Risk Parity and Beyond - From Asset Allocation to Risk Allocation Decisions

Romain Deguest, Lionel Martellini and Attilio Meucci*

August 16, 2013

Abstract

While it is often argued that allocation decisions can be best expressed in terms of exposure to rewarded risk factors, as opposed to somewhat arbitrary asset class decompositions, the practical implications of this paradigm shift for the optimal design of the policy portfolio still remain largely unexplored. This paper aims at analyzing whether the use of uncorrelated underlying risk factors, as opposed to correlated asset returns, can lead to a more efficient framework for measuring and managing portfolio diversification. Following Meucci (2009), we use the entropy of the factor exposure distribution as the number of uncorrelated bets (also known as the *effective number of bets*, or ENB in short), implicitly embedded within a given asset allocation decision. We present a set of formal results regarding the existence and unicity of portfolios designed to achieve the maximum effective number of bets. We also provide empirical evidence that incorporating constraints, or target levels, on a portfolio effective number of bets generates an improvement in out-of-sample risk-adjusted performance with respect to standard mean-variance analysis.

JEL Classification: G11, C44, C61

Keywords: portfolio choice; risk parity; diversification; concentration; principal component analysis.

*Romain Deguest is a senior research engineer at EDHEC-Risk Institute. Lionel Martellini is a professor of finance at EDHEC Business School and the scientific director of EDHEC-Risk Institute. Attilio Meucci is the Chief Risk Officer and Director of Portfolio Construction at Kepos Capital LP. The author for correspondence is Lionel Martellini. He can be reached at EDHEC-Risk Institute, 400 Promenade des Anglais, BP 3116, 06202 Nice Cedex 3 France. Ph: +33 (0)4 93 18 78 24. Fax: +33 (0)4 93 18 78 44. E-mail: lionel.martellini@edhec.edu.

1 Introduction

Recent research (e.g., Ang et al. (2009)) has highlighted that risk and allocation decisions could be best expressed in terms of rewarded risk factors, as opposed to standard asset class decompositions, which can be somewhat arbitrary. For example, convertible bond returns are subject to equity risk, volatility risk, interest rate risk and credit risk. As a consequence, analyzing the optimal allocation to such hybrid securities as part of a broad bond portfolio is not likely to lead to particularly useful insights. Conversely, a seemingly well-diversified allocation to many asset classes that essentially load on the same risk factor (e.g., equity risk) can eventually generate a portfolio with very concentrated risk exposure. More generally, given that security and asset class returns can be explained by their exposure to pervasive systematic risk factors, looking through the asset class decomposition level to focus on the underlying factor decomposition level appears to be a perfectly legitimate approach, which is also supported by standard asset pricing models such as the intertemporal CAPM (Merton, 1973) or the arbitrage pricing theory (Ross, 1976). Two main benefits can be expected from shifting to a representation expressed in terms of risk factors, as opposed to asset classes. On the one hand, allocating to risk factors may provide a cheaper, as well as more liquid and transparent, access to underlying sources of returns in markets where the value added by existing active investment vehicles has been put in question. For example, Ang et al. (2009) argue in favor of replicating mutual fund returns with suitably designed portfolios of factor exposures such as the value, small cap and momentum factors.[1] Similar arguments have been made for private equity and real estate funds, for example. On the other hand, allocating to risk factors should provide a better risk management mechanism, in that it allows investors to achieve an ex-ante control of the factor exposure of their portfolios, as opposed to merely relying on ex-post measures of such exposures.

While working at the level of underlying risk factors that impact/explain the returns on all asset classes is an intuitively meaningful approach, the practical implications of this paradigm shift for the organization of the asset allocation process still remain largely unexplored. This paper aims at analyzing whether using uncorrelated principal component factors, as opposed to correlated asset returns, effectively allows for a more efficient framework for measuring and managing policy portfolio diversification. In this paper, we divide the analysis into two steps, namely measuring portfolio diversification and managing portfolio diversification. In the first step, we follow Meucci (2009) and use the entropy of the factor exposure distribution as a measure of the number of (equal-size) uncorrelated bets (also known as effective number of bets, or ENB in short) implicitly embedded within a given portfolio. Then we provide evidence, using the policy portfolio of a large state pension fund as an example, that even a seemingly well-diversified portfolio may end up loading on a very limited number of independent risk factors. In the second step, we aim at building a better diversified portfolio, also known as

[1] In the same vein, Hasanhodzic and Lo (2007) argue for the passive replication of hedge fund vehicles, even though Amenc et al. (2010) found that the ability of linear factor models to replicate hedge fund performance is modest at best.

factor risk parity (or FRP) portfolio, by maximizing the effective number of bets. We first show that starting with an investment universe of N assets, there exist 2^{N-1} distinct portfolios that achieve the maximum effective number of bets. In this context, we introduce a procedure for identifying two remarkable FRP portfolios, the one with the lowest volatility (denoted by FRP-MV) and the one with the highest Shape ratio (denoted by FRP-MSR). In an empirical application, we use principal component analysis (PCA) to extract uncorrelated factors from correlated asset returns, we provide evidence that incorporating constraints (or target levels) on a portfolio effective number of bets generates an improvement in out-of-sample risk-adjusted performance with respect to standard mean-variance analysis or ad-hoc benchmarks. We also outline a number of weaknesses of PCA in the context of the design of well-diversified factor risk parity portfolios, including the lack of clear interpretation for the factors, a concern over factor stability, as well as the difficulty to estimate the sign of the factor risk premium, which is needed for the derivation of the FRP-MSR portfolio. In a follow-up paper (Meucci et al. (2013)), we explore a competing approach, known as *minimal torsion* approach, for extracting uncorrelated factors from correlated asset returns, which is shown to alleviate the aforementioned concerns raised by the use of PCA.

