fama, 1965) and despite the presence of financial crises in our sample period, not all portfolios exhibit negatively skewed gross returns. RP and $MVID$ pose virtually no skewness, while MV reveals a positive skewness of 0.51. These results imply that in the absence of transaction costs, these allocation strategies are effective in mitigating the extreme negative returns that occur during market turbulences. Nonetheless, all portfolio returns are negatively skewed when taking transaction costs into account.

Regarding the maximum drawdown (MDD), which is the maximum loss during the sample period, $MVCCM$ and $MVNLS$ perform best in both gross and net terms. MV shows a quite low gross MDD (23.98%), but due to the high transaction costs, it results in a net drawdown of 35.83%. Heuristic strategies, namely EW and $VOLA$, exhibit the MDD along with $MVID$ of roughly 50%, while RP results in a drawdown of 40%. Finally, albeit MV yielding the highest Sharpe ratio of 0.66 in gross terms, it yields the lowest ratio in net terms of 0.17. $MVCCM$ has the highest ratio of 0.39, whereas $MVNLS$ and RP result in a ratio of roughly 0.30. The observed outperformance of $MVCCM$ coincides with the findings of Ledoit and Wolf (2004a), who find that $MVCCM$ performs best for samples with $N < 100$ assets.

To better understand the portfolio dynamics we also assess the performance measures over time in Figure B.1. In terms of historical performance, MV shows the highest terminal wealth. However, the only strategy to show positive performance during the 2008 financial crisis is $MVID$, an indication that naive diversification can be beneficial during market turmoils. While the differences in Sharpe ratio dynamics are not very profound, only MV based portfolios reduce portfolio volatility significantly. The same is true for the realized CVaR, where the magnitude for the MV shrinkage and non-shrinkage strategies is much smaller than for EW, $VOLA$, and the RP variants. RP dominates the variance regularization completely in the tails of the portfolio return distribution as the CVaR is identical for all RP strategies.

6 Applications in high-dimensional portfolios

Next, we assess the strategies in equity portfolios and the effect of dimensionality on portfolio performance, where we randomly select three cases with $N = 50, 100, 250$ equities from all stocks listed in the NYSE, AMEX, and NASDAQ exchanges. The data set is nominated in USD, consists of total return price indices, and ranges from 01.01.1994 to 31.12.2019. Tables A.3–A.5 report the results of the out-of-sample performance of each strategy for $N = 50, 100,$ and 250, respectively.

The distribution of portfolio returns remain unchanged to Section 5, where almost all strategies exhibit negatively skewed returns. MV is the most concentrated portfolio in all cases of N and exhibits the largest turnover rate, supporting Michaud (1989) that the mean-variance optimization using plug-in estimates results in highly concentrated and unstable portfolios. As the asset universe increases, not only does the concentration and turnover of MV increase substantially, but the return of the portfolio also drastically decreases whilst its volatility spikes. Again, this is in line with Michaud (1989) and Best and Grauer (1991), among others, that mean-variance optimized portfolios behave as error maximizers instead of achieving the mean-variance efficient portfolio, thereby deteriorating their risk-adjusted performance. As a result, MV performs well only for the case of $N = 50$ where it yields a Sharpe ratio of 0.63 net of transaction costs, thereby being suboptimal to only the covariance shrinkage strategies of Ledoit and Wolf (2004a,b), whereas its net Sharpe ratio declines with N yielding the lowest net Sharpe ratio of 0.25 for $N = 250$.

Regarding heuristic weighting schemes, EW yields stable albeit modest risk-adjusted returns for all N. As the number of assets increases, the risk-adjusted return of EW improves while the turnover remains quite low at around 0.35% of portfolio volume, yielding a higher net Sharpe ratio than MV for $N = 250$. $VOLA$ as well as RP display very similar properties, which are yet better than EW. The Sharpe ratio of both strategies increases with the number of assets and is almost identical for both strategies, although $VOLA$ exhibits a significantly lower turnover ratio of roughly 0.34.