Our paper is closely related to the literature on *risk budgeting*, which advocates allocating to various constituents in a portfolio so as to achieve a given target in terms of contribution of these constituents to the total risk of the overall portfolio (see Roncalli (2013) for a comprehensive analysis of risk budgeting techniques). When the target is taken to be an equal contribution for all constituents, then the risk parity portfolio is sometimes called an *equal risk contribution* (ERC) portfolio or a *risk parity* (RP) portfolio (see Maillard et al. (2010), Bruder and Roncalli (2012) or Lee (2011)). Particularly related to ours are a series of recent papers that have proposed to apply the ERC approach to uncorrelated underlying factor returns, as opposed to correlated asset returns. In this vein, Lohre et al. (2012, 2011), or Poddig and Unger (2012) use principal component analysis to extract uncorrelated factors and analyze the out-of-sample performance of factor risk parity (FRP) portfolios, that is portfolios achieving the maximum effective number of bets (see also Roncalli (2013) for applications to fixed-income portfolio construction and asset allocation decisions). Our contribution with respect to these papers is first to demonstrate that maximizing the effective number of bets does not lead to a unique solution, and to propose a natural procedure for choosing one particular portfolio amongst the 2^{N-1} distinct FRP portfolios. We also complement these papers by providing a formal comparison of FRP portfolios with respect to mean-variance analysis. Our paper is also related to the literature on the benefits of weight constraints (see Jagannathan and Ma (2003) for *hard* minimum and maximum weight constraints and DeMiguel et al. (2009) for *flexible* constraints applying to the norm of the weight vector), which are known to be critically useful ingredients in portfolio optimization models since they imply a minimum level of naive diversification in the portfolio, and which can be formally interpreted as providing an implicit form of statistical shrinkage similar to the one discussed in Ledoit and Wolf (2003, 2004).[2] In particular, we show

[2] See also Jurczenko et al. (2013) for a more general approach nesting both hard and flexible weight constraints.

that mean-variance analysis with constraints on the effective number of independent bets is equivalent to a form of shrinkage towards a target portfolio that minimizes the factor exposure, as opposed to the weight vector as in DeMiguel et al. (2009).

The rest of the paper is organized as follows. In Section 2, we discuss various approaches that can be used to *measure* the diversification/concentration of a given portfolio, and provide an empirical illustration of the usefulness of these measures for a pension fund. In Section 3, we turn to the management of portfolio diversification in an asset allocation framework. In Section 4, we provide an empirical analysis of the performance of various diversification-optimized portfolios. Section 5 concludes. Technical details are relegated to a dedicated Appendix.

2 Measuring Portfolio Diversification - The Effective Number of Bets

A key distinction exists between weight-based measures of portfolio concentration, which are based on the analysis of the portfolio weight distribution independently of the risk characteristics of the constituents of the portfolio, and risk-based measures of portfolio concentration, which incorporate information about the correlation and volatility structure of the return on the portfolio constituents. In a nutshell, weight-based measures can be regarded as a measure of naive diversification, while risk-based measures can be regarded as measures of scientific diversification.

2.1 Notations

In this paper, we use the following notation for the characteristics of the portfolio constituents:

- $\boldsymbol{\Sigma}$ for the covariance matrix of the assets;
- $\boldsymbol{\Omega}$ for the correlation matrix of the assets;
- $\boldsymbol{V}$ for the matrix of assets' volatilities on the diagonal and 0 elsewhere;

$$\boldsymbol{\Sigma} = \boldsymbol{V}\boldsymbol{\Omega}\boldsymbol{V}$$

- $\boldsymbol{\mu}$ for the vector of expected excess returns of the assets.

We assume that all portfolios contain N constituents, where $\boldsymbol{w}$ denotes the weight vector representing the percentage invested in each constituent. The vector of ones will be denoted by $\mathbf{1}_N$, and the identity matrix by $\boldsymbol{I}_N$. In the following, we will also always assume that all portfolios satisfy the budget condition, i.e., $\sum_{k=1}^{N} w_k = 1$ (or equivalently $\mathbf{1}_N'\boldsymbol{w} = 1$).

2.2 Benefits and Limits of Traditional Weight-Based and Risk-Based Measures

Looking at the *nominal* number of constituents in a portfolio, as what was done in early studies such as Evans and Archer (1968) or Statman (1987), is a simplistic indication of how well-diversified a portfolio is, since it does not convey any information about the relative dollar contribution of risk contribution of each constituent to the portfolio. Weight-based measures of portfolio concentration, which can be regarded as measures of the effective number of constituents (ENC) in a portfolio, provide a quantitative estimate of the concentration of a portfolio. The academic and practitioner literatures have considered various such measures, which can be seen as a naive way to quantify the diversification of a portfolio. These measures are built from the norm of the portfolio weight vector (see for example DeMiguel et al. (2009)):

$$\mathrm{ENC}_\alpha(\boldsymbol{w}) = \|\boldsymbol{w}\|_\alpha^{\frac{\alpha}{1-\alpha}} = \left(\sum_{k=1}^{N} w_k^\alpha\right)^{\frac{1}{1-\alpha}}, \qquad \alpha \geq 0, \quad \alpha \neq 1. \tag{2.1}$$

Taking $\alpha = 2$ leads to a diversification measure defined as the inverse of Herfindahl Index, which is itself a well-known measure of portfolio concentration, or $\mathrm{ENC}_2(\boldsymbol{w}) = \frac{1}{\sum_{k=1}^{N} w_k^2}$. It can be shown that when α converges to 1, then the norm-based distance converges to the *entropy* of the distribution of the portfolio weight vector:[3]

$$\mathrm{ENC}_1(\boldsymbol{w}) = \exp\left(-\sum_{k=1}^{N} w_k \ln(w_k)\right). \tag{2.2}$$

It is straightforward to check that for positive weights the ENC_α measures reach a minimum equal to 1 if the portfolio is fully concentrated in a single constituent, and a maximum equal to N, the nominal number of constituents, attained by the equally-weighted portfolio. These properties justify using these measures to compute the effective number of constituents. In spite of their intuitive appeal, these weight-based measures suffer however from a number of major shortcomings, and most notably from the following two main limits. On the one hand ENC measures can be deceiving when applied to assets with non homogenous risks. Consider for example a position invested for 50% in a 1% volatility bond, and the other 50% in a 30% volatility stock. The weights are perfectly distributed, but the risk is highly concentrated. This is due to differences in the total variance of each constituent, with $(50\%)^2 \times (30\%)^2$ being much larger than $(50\%)^2 \times (1\%)^2$, thus implying that the equity allocation has a much larger contribution to portfolio risk compared to the bond allocation. On the other hand, ENC measures can be deceiving when applied to assets with correlated risks. For instance, consider a portfolio with equal weights invested in two bonds with similar duration and volatility. Despite the fact that weights and risks are homogeneously distributed between both bonds, risk is still very concentrated because of the high correlation between the two bonds.

[3] This result is known in information theory under the following statement: the Rényi entropy converges to the Shannon entropy.