Covariance shrinkage techniques provide mixed results. $MVID$ and $MVCCM$ yield the highest Sharpe ratios before and after transaction costs for all cases of N. The difference in risk-adjusted performance is minimal for $N = 50$ and 100, but the gap widens significantly for $N = 250$ with $MVID$ and $MVCCM$ respectively yielding a Sharpe ratio of 0.73 and 0.91 net transaction costs. Meanwhile, $MVNLS$ seems to shrink the portfolio towards EW as the number of assets increases, where $MVNLS$ is identical to EW for $N = 250$ in all portfolio characteristics. These results contradict the findings of Ledoit and Wolf (2019), who using Monte Carlo simulations find that non-linear shrinkage of the covariance matrix performs better than linear shrinkage in high dimensional portfolios. Our results indicate that a more structured shrinkage target, that is linear shrinkage of all parameters to a unified shrinkage target, results in a better risk-adjusted performance for MV portfolios than imposing less structure as in individual shrinkage targets via non-linear shrinkage.

Finally, shrinking the covariance matrix in RP results in a deterioration of portfolio performance for all cases of N, regardless of the approach. This result is in line with those of Section 5 and further reinforces the notion of the inflexibility of the asset weights in RP portfolios due to the risk budgeting constraint. Altogether, shrinkage techniques seem to exhibit some merits when applied in MV optimized portfolios, but do not yield significant benefits for variations thereof such as RP.[10]

[10] We refrain from including the rolling performance measures for the high dimensional portfolios in this study for parsimony. These results are available upon request from the authors.

7 Conclusion

In this study we investigate the impact of shrinkage estimation techniques for the variance-covariance matrix of asset returns in risk parity portfolios for different asset universe dimensions. While we provide theoretical and empirical evidence that risk budgeting on the asset level serves as a regularization mechanism, most importantly this chapter combines two strands of literature: Using mean-variance linear and non-linear shrinkage techniques in risk parity portfolios. In this, we provide a unique and novel perspective on risk parity portfolios.

We find that: (i) RP is a shrinkage variant that works well in lower dimensions; (ii) Combining RP with shrinkage estimators of the variance-covariance matrix yields even better returns in lower dimensions, but seems suboptimal for larger asset universes; (iii) pure shrinkage estimators for MV portfolios dominate all RP variants in higher dimensions, indicating that fixing the risk contribution of individual assets does not mitigate the effects of estimation errors and even dominates regularization effects from more stable variance estimators. A further argument for this is that for large N, the equal risk budgeting constraint does not decrease portfolio variance.

Future research should consider combinations of mean-variance and risk parity portfolios that specifically account for a better portfolio allocation rule that captures the best of both worlds; for example through an optimal shrinkage intensity similar to DeMiguel et al. (2013). Applications should also include more asset classes such as currencies or private equity, account for fat-tail distribution through a typical GARCH model as in Ardia et al. (2017) or even be concerned with more practical considerations in asset liability and risk management.

References

Abadir, K. M., W. Distaso, and F. Žikeš (2014). Design-free estimation of variance matrices. *Journal of Econometrics 181*(2), 165–180.

Amenc, N., F. Goltz, and A. Lodh (2012). Choose Your Betas: Benchmarking Alternative Equity Index Strategies. *Journal of Portfolio Management 39*(1), 88–111.

Anderson, R. M., S. W. Bianchi, and L. R. Goldberg (2012). Will My Risk Parity Strategy Outperform? *Financial Analysts Journal 68*(6), 75–93.

Ardia, D., G. Bolliger, K. Boudt, and J.-P. Gagnon-Fleury (2017). The impact of covariance misspecification in risk-based portfolios. *Annals of Operations Research 254*(1), 1–16.

Baker, M., B. Bradley, and J. Wurgler (2011). Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly. *Financial Analysts Journal 67*(1), 40–54.

Barry, C. B. (1974). Portfolio Analysis Under Uncertain Means, Variances, and Covariances. *Journal of Finance 29*(2), 515–522.

Benartzi, S. and R. H. Thaler (2001). Naive Diversification Strategies in Defined Contribution Saving Plans. *American Economic Review 91*(1), 79–98.

Best, M. J. and R. R. Grauer (1991). On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results. *Review of Financial Studies 4*(2), 315–342.

Black, F. and R. Litterman (1992). Global Portfolio Optimization. *Financial Analysts Journal 48*(5), 28–43.

Chan, L. K. C., J. Karceski, and J. Lakonishok (1999). On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model. *Review of Financial Studies 12*(5), 937–974.