Intuitively, diversification measures should also account for information about the covariance matrix $\boldsymbol{\Sigma}$ of asset returns. A number of such risk-based measures of diversification have been introduced by various authors. For example, one can consider the ratio of the variance of some general portfolio of investable assets to the weighted average variance of these investable assets (see Goetzmann et al. (2005) for the use of this measure for equally-weighted portfolios):

$$\mathrm{GLR}(\boldsymbol{w}) = \frac{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}{\sum_{k=1}^{N} w_k \sigma_k^2} \tag{2.3}$$

This measure takes into account not only the number of available assets but also the correlation properties. More specifically, a portfolio that concentrates weights in assets with high correlation will tend to have portfolio risk higher than the average standalone risk of each of its constituents. Thus it will have a high Goetzmann-Li-Rouwenhorst measure, that is, high correlation-adjusted concentration.

2.3 The Effective Number of Bets as an Operational Measure of Portfolio Diversification

While afore-mentioned weight-based measures of the effective number of constituents and risk-based measures (of distance between the whole portfolio and the -weighted- sum of the components) are meaningful diversification measures, they nonetheless provide very little information about the effective number of bets (ENB) in a portfolio. Assuming that one can decompose the portfolio return as the sum of N uncorrelated factors, Meucci (2009) proposes to use the entropy of the portfolio factor exposure distribution as an operational measure of risk diversification/concentration. In simple terms, if diversification is the art and science of avoiding to have all eggs in the same basket, one could say that ENB measures the effective number of (uncorrelated) baskets, while ENC merely measures the effective number of eggs.

The ENB measure relies on the choice of N uncorrelated factors, whose returns can easily be expressed as $\boldsymbol{r}_F = \boldsymbol{A}'\boldsymbol{r}$, where $\boldsymbol{r}$ is the vector of constituents' returns and $\boldsymbol{A}$ is a square matrix of size N. The important assumptions on the factors, hence on matrix $\boldsymbol{A}$, are of two kinds:

- The factors are uncorrelated, namely $\boldsymbol{\Sigma}_F = \boldsymbol{A}'\boldsymbol{\Sigma}\boldsymbol{A} \in \mathcal{D}^+$, where $\mathcal{D}^+$ is the set of diagonal matrices of size N with strictly positive entries;
- The transition matrix $\boldsymbol{A}$ from the constituents to the factors is invertible.

The advantage of introducing uncorrelated factors is that the total variance of the portfolio can be written as a sum of the contributions of each factor's variance:

$$\begin{aligned}\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} &= \boldsymbol{w}'\left(\boldsymbol{A}'\right)^{-1}\boldsymbol{\Sigma}_F\boldsymbol{A}^{-1}\boldsymbol{w}\\ &= \boldsymbol{w}_F'\boldsymbol{\Sigma}_F\boldsymbol{w}_F\\ &= \sum_{k=1}^{N}\left(\sigma_{F_k}w_{F_k}\right)^2,\end{aligned}$$

where $\boldsymbol{w}_F = \boldsymbol{A}^{-1}\boldsymbol{w}$ can be interpreted as a portfolio of factors, and $\boldsymbol{\Sigma}_F$ is the diagonal matrix containing the factors' variances. Then, one can define the contribution of the k^{th} factor to the total portfolio variance as $p_k = \frac{\left(\sigma_{F_k}w_{F_k}\right)^2}{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}$. Each p_k is positive and the sum of all p_k's is equal to 1. Therefore, one can use the same dispersion measures (2.1) as for the ENC and define the ENB as the distribution of factor's contributions to the total variance:

$$\text{ENB}_\alpha(\boldsymbol{w}) = \|\boldsymbol{p}\|_\alpha^{\frac{\alpha}{1-\alpha}} = \left(\sum_{k=1}^{N} p_k^\alpha\right)^{\frac{1}{1-\alpha}}, \qquad \alpha \geq 0, \quad \alpha \neq 1. \tag{2.4}$$

As in the case of the ENC measure, an entropy-based measure of the effective number of bets can be defined by taking the limit $\alpha \to 1$:

$$\text{ENB}_1(\boldsymbol{w}) = \exp\left(-\sum_{k=1}^{N} p_k \ln(p_k)\right). \tag{2.5}$$

We write ENB measures as direct functions of the constituent weights $\boldsymbol{w}$, since it is clear from the definition of each $p_k = \frac{\left(\sigma_{F_k}\times\left[\boldsymbol{A}^{-1}\boldsymbol{w}\right]_k\right)^2}{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}$ that the vector $\boldsymbol{p}$ is an explicit function of $\boldsymbol{w}$. Again, ENB measures reach a minimum equal to 1 if the portfolio is loaded in a single risk factor, and a maximum equal to N, the nominal number of constituents, if the risk is evenly spread amongst the factors.

2.4 Empirical Illustration: Analyzing the Policy Portfolio of a State Pension Fund

In this section, we provide evidence, using the actual policy portfolio of a large US state pension fund as an example, that even seemingly well-diversified portfolios may end up loading on a very limited number of independent risk factors. In our empirical study, we consider 7 asset classes: US Treasury bonds (using the Bank of America Merrill Lynch US master treasury bond index as a proxy), US corporate bonds (using the Bank of America Merrill Lynch US master corporate bond index as a proxy), US large cap stocks (using S&P 500 equity index as a proxy), US private equity (using the S&P 600 small cap equity index as a proxy, in the absence of liquid high-frequency benchmarks for private equity), international equities (using the FTSE World Ex-US equity index as a proxy), real estate (using the Dow-Jones total market US real estate index as a proxy), and commodities (using the S&P Goldman Sachs commodity index as a proxy).[4] The period considered for our analysis is from 3 January 1992 to 29 June 2012, hence includes 1,069 weekly returns. In Table 1, we give the main descriptive statistics of each asset class over the entire sample period. We notice that the lowest volatility is achieved by Treasury bonds (4.63%), with corporate bonds as the second lowest volatility asset class (with a volatility of 5.20%). Overall, we see that bonds are at least three times less volatile than any other asset class considered in the empirical analysis. Private equity, real estate and commodities show volatilities above 20%. Not only do corporate bonds have a low volatility, but they also exhibit the highest Sharpe ratio of 0.76 on the sample period (treasury bonds come second with 0.69). US, private equity and real estate exhibit higher average returns over the period, but also significantly higher volatilities, resulting in lower Sharpe ratios (around 0.30). Looking at panel (b) of Table 1, we also see that the correlations among the various equity indexes considered are high (above 69.40%). We notice that Treasury bonds are negatively correlated with all the other asset classes, while corporate bonds are positively correlated with US and non-US equities. Finally, we note that real estate is also highly correlated with equities over the sample period (between 50% and 75% depending on the equity index). In order to better understand what happened over the period of analysis, we have estimated the volatilities over rolling windows of two years (104 weekly returns) in Figure 1. Hence, we notice that bonds are the smallest contributors to the sum of all volatilities among the asset classes whatever the time period. On the other hand, the commodity index consistently seems to be the highest contributor to portfolio risk.