Chaves, D., J. Hsu, F. Li, and O. Shakernia (2011). Risk Parity Portfolio vs. Other Asset Allocation Heuristic Portfolios. *Journal of Investing 20*(1), 108–118.

Chopra, V. K. and W. T. Ziemba (1993). The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice. *Journal of Portfolio Management 19*(2), 6–11.

Chow, T.-M., J. Hsu, V. Kalesnik, and B. Little (2011). A Survey of Alternative Equity Index Strategies. *Financial Analysts Journal 67*(5), 37–57.

Clarke, R., H. de Silva, and S. Thorley (2006). Minimum-Variance Portfolios in the U.S. Equity Market. *Journal of Portfolio Management 33*(1), 10–24.

Clarke, R., H. de Silva, and S. Thorley (2011). Minimum-Variance Portfolio Composition. *Journal of Portfolio Management 37*(2), 31–45.

Clarke, R., H. de Silva, and S. Thorley (2013). Risk Parity, Maximum Diversification, and Minimum Variance: An Analytic Perspective. *Journal of Portfolio Management 39*(3), 39–53.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *Review of Financial Studies 22*(5), 1915–1953.

DeMiguel, V., A. Martin-Utrera, and F. J. Nogales (2013). Size matters: Optimal calibration of shrinkage estimators for portfolio selection. *Journal of Banking & Finance 37*(8), 3018–3034.

DeMiguel, V. and F. J. Nogales (2009). Portfolio Selection with Robust Estimation. *Operations Research 57*(3), 560–577.

Fabozzi, F. J., P. N. Kolm, D. A. Pachamanova, and S. M. Focardi (2007). *Robust Portfolio Optimization and Management.* John Wiley & Sons.

Fama, E. F. (1965). The Behavior of Stock-Market Prices. *Journal of Business 38*(1), 34–105.

Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics 111*(1), 1–25.

Garlappi, L., R. Uppal, and T. Wang (2007). Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach. *Review of Financial Studies 20*(1), 41–81.

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

James, W. and C. Stein (1961). Estimation with Quadratic Loss. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pp. 361–379. Berkeley, CA: University of California Press.

Jobson, J. D. and B. Korkie (1980). Estimation for Markowitz Efficient Portfolios. *Journal of the American Statistical Association 75*(371), 544–554.

Lam, C. (2016). Nonparametric eigenvalue-regularized precision or covariance matrix estimator. *Annals of Statistics 44*(3), 928–953.

Ledoit, O. and S. Péché (2011). Eigenvectors of some large sample covariance matrix ensembles. *Probability Theory and Related Fields 151*(1-2), 233–264.

Ledoit, O. and M. Wolf (2004a). Honey, I Shrunk the Sample Covariance Matrix. *Journal of Portfolio Management 30*(4), 110–119.

Ledoit, O. and M. Wolf (2004b). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis 88*(2), 365–411.

Ledoit, O. and M. Wolf (2017). Nonlinear Shrinkage of the Covariance Matrix for Portfolio Selection: Markowitz Meets Goldilocks. *Review of Financial Studies 30*(12), 4349–4388.

Ledoit, O. and M. Wolf (2019). The Power of (Non-)Linear Shrinking: A Review and Guide to Covariance Matrix Estimation. *University of Zurich, Department of Economics, Working Paper No. 323, Revised Edition.*

Lee, W. (2014). Constraints and Innovations for Pension Investment: The Cases of Risk Parity and Risk Premia Investing. *Journal of Portfolio Management 40*(3), 12–20,6.

Maillard, S., T. Roncalli, and J. Teïletche (2010). The Properties of Equally Weighted Risk Contribution Portfolios. *Journal of Portfolio Management 36*(4), 60–70.

Markowitz, H. (1952). Portfolio Selection. *Journal of Finance 7*(1), 77–91.

Michaud, R. O. (1989). The Markowitz Optimization Enigma: Is 'Optimized' Optimal? *Financial Analysts Journal 45*(1), 31–42.

Pástor, L. (2000). Portfolio Selection and Asset Pricing Models. *Journal of Finance 55*(1), 179–223.

Wang, Z. (2005). A Shrinkage Approach to Model Uncertainty and Asset Allocation. *Review of Financial Studies 18*(2), 673–705.

Windcliff, H. and P. P. Boyle (2004). The 1/n Pension Investment Puzzle. *North American Actuarial Journal 8*(3), 32–45.

A Tables

Table A.1: List of Employed Portfolio Strategies

№	Strategy Definition	Abbreviation
Panel A: Heuristic Weighting Techniques		
1	Equally Weighted Portfolio	EW
2	Inverse Volatility Portfolio	$VOLA$
Panel B: Minimum Variance Portfolios		
3	Minimum Variance Portfolio	MV
4	Minimum Variance Portfolio with Linear Covariance Matrix Shrinkage using the Identity Matrix Model	$MVID$
5	Minimum Variance Portfolio with Linear Covariance Matrix Shrinkage using the Constant Correlation Model	$MVCCM$
6	Minimum Variance Portfolio with Non-Linear Covariance Matrix Shrinkage	$MVNLS$
Panel C: Risk Parity Portfolios		
7	Risk Parity Portfolio	RP
8	Risk Parity Portfolio with Linear Covariance Matrix Shrinkage using the Identity Matrix Model	$RPID$
9	Risk Parity Portfolio with Linear Covariance Matrix Shrinkage using the Constant Correlation Model	$RPCCM$
10	Risk Parity Portfolio with Non-Linear Covariance Matrix Shrinkage	$RPNLS$

Table A.2: Global Portfolio Summary Statistics

	EW	*IV*	*MV*	*MVID*	*MVCCM*	*MVNLS*	*RP*	*RPID*	*RPCCM*	*RPNLS*
Panel A: General Performance Metrics										
Gross Wealth	232.80	244.61	340.40	242.64	283.96	236.32	290.09	244.61	244.61	299.57
Net Wealth	224.97	235.11	156.81	200.31	190.69	174.69	259.84	235.11	227.94	269.97
TC ($)	7.82	9.50	183.59	42.33	93.27	61.63	30.26	9.50	16.67	29.60
TC (%)	3.36	3.88	53.93	17.45	32.85	26.08	10.43	3.88	6.81	9.88
Turnover	0.22	0.25	4.94	1.22	2.54	1.93	0.70	0.25	0.45	0.66
HHI	2.44	2.63	30.25	44.49	34.27	19.08	3.45	2.63	3.95	3.06
Panel B: Gross Performance Characteristics										
Annual Return	4.37	4.41	5.27	3.96	4.26	3.62	4.98	4.41	4.41	5.12
Annual Volatility	14.51	13.48	10.26	10.27	6.45	7.48	13.02	13.48	13.48	13.20
Maximum Drawdown	51.12	46.51	23.98	45.77	18.61	22.68	40.62	46.51	46.51	41.04
Skewness	−1.30	−1.14	0.52	0.08	−0.86	−0.81	0.06	−1.14	−1.14	0.05
Kurtosis	13.68	13.06	9.32	5.52	7.61	6.78	9.54	13.06	13.06	8.98
Sharpe Ratio	0.30	0.33	0.51	0.39	0.66	0.48	0.38	0.33	0.33	0.39
Panel C: Net Performance Characteristics										
Annual Return	3.14	3.31	1.74	2.69	2.50	2.16	3.70	3.31	3.19	3.84
Annual Volatility	15.04	13.87	10.18	10.24	6.46	7.54	13.04	13.87	13.87	13.21
Maximum Drawdown	51.24	46.63	35.83	49.34	20.11	24.87	40.90	46.63	46.68	41.30
Skewness	−2.18	−1.91	−0.15	−0.12	−1.04	−0.94	−0.46	−1.91	−1.91	−0.46
Kurtosis	21.44	19.21	8.21	5.44	8.40	7.21	10.04	19.21	19.22	9.62
Sharpe Ratio	0.21	0.24	0.17	0.26	0.39	0.29	0.28	0.24	0.23	0.29

Notes: This table provides the out-of-sample portfolio performance measures for the portfolio allocation schemes summarized in table A.1. The data set spans the period from January 1994 to December 2019, which consists of 6782 daily portfolio return observations. Portfolios are rebalanced on a monthly basis and all measures are annualized. *Gross Wealth* is the terminal wealth accumulated over the whole sample period on $100 investment while *Net Wealth* represents the terminal wealth net of transaction costs (*TC*), which are listed in nominal terms (*TC ($)*) as well as percentage points of Terminal Gross Wealth (*TC (%)*). *HHI* denotes the Herfindahl-Hirschman Index for portfolio concentration.