2.4.1 Extracting Uncorrelated Factors

In this section, we perform a Principal Component Analysis (PCA) on the sample covariance matrix $\mathbf{\Sigma}$ estimated over the entire sample period in order to identify the statistical factors that drive asset classes' returns. Table 2 shows the exposures of each of these factors with respect to the asset classes. The exposures provided in Table 2 lie between -100% and 100%, since by convention, the L^2-norm of each factor's exposures is equal to 1.[5] The resulting eigenvalues obtained from the PCA are always positive, and represent the variances of each factor. The first factor can be interpreted as an equity risk since it is mostly exposed to the various equity indices, with an exposure to bonds close to 0 (less than 3%). We also notice that its exposure to the real estate is significant (52.78%), which can be explained by the high correlation observed in panel (b) of Table 1 between equities and the REITS index which we use as an imperfect proxy for real estate. The second factor can clearly be interpreted as commodity risk since its exposure to commodities is equal to 94.29% while the other weights remain below 26.89% in absolute value. The third factor represents pure real estate risk since it is massively loaded on the real estate index (74.24%), with a negative exposure to equities, thus capturing the idiosyncratic variation of the real estate that is not explained by equities. The fourth factor is long non-US equities (70.82%) and short private and US equities (-53.25% and -33.27%

[4] We also use the 3-month US T-bill index as a proxy for the risk-free asset.

[5] This does not imply that the factor exposures sum to 1.

respectively), which shows that it can naturally be interpreted as an international equity risk factor. The fifth and sixth factors seem to be both related to interest rate factors. More precisely the fifth factor is heavily loaded on both bonds (48.41% for treasury and 56.13% for corporate) and US equities (53.74%), which shows that it focuses on the interest rate risk that is correlated with the US equity market. On the other hand, the sixth factor is capturing the interest rate risk that is not explained by the US equity market since it is negatively exposed to bonds (-42.89% for treasury and -50.69% for corporate) and positively exposed to US equities (53.75%). Finally, the last factor clearly represents a proxy for credit risk since it is long the Treasury bond index (76.06%) and short the corporate bond index (-64.83%).

By construction, all these factors are uncorrelated, and explain the entire variability of the return on the 7 asset classes. However, their explanatory power is not homogeneous, as illustrated by the variances of each factor given in Table 2. Indeed, the first factor explains 60.84% of the total variance, the first two factors explain 81.30%, and the first three explain already more than 90%. As already explained, these factors are related to equities, commodities, and real estate, which tend to be very volatile asset classes (see again Figure 1 for an analysis of the volatility of each asset class over time).

2.4.2 Analyzing the Level of Diversification/Concentration of a Policy Portfolio

We now consider a specific allocation in the asset classes introduced previously. This allocation can be seen as a particular example of a policy portfolio of a large US state pension fund. The weights are given in the top panel of Table 3. If one were to use a weight-based measure such as the ENC_1 in order to quantify the portfolio concentration, we would obtain $\frac{ENC_1}{N} = \frac{5.90}{7} =$ 84.29%. Even though this weight-based measure suggests a rather well-balanced portfolio, it fails to reflect the real diversification with respect to underlying sources of risk, because of the shortcomings of the ENC measure discussed in Section 2.2, in particular in the light of highly correlated returns for some asset classes, as well as inhomogeneities in volatility levels. To address this concern, we use the ENB measure instead, which replaces the asset weight the distribution of factor risk contributions. In order to compute the ENB, we use the factors extracted from the PCA described in Section 2.4.1 run on the sample covariance matrix $\boldsymbol{\Sigma}$ of the risky asset classes returns:

$$\boldsymbol{\Sigma} = \boldsymbol{P}\boldsymbol{\Lambda}\boldsymbol{P}' \qquad \text{where} \qquad \boldsymbol{P}\boldsymbol{P}' = \boldsymbol{I}_N, \quad \text{and} \quad \boldsymbol{\Lambda} = \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_N \end{pmatrix} \tag{2.6}$$

Then, the relation between the asset class returns $\boldsymbol{r}$ and the factor returns is defined as:

$$\boldsymbol{r}_F = \boldsymbol{A}'\boldsymbol{r} \quad \text{where } \boldsymbol{A} = \boldsymbol{P}, \tag{2.7}$$

and the covariance matrix $\boldsymbol{\Sigma_F}$ of the factor returns is equal to:

$$\boldsymbol{\Sigma_F} = \boldsymbol{A}'\boldsymbol{\Sigma}\boldsymbol{A} = \boldsymbol{P}'\boldsymbol{P}\boldsymbol{\Lambda}\boldsymbol{P}'\boldsymbol{P} = \boldsymbol{\Lambda} \in \mathcal{D}^{+}. \tag{2.8}$$

Computing the ENB measure with the policy portfolio and using the PCA factors described previously leads to: $\frac{ENB_1}{N} = \frac{1.20}{7} = 17.14\%$, which shows that the portfolio is highly concentrated in terms of factor exposure. In order to better analyze the exposition of the fund to each risk factor, we compute in the bottom panel of Table 3 the factor weights resulting from the above asset allocation. From these results, it clearly appears that the policy portfolio is positively exposed to the first factor (36.20%) which is equity risk, and also to the third (-12.97%) and fifth (16.79%) factors which represent real estate and domestic interest rate risks. We further notice that 96.69% of the total policy portfolio variance is explained by the first factor, which illustrates that the variations in returns on the policy portfolio are almost exclusively driven by equity risk.

One might wonder at this stage whether this extremely concentrated factor allocation could be supported by the investor's views regarding a higher reward on the equity risk factor with respect to other factors. To see this, in the spirit of the first step in the model of Black and Litterman (1992, 1991), we take as given the covariance matrix and compute what should be the implicit views regarding the Sharpe ratios for the factors so as to rationalize the policy allocation as a maximum Sharpe ratio (MSR) allocation.[6] The result given in the last row of the bottom panel of Table 2.4.1 shows an arguably unreasonable set of views regarding how each factor is rewarded. Indeed, if the first one is set to 100% by convention, we see that it implies a negative expected reward for real estate and commodity risks, and a very small positive reward for domestic interest rate risk (9.16%). In this context, one might wonder if the ENB diversification measure can be used as a target of a constraint in the process of designing better diversified policy portfolios. This is the purpose of the following section.