Table A.3: High Dimensional Portfolio Summary Statistics for N=50

	EW	*IV*	*MV*	*MVID*	*MVCCM*	*MVNLS*	*RP*	*RPID*	*RPCCM*	*RPNLS*
Panel A: General Performance Metrics										
Gross Wealth	572.67	709.89	557.18	1190.40	1575.46	1325.46	1042.98	683.87	710.02	710.02
Net Wealth	559.83	694.80	541.76	962.15	1468.18	1142.71	903.75	666.26	694.93	693.45
TC ($)	12.85	15.09	15.42	228.26	107.28	182.75	139.23	17.61	15.09	16.57
TC (%)	2.24	2.13	2.77	19.17	6.81	13.79	13.35	2.57	2.13	2.33
Turnover	0.36	0.34	0.45	3.40	1.13	2.37	2.29	0.42	0.34	0.38
HHI	2.00	2.15	2.08	6.84	9.80	6.97	5.54	2.25	2.15	2.21
Panel B: Gross Performance Characteristics										
Annual Return	8.87	9.13	8.72	10.57	11.72	10.90	9.95	8.92	9.14	9.14
Annual Volatility	19.72	16.92	19.47	13.54	14.16	12.91	12.84	16.52	16.92	16.92
Maximum Drawdown	63.88	57.25	64.22	40.54	34.68	33.37	37.37	57.34	57.25	57.25
Skewness	−1.66	−1.74	−1.68	−1.21	−0.30	−1.13	−1.05	−1.82	−1.74	−1.74
Kurtosis	12.98	13.19	13.64	9.90	6.16	10.53	8.42	14.16	13.19	13.19
Sharpe Ratio	0.45	0.54	0.45	0.78	0.83	0.84	0.77	0.54	0.54	0.54
Panel C: Net Performance Characteristics										
Annual Return	6.67	7.50	6.54	8.76	10.40	9.43	8.52	7.34	7.50	7.50
Annual Volatility	20.98	17.79	20.74	13.83	14.17	13.14	13.05	17.40	17.79	17.79
Maximum Drawdown	63.96	57.33	64.30	41.49	35.00	34.12	38.05	57.42	57.33	57.34
Skewness	−2.76	−2.62	−2.84	−1.70	−0.63	−1.65	−1.43	−2.76	−2.62	−2.62
Kurtosis	23.34	21.43	24.98	12.53	6.78	13.28	9.94	23.37	21.43	21.43
Sharpe Ratio	0.32	0.42	0.32	0.63	0.73	0.72	0.65	0.42	0.42	0.42

Notes: This table provides the out-of-sample portfolio performance measures for the portfolio allocation schemes summarized in table A.1. The data set spans the period from January 1994 to December 2019, which consists of 6782 daily portfolio return observations. Portfolios are rebalanced on a monthly basis and all measures are annualized. *Gross Wealth* is the terminal wealth accumulated over the whole sample period on $100 investment while *Net Wealth* represents the terminal wealth net of transaction costs (*TC*), which are listed in nominal terms (*TC ($)*) as well as percentage points of Terminal Gross Wealth (*TC (%)*). *HHI* denotes the Herfindahl-Hirschman Index for portfolio concentration.

Table A.4: High Dimensional Portfolio Summary Statistics for N=100

	EW	*IV*	*MV*	*MVID*	*MVCCM*	*MVNLS*	*RP*	*RPID*	*RPCCM*	*RPNLS*
Panel A: General Performance Metrics										
Gross Wealth	713.34	822.52	667.58	863.46	1599.21	1247.03	855.59	769.96	822.63	822.63
Net Wealth	697.79	805.43	648.31	544.59	1463.99	972.73	689.18	749.43	805.54	803.67
TC ($)	15.55	17.09	19.27	318.87	135.21	274.29	166.41	20.53	17.10	18.97
TC (%)	2.18	2.08	2.89	36.93	8.46	22.00	19.45	2.67	2.08	2.31
Turnover	0.35	0.34	0.47	7.36	1.41	3.97	3.45	0.43	0.34	0.37
HHI	1.00	1.07	1.04	5.16	5.47	4.60	3.32	1.10	1.07	1.08
Panel B: Gross Performance Characteristics										
Annual Return	9.43	9.55	9.17	9.36	11.73	10.54	9.07	9.24	9.55	9.55
Annual Volatility	18.38	16.09	18.31	13.99	13.89	12.06	12.03	15.75	16.09	16.09
Maximum Drawdown	61.02	54.36	61.89	34.97	36.94	28.18	37.20	54.90	54.36	54.36
Skewness	−1.62	−1.65	−1.67	−0.70	0.18	−0.44	−0.79	−1.78	−1.65	−1.65
Kurtosis	12.82	12.73	13.60	6.20	5.80	5.28	5.74	13.95	12.73	12.73
Sharpe Ratio	0.51	0.59	0.50	0.67	0.84	0.87	0.75	0.59	0.59	0.59
Panel C: Net Performance Characteristics										
Annual Return	7.52	8.08	7.24	6.56	10.39	8.81	7.47	7.80	8.08	8.07
Annual Volatility	19.40	16.81	19.37	14.13	13.76	12.07	12.14	16.51	16.81	16.81
Maximum Drawdown	61.09	54.44	61.97	37.56	37.29	28.99	38.00	54.99	54.44	54.45
Skewness	−2.61	−2.46	−2.75	−1.01	−0.09	−0.66	−1.03	−2.66	−2.46	−2.46
Kurtosis	21.91	19.99	23.99	7.36	5.23	5.70	6.48	22.47	19.99	19.98
Sharpe Ratio	0.39	0.48	0.37	0.46	0.75	0.73	0.62	0.47	0.48	0.48

Notes: This table provides the out-of-sample portfolio performance measures for the portfolio allocation schemes summarized in table A.1. The data set spans the period from January 1994 to December 2019, which consists of 6782 daily portfolio return observations. Portfolios are rebalanced on a monthly basis and all measures are annualized. *Gross Wealth* is the terminal wealth accumulated over the whole sample period on $100 investment while *Net Wealth* represents the terminal wealth net of transaction costs (*TC*), which are listed in nominal terms (*TC ($)*) as well as percentage points of Terminal Gross Wealth (*TC (%)*). *HHI* denotes the Herfindahl-Hirschman Index for portfolio concentration.

Table A.5: High Dimensional Portfolio Summary Statistics for N=250

	EW	*IV*	*MV*	*MVID*	*MVCCM*	*MVNLS*	*RP*	*RPID*	*RPCCM*	*RPNLS*
Panel A: General Performance Metrics										
Gross Wealth	701.17	811.03	1422.36	1505.93	2184.44	701.17	820.75	810.98	810.98	901.11
Net Wealth	685.68	794.10	277.81	1355.57	1306.38	685.68	796.64	794.06	791.73	882.75
TC ($)	15.49	16.92	1144.54	150.36	878.06	15.49	24.11	16.92	19.25	18.36
TC (%)	2.21	2.09	80.47	9.98	40.20	2.21	2.94	2.09	2.37	2.04
Turnover	0.36	0.34	26.01	1.68	8.20	0.36	0.48	0.34	0.38	0.33
HHI	0.40	0.43	3.49	2.68	2.65	0.40	0.44	0.43	0.43	0.45
Panel B: Gross Performance Characteristics										
Annual Return	9.36	9.49	11.57	11.52	12.60	9.36	9.48	9.49	9.49	10.07
Annual Volatility	18.37	16.05	15.75	14.05	10.94	18.37	15.71	16.05	16.05	17.04
Maximum Drawdown	60.74	54.47	42.21	45.02	32.06	60.74	55.48	54.47	54.47	57.04
Skewness	−1.62	−1.63	−0.32	0.25	−0.58	−1.62	−1.77	−1.63	−1.63	−1.55
Kurtosis	12.24	12.16	8.69	5.48	6.38	12.24	13.25	12.16	12.16	11.82
Sharpe Ratio	0.51	0.59	0.73	0.82	1.15	0.51	0.60	0.59	0.59	0.59
Panel C: Net Performance Characteristics										
Annual Return	7.45	8.02	3.96	10.09	9.95	7.45	8.03	8.02	8.01	8.43
Annual Volatility	19.36	16.74	15.81	13.89	10.95	19.36	16.43	16.74	16.74	17.80
Maximum Drawdown	60.81	54.54	52.77	45.32	33.66	60.81	55.56	54.54	54.55	57.14
Skewness	−2.52	−2.38	−0.92	0.01	−0.83	−2.52	−2.56	−2.38	−2.38	−2.34
Kurtosis	20.00	18.43	11.09	4.95	7.41	20.00	20.55	18.43	18.42	18.27
Sharpe Ratio	0.38	0.48	0.25	0.73	0.91	0.38	0.49	0.48	0.48	0.47