3 Managing Portfolio Diversification - Factor Risk Parity Portfolios

The operational definition of a well-diversified portfolio is unambiguous in modern portfolio theory, which states that all investors aim at maximizing the risk/reward ratio. In the presence of estimation risk, however, which road should be taken to reach the highest risk/reward is less straightforward. In a first section, we present four popular strategies that are used in practice to construct policy portfolios, namely the Equally Weighted (EW) portfolio, the Global Minimum Variance (GMV) portfolio, the Equal Risk Contribution (ERC) portfolio, and the Maximum Sharpe Ratio (MSR) portfolio. We also recall the conditions for the heuristic strategies to be optimal in the sense of maximizing the Sharpe ratio. Then, in a second section we consider an approach based on the maximization of the ENB measure as an additional heuristic strategy

[6] The factors' risk premiums have been normalized so that the first one is equal to 100%.

which we call Factor Risk Parity (FRP) portfolio strategy.

3.1 Some Popular Portfolio Allocation Strategies

We first review a couple of heuristic portfolio strategies focusing on equal allocation of dollar budgets (equally-weighted portfolio strategy) or risk budgets (equal-risk contribution strategies), before turning to portfolio strategies grounded in modern portfolio theory.

3.1.1 Heuristic Portfolio Strategies

In the absence of any information on risk $\boldsymbol{V}$, return $\boldsymbol{\mu}$ parameters of portfolio constituents and on their correlations $\boldsymbol{\Omega}$, the most natural approach to portfolio diversification is the so-called "naive" diversification approach, which suggests investing an equal amount in each constituent of the portfolio:

$$\boldsymbol{w}^{\text{EW}} = \frac{1}{N}\mathbf{1}_N. \tag{3.1}$$

The out-of-sample performance of this naive diversification is extensively studied in DeMiguel et al. (2009), where the authors find that it leads to higher Sharpe ratios than following an optimal diversification of equity stocks. This heuristic portfolio strategy can also be regarded as the solution to the simple following optimization problem[7]:

$$\max_{\boldsymbol{w}} \text{ENC}_\alpha(\boldsymbol{w}), \quad \text{subject to } \mathbf{1}_N'\boldsymbol{w} = 1. \tag{3.2}$$

While maximizing the ENC, the EW may lead to a concentrated risk exposure because it disregards the differences in volatilities and correlations of the constituents. In particular, while they are assigned equal *dollar* contributions, their *risk* contributions can be very different due to differences in volatilities. For instance, because some constituents have higher volatilities, they would contribute more to the total risk of an EW portfolio. The idea of the so-called ERC method is precisely to equalize the contributions to overall risk. This approach, which is formally analyzed in Maillard et al. (2010), is based on the following additive decomposition for the total portfolio volatility:

$$\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}} = \sum_{k=1}^{N} w_k \frac{\partial\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}{\partial w_k}.$$

The weights are then computed so as to equalize the relative contributions of each constituent to the volatility:

$$w_k \frac{\partial\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}{\partial w_k} = \frac{1}{N}\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}, \quad \text{for all } i = 1, ..., N.$$

[7] If one wants to consider more weight constraints in the portfolio specification, such as linear constraints (e.g., constraints on factor exposure), or quadratic constraints (e.g., tracking-error constraints), then it is simple to add these constraints to Problem (3.2).

While this portfolio strategy does not lead to analytical expressions for portfolio weights, Maillard et al. (2010) show that it can be obtained from the solution to the following optimization problem[8]:

$$\boldsymbol{y}^*(c) = \underset{\boldsymbol{y}}{\operatorname{argmax}}\, \sigma(\boldsymbol{y}) \quad \text{subject to } \mathbf{1}_N' \ln(\boldsymbol{y}) \geq c, \text{ and } \boldsymbol{y} \geq \mathbf{0}_N, \tag{3.3}$$

by setting[9]:

$$\boldsymbol{w}^{\text{ERC}} = \frac{\boldsymbol{y}^*(c)}{\mathbf{1}_N' \boldsymbol{y}^*(c)}. \tag{3.4}$$

One drawback of this otherwise intuitively appealing approach is that it disregards the fact that large portfolios may be driven by a small number of factors. In this context, a seemingly well-diversified portfolio, with an equal dollar or risk contribution of each constituent, may end up being heavily exposed to a very limited number of factor exposures, thus making it a not so well-diversified portfolio from an intuitive standpoint. This limitation of the ERC approach can be addressed with the factor risk parity methodology. More generally, it can be shown that the ERC portfolio verifies the duplication invariance property, which states that the allocation should not be impacted if an asset is artificially duplicated, if the risk budgets are expressed with respect to uncorrelated factors as opposed to correlated assets (see Roncalli (2013)).

3.1.2 Efficient Frontier Portfolio Strategies

Modern portfolio theory, starting with the seminal work of Markowitz (1952), states that all mean-variance investors rationally seek to maximize the Sharpe ratio, subject to the constraint that the portfolio is fully invested in the N assets. This program reads:

$$\max_{\boldsymbol{w}} \lambda(\boldsymbol{w}) = \frac{\boldsymbol{w}'\boldsymbol{\mu}}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}}}, \quad \text{subject to } \mathbf{1}_N'\boldsymbol{w} = 1. \tag{3.5}$$

If $\mathbf{1}_N'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} > 0$, this problem has a well-known solution[10] (Merton, 1971):

$$\boldsymbol{w}^{\text{MSR}} = \frac{\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{\mathbf{1}_N'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}. \tag{3.6}$$

In practice, to compute the MSR portfolio one thus needs to estimate the vector of expected excess returns $\boldsymbol{\mu}$, and the covariance matrix $\boldsymbol{\Sigma}$, which are unobservable. Another efficient frontier portfolio is the global minimum variance (GMV) portfolio, namely the portfolio strategy

[8] In the design of ERC strategies, one minimizes the volatility of the portfolio subject to an additional constraint ($\mathbf{1}_N' \ln(\boldsymbol{y}) \geq c$), where the target c is arbitrary. Hence, the ERC portfolio can be seen as an intermediary portfolio between the global minimum variance (GMV) and the EW portfolios, since it is a variance-minimizing strategy subject to a lower bound on diversification in terms of constituent weights.

[9] We notice that for the solution of Problem (3.3) to be well-defined, one has to consider long-only portfolios ($\boldsymbol{y} \geq \mathbf{0}_N$), and we will only consider long-only ERC portfolios in the empirical study. Indeed, if long-only constraints are not imposed, multiple solutions may exist with equal relative contributions. Imposing long-only constraints enables pinpointing a unique solution with equal risk contributions. A similar (lack of) unicity problem exists for the so-called factor risk parity portfolios that will be also analyzed in this paper.

[10] This condition implies that the expected excess return of the GMV is greater than the risk-free interest rate.

that minimizes the portfolio variance $\sigma^2(\boldsymbol{w}) = \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}$:

$$\min_{\boldsymbol{w}} \sigma^2(\boldsymbol{w}), \quad \text{subject to } \mathbf{1}_N'\boldsymbol{w} = 1. \tag{3.7}$$

a program which admits the following closed-form solution:

$$\boldsymbol{w}^{\text{GMV}} = \frac{\boldsymbol{\Sigma}^{-1}\mathbf{1}_N}{\mathbf{1}_N'\boldsymbol{\Sigma}^{-1}\mathbf{1}_N}. \tag{3.8}$$

One often advocated advantage of the use of the GMV is that it does not require any estimate for expected returns, and as such is well-suited in any situation where such estimates are unreliable. However, it tends to overweight low-volatility constituents, and as a result to lead to highly concentrated portfolios.[11] In this context, DeMiguel et al. (2009) proposes to introduce a constraint on the minimum ENC, and show that the introduction of this constraint, which can be seen as an extension of the shrinkage approach (Jagannathan and Ma (2003), Ledoit and Wolf (2003, 2004)), leads to a substantial improvement in out-of-sample Sharpe ratios. Relationships exist between heuristic portfolios and the maximum Sharpe ratio portfolio. In particular, it is straightforward to show that (i) the MSR portfolio coincides with the EW portfolio if expected excess returns, volatilities, and pair of correlations are assumed to be equal, (ii) the MSR portfolio coincides with the ERC portfolio if all portfolio constituents have the same Sharpe ratio, and if all pairs of correlation are identical; and (iii) the MSR portfolio coincides with the GMV portfolio if expected excess returns are all identical.

3.2 Introducing FRP Portfolios

In this section, we adopt the portfolio variance decomposition suggested by the ENB measure definition. In other words, we decompose the total variance as the sum of contributions of each risk factor and not as the sum of contributions of each constituent. The factor weights obtained with FRP portfolios are therefore defined so as to equalize the relative contribution of each factor to the total portfolio variance[12]:

$$p_k = \frac{1}{N} \quad \Leftrightarrow \quad (\sigma_{F_k} w_{F_k})^2 = \frac{1}{N}\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}, \quad \text{for all } k = 1, ..., N.$$

This implies that each factor weight w_{F_k} can be written as $\pm\frac{\gamma}{\sigma_{F_k}}$ for $k = 1, ..., N$ where γ

[11] By "concentrated portfolios", we mean here concentrated in terms of dollar allocation, not in terms of risk allocation. In fact, as argued before, overweighting low volatility components would be consistent with having a more balanced portfolio in terms of risk contributions.

[12] Roncalli and Weisang (2012) explore the construction of FRP-like portfolios with a number of correlated factors less than the number of constituents. To build FRP portfolios with a fully arbitrary number of factors, correlated or uncorrelated, the reader could use the framework developed in Meucci et al. (2013).

is a constant. In matrix form, this leads to the following multiple solutions:

$$\boldsymbol{w}_F = \gamma \boldsymbol{\Sigma}_F^{-\frac{1}{2}} \begin{pmatrix} \pm 1 \\ \vdots \\ \pm 1 \end{pmatrix}.$$

We immediately notice that the FRP portfolio strategy is not uniquely defined. Indeed, there are 2^N weight vectors $\boldsymbol{w}_F$ that lead to equal contributions of each factor. The number of solutions depends on the choice of the sign for each of the N entries of the vector of ones. In order to go back to the original constituent weight vectors we use the following identity $\boldsymbol{w} = \boldsymbol{A}\boldsymbol{w}_F$. Then, adding the budget constraint $\mathbf{1}_N'\boldsymbol{w} = 1$ finally leads to the following closed-form expression for the FRP portfolios:

$$\boldsymbol{w}^{\text{FRP}} = \frac{\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}} \begin{pmatrix} \pm 1 \\ \vdots \\ \pm 1 \end{pmatrix}}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}} \begin{pmatrix} \pm 1 \\ \vdots \\ \pm 1 \end{pmatrix}}. \tag{3.9}$$

Note that we are left with 2^{N-1} solutions for the constituents' weights.[13] Moreover, we see from the expression of the solutions that flipping the sign of the k^{th} entry of vector $\boldsymbol{w}_F$ together with the sign of the k^{th} column of matrix $\boldsymbol{A}$ has no impact on the weight vector $\boldsymbol{w}$. Since changing the sign of the k^{th} column of matrix $\boldsymbol{A}$ is equivalent to changing the sign of the k^{th} factor, and since the choice of the factor signs is arbitrary, we can choose the vector $\boldsymbol{w}_F$ with all positive entries and span the FRP solutions by flipping the signs of the uncorrelated factors in matrix $\boldsymbol{A}$. It is straightforward to see that the FRP portfolios given in Equation (3.9) are all solutions to the following optimization problem:

$$\max_{\boldsymbol{w}} \text{ENB}_\alpha(\boldsymbol{w}), \quad \text{subject to } \mathbf{1}_N'\boldsymbol{w} = 1. \tag{3.10}$$

The FRP portfolios with positive signs coincide with the ERC portfolios of Maillard et al. (2010) when the constituents are uncorrelated and the factors are chosen to be the original constituents. The following proposition gives the conditions on the constituents' expected excess returns and covariance structures for the FRP portfolios to be seen as the result of a MSR strategy.

Proposition 1 *Consider the set of FRP portfolios that are solutions to the optimization problem* (3.10)*. We arbitrarily pick a sequence of signs for the vector of ones and call this vector*

[13] There are 2^N different vectors $\boldsymbol{w}_F$, but the final solution in (3.9) is normalized by $\frac{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{w}_F}{\gamma}$ giving a total of $\frac{2^N}{2}$ solutions since both $\boldsymbol{w}_F$ and $-\boldsymbol{w}_F$ lead to the same solution.