Notes: This table provides the out-of-sample portfolio performance measures for the portfolio allocation schemes summarized in table A.1. The data set spans the period from January 1994 to December 2019, which consists of 6782 daily portfolio return observations. Portfolios are rebalanced on a monthly basis and all measures are annualized. *Gross Wealth* is the terminal wealth accumulated over the whole sample period on $100 investment while *Net Wealth* represents the terminal wealth net of transaction costs (*TC*), which are listed in nominal terms (*TC ($)*) as well as percentage points of Terminal Gross Wealth (*TC (%)*). *HHI* denotes the Herfindahl-Hirschman Index for portfolio concentration.

B Figures

Figure B.1: Portfolio Sharpe Ratio under Estimation Error

(a) Minimum Variance

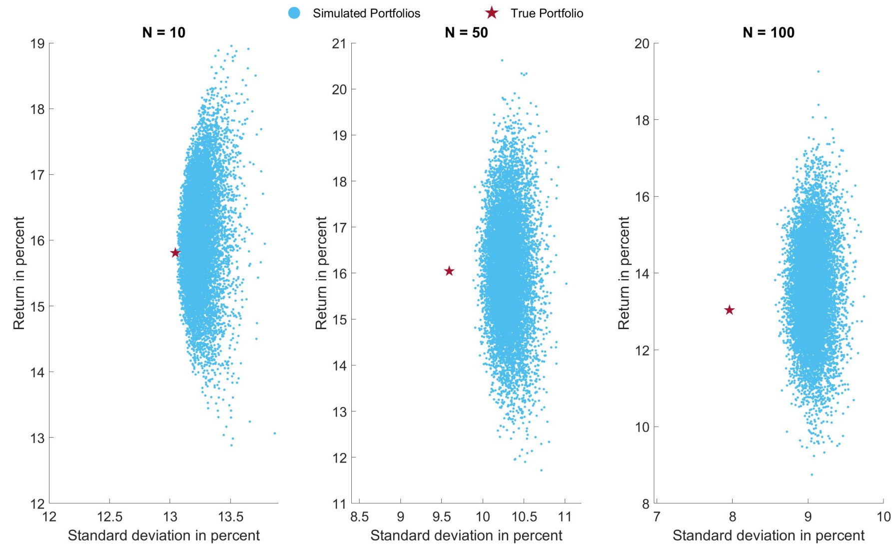

(b) Risk Parity

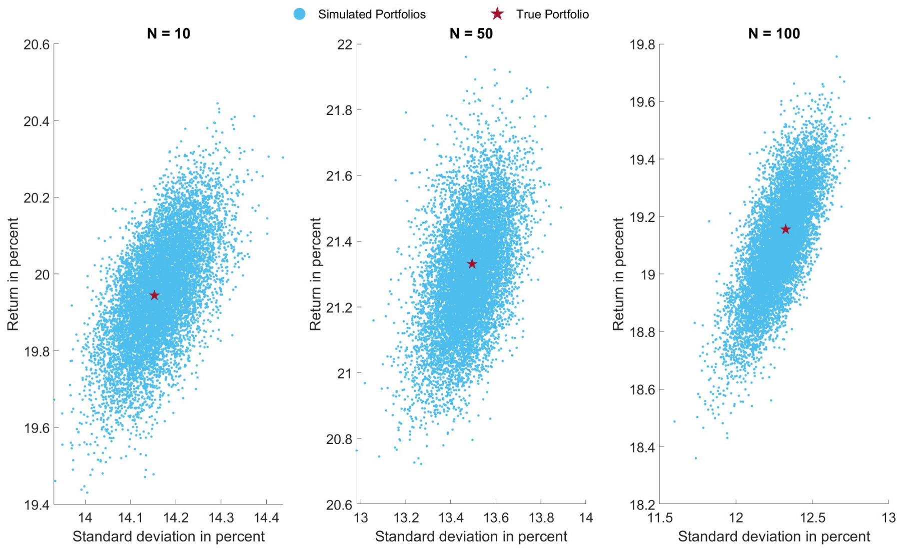

Notes: This figure illustrates the difference in the risk-return properties between the true portfolio and 10,000 simulated portfolios for the cases of $N = 10, 50, 100$ assets. The true portfolio is determined by the ex post sample moments, whereas the simulated portfolios are determined by plug-in estimates for 250 simulated returns.

Figure B.2: Global Portfolio Performance Measures over Time

(a) Historical Performance

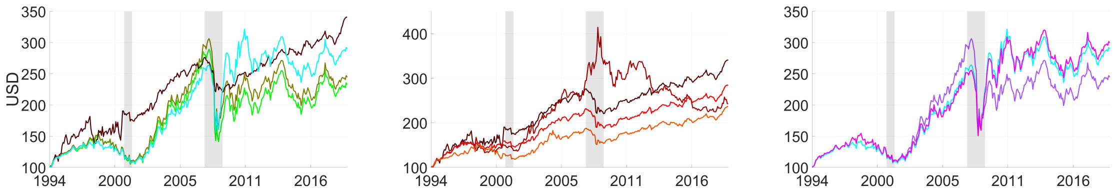

(b) Realized Sharpe

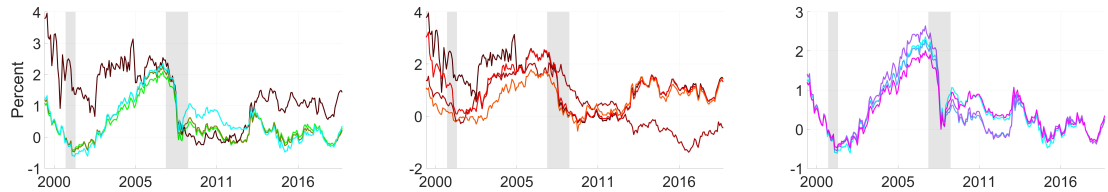

(c) Realized Volatility

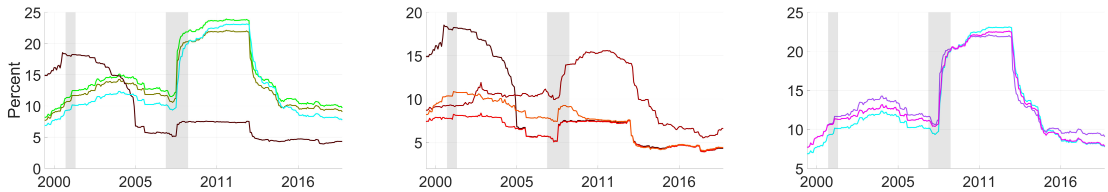

(d) Realized CVaR

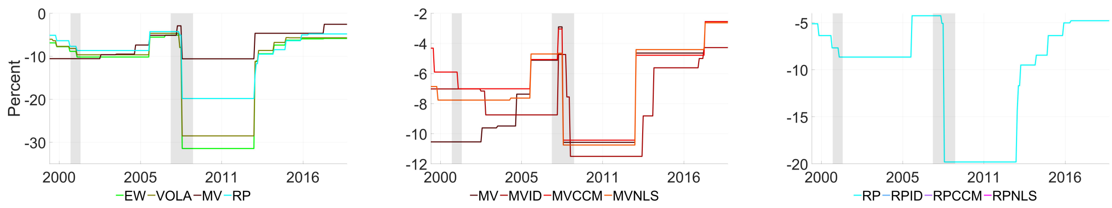

Notes: This figure illustrates the differences in gross historical performance, realized Sharpe ratio, realized volatility and realized Conditional Value at Risk (CVaR) between the different portfolio strategies outlined in Table A.1. The gray shaded areas indicate U.S. recessions as reported by NBER. The results in Panels (b) to (d) are calculated based on a rolling window of 60 monthly portfolio realized return observations.