(resp. the corresponding FRP portfolio) $\mathbf{1}_N^j$ *(resp.* $\boldsymbol{w}^{FRP,j}$*).*

An MSR portfolio that coincides with $\boldsymbol{w}^{FRP,j}$ *exists if the estimates for the uncorrelated factors' expected excess returns are assumed to be equal to*[14]*:*

$$\boldsymbol{\mu}_F^j = \kappa\, \boldsymbol{\Sigma}_F^{\frac{1}{2}}\, \mathbf{1}_N^j,$$

implying the following estimates for the assets' expected excess returns:[15]

$$\boldsymbol{\mu}^j = \left(\boldsymbol{A}'\right)^{-1} \boldsymbol{\mu}_F^j = \kappa \left(\boldsymbol{A}'\right)^{-1} \boldsymbol{\Sigma}_F^{\frac{1}{2}} \mathbf{1}_N^j.$$

Proof. See Appendix A.1. ■

Hence, an FRP portfolio can be seen as the proper starting point if one has access to information about the volatility and the correlation structure of the constituents, and if one uses the agnostic prior that all factors have the same Sharpe ratio in absolute value. In practice, one needs to address the issue raised by the multiplicity of FRP portfolios that are solutions to (3.10). In this context, a natural way to select one FRP portfolio amongst the 2^{N-1} FRP portfolios consists in selecting the FRP that has the highest Sharpe ratio among all the FRP solutions to (3.10). The following proposition gives an explicit expression for the weights of this particular FRP portfolio, which we call the FRP-MSR portfolio (for factor risk parity portfolio with maximum Sharpe ratio).

Proposition 2 *Consider the set of all FRP portfolios that are solutions to the optimization problem* (3.10)*. Define one of the FRP as follows:*

$$\boldsymbol{w}^{FRP\text{-}MSR} = \frac{\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{MSR}}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{MSR}}, \tag{3.11}$$

where the sign of each entry of vector $\mathbf{1}_N^{MSR}$ *is the same as the sign of its corresponding factor's Sharpe ratio:*

$$\mathbf{1}_N^{MSR} = sign\left(\boldsymbol{\lambda}_F\right) = sign\left(\boldsymbol{A}'\boldsymbol{\mu}\right).$$

- *If* $\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{MSR} > 0$*, then* $\boldsymbol{w}^{FRP\text{-}MSR}$ *has the highest Sharpe ratio among all FRP solutions to* (3.10)*.*
- *In that case, its Sharpe ratio is equal to:* $\lambda\left(\boldsymbol{w}^{FRP\text{-}MSR}\right) = \frac{\sum_{k=1}^N |\lambda_{F_k}|}{\sqrt{N}}$.

Proof. See Appendix A.2. ■

One key insight is that to estimate the MSR portfolio subject to $ENB = N$, one only needs to know the sign of excess returns on factors. This stands in contrast with the MSR portfolio

[14]This condition implies that the factors' Sharpe ratios are all equal in absolute value, but the reverse does not hold.

[15]The choice of the sign of the multiplicative factor κ is such that the MSR exists, i.e., such that $\mathbf{1}_N'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^j > 0$. Notice that κ is therefore of the same sign as the normalization factor of $\boldsymbol{w}^{\text{FRP},j}$, i.e. $\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j$.

with no constraints on ENB, for which the absolute value of expected return on factors are also needed. Since the factors' volatilities are positive, we have $\text{sign}(\boldsymbol{\lambda}_F) = \text{sign}(\boldsymbol{\mu}_F)$. In practice, one should distinguish between two different contexts. In case factors can be easily interpreted, one may have a direct view on the sign of each factor's expected excess return, and the FRP-MSR portfolio can be obtained with additional information about the covariance matrix. In case factors are statistical factors, then one can try and use views on assets' expected excess returns $\boldsymbol{\mu}$ to derive information on factors's expected excess returns using $\boldsymbol{\mu}_F = \boldsymbol{A}'\boldsymbol{\mu}$. While this approach requires some views on the assets' expected excess returns, only the signs of $\boldsymbol{\mu}_F$ are eventually used to implement the FRP-MSR. In the empirical illustration, we will use a somewhat agnostic approach for the estimation of the sign of factors by assuming that all asset classes have the same Sharpe ratios, which will allow us to obtain estimates for the sign of the excess expected return on the factors as explained above.[16]

In the absence of any view on assets' expected excess returns, then one may instead select the FRP portfolio with the lowest volatility, as opposed to highest Sharpe ratio. The following proposition provides an explicit characterization for this portfolio, which we call the FRP-MV (for factor risk parity portfolio with minimum variance).

Proposition 3 *Consider the set of all FRP portfolios that are solutions to the optimization problem* (3.10)*. Define one of the FRP as follows:*

$$\boldsymbol{w}^{FRP\text{-}MV} = \frac{\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{GMV}}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{GMV}}, \tag{3.12}$$

where the sign of each entry of vector $\mathbf{1}_N^{GMV}$ is the same as the sign of its corresponding entry in $\boldsymbol{A}'\mathbf{1}_N$:

$$\mathbf{1}_N^{GMV} = sign\left(\boldsymbol{A}'\mathbf{1}_N\right).$$

- *Then $\boldsymbol{w}^{FRP\text{-}MV}$ has the lowest volatility among all FRP solutions to* (3.10)*.*
- *Its volatility is equal to: $\sigma\left(\boldsymbol{w}^{FRP\text{-}MV}\right) = \frac{\sqrt{N}}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{GMV}}$.*

Proof. The proof of this proposition directly follows from the proof of Proposition 2. ■

Note that in this case the choice of the sign of the position in factor k depends on the sign of all the elements of column k of matrix $\boldsymbol{A}$. This means that the sign of the position in a given factor is the same as the sign of the total exposure to underlying asset classes needed to replicate that factor k. In other words, if factor k is globally long (respectively, short) in the asset classes, then its sign in the FRP-MV portfolio will be positive (respectively, negative).

3.3 Portfolio Strategies under ENB Constraints

Another way to take into account our operational measure of diversification in a portfolio optimization procedure consists in using it as a constraint, as opposed to using it as a target. Hence,

[16]The same assumption will be maintained for the construction of the MSR portfolio without any restriction on the ENB.

one could perform mean-variance analysis (e.g., GMV or MSR) with a given target/constraint on the portfolio ENB. This approach is somewhat similar to portfolio optimization with "norm" (or ENC) constraints studied in DeMiguel et al. (2009), except that the FRP portfolio as opposed to the EW portfolio is now used as a anchor point for naive diversification. We therefore consider the following program:

$$\min_{\boldsymbol{w}} \sigma^2(\boldsymbol{w}) \quad \text{subject to } \mathbf{1}_N' \boldsymbol{w} = 1, \text{ and } \mathrm{ENB}_\alpha(\boldsymbol{w}) \geq \mathrm{ENB}^*, \tag{3.13}$$

where ENB^* is a minimal target effective number of bets. This generalizes program (3.7), which is the special case where $\mathrm{ENB}^* = 1$. In the following, we will focus on the L^2-norm for the definition of the ENB measure in (2.4)[17]. The following proposition shows that we can interpret the ENB_2-constrained GMV portfolios as GMV portfolios that result from shrinking the elements of the sample covariance matrix $\boldsymbol{\Sigma}$.

Proposition 4 *If the ENB_2-constrained GMV problem* (3.13) *has a solution, then it is also solution to the unconstrained GMV problem* (3.7) *where the sample covariance matrix* $\boldsymbol{\Sigma}$ *is replaced by the matrix:*

$$\bar{\boldsymbol{\Sigma}} = \boldsymbol{\Sigma} + \nu \left(\boldsymbol{A}'\right)^{-1} \boldsymbol{M} \boldsymbol{A}^{-1}. \tag{3.14}$$

Here $\nu \geq 0$ is the Lagrange multiplier for the ENB_2 constraint at the solution to the ENB_2-constrained GMV problem and $\boldsymbol{M} = [m_{ij}]_{1 \leq i,j \geq N}$ is a target matrix that depends on the factor weights that are solutions to the ENB_2-constrained GMV problem:

$$\begin{aligned} m_{ij} &= -w_{F_i} w_{F_j} \sigma^2_{F_i} \sigma^2_{F_j} && \textit{if } i \neq j \\ m_{ii} &= (ENB^* - 1)\, w^2_{F_i} \sigma^4_{F_i} && \textit{otherwise} \end{aligned}$$

Proof. See Appendix A.3. ■

This proposition illustrates that the solution to the ENB_2-constrained GMV problem (3.13) can be seen as a GMV portfolio for which each element of the factors' covariance matrix have been shrunk towards $\boldsymbol{M}$. Note that the diagonal is always positively shrunk while the shrinkage of the other entries depends on the signs of the factors' weights. This result is very similar to Proposition 6 of DeMiguel et al. (2009) which looks at GMV portfolios under ENC constraints. In particular, we show that mean-variance analysis with constraints on the effective number of independent bets is equivalent to a form of shrinkage of the shortsale-unconstrained sample-based minimum variance portfolio towards a target portfolio that minimizes the norm of the factor exposure, namely the FRP-MV portfolio, versus a portfolio that maximizes the norm of the weight vector, namely the EW portfolio.

[17] If we had considered the L^1-norm, the constraint would have translated into a redundant constraint on portfolio variance. Alternatively, one could have used the L^1-norm with a number of factors lower than the number of constituents.

4 Empirical Analysis

In this section, we run backtests of the portfolio strategies described in Section 3 and compare the out-of-sample risk-adjusted performance of these strategies. In particular, we implement the EW, ERC, GMV, MSR portfolios as well as the two FRP portfolios that are defined in Propositions 2 and 3. The calibration period for the computation of the sample covariance matrix in our empirical study is two years (104 weekly returns), and the rebalancing of the portfolios occurs every quarter (we apply a buy-and-hold allocation between two consecutive rebalancing dates). For the implementation of FRP portfolios without look-ahead bias, we use the N uncorrelated factors resulting from the PCA (see Section 2.4.2 for more details) applied to the sample covariance matrix estimated over rolling windows of the two previous years of weekly data.

4.1 FRP Strategies in the Absence of Shortsale Constraints

Table 4 provides all the descriptive statistics for the two popular portfolios GMV and MSR, together with the two FRP portfolios defined in Proposition 2 and 3. The EW and ERC portfolios, which give positive weights, will be discussed in the next Section. The construction of the MSR and the FRP-MSR portfolio requires the estimation of expected excess returns. Here we use the arbitrary agnostic prior that all assets have the same Sharpe ratios, leading to expected excess returns that are proportional to the assets' volatilities.[18]

From Table 4, we see that the FRP-MSR portfolio defined in Proposition 2 exhibits a higher Sharpe ratio than the FRP-MV portfolio defined in Proposition 3. This out-of-sample result supports the result of Proposition 2 which states that the FRP-MSR portfolio is, amongst all FRP portfolios that are solutions to Problem 3.10, the one with the highest ex-ante Sharpe ratio. On the other hand, the FRP-MV portfolio shows a lower volatility than the FRP-MSR portfolio which also supports the findings of Proposition 3 stating that the FRP-MV portfolio should have the lowest ex-ante volatility among all FRP portfolios. The difference of average return (7.92% for the FRP-MSR versus 6.36% for the FRP-MV), of volatility (5.88% for the FRP-MSR versus 5.33% for the FRP-MV), and of Sharpe ratio levels (0.81 for the FRP-MSR versus 0.60 for the FRP-MV) shows that the characteristics of two FRP portfolios can be significantly different, hence enlightening robustness issues in the implementation of FRP strategies. We also check that both portfolios have an effective number of bets equal to the $N = 7$.

Comparing the Sharpe ratio obtained by the FRP-MSR and MSR portfolios, we see that forcing the effective number of bets to be maximal (equal to the number of asset classes N) leads to a slight decrease in out-of-sample Sharpe ratio (0.83 for the MSR, and 0.81 for the FRP-MSR). On the other hand, the FRP-MSR portfolio has a much higher average ENB measure compared to the unconstrained MSR portfolio (7 for the FRP-MSR portfolio versus

[18]Note that the FRP-MSR uses the expected excess returns to compute the signs of factors' expected excess returns, while the MSR portfolios uses the assets' expected excess returns directly.

3.77 for the unconstrained MSR portfolio). A competing approach to the selection of the FRP strategy is to implement the FRP-MV portfolio. The advantage of this approach is that it does not rely on $\boldsymbol{\mu}$ since the choice of the sign in the closed-form expression of the FRP-MV is simply obtained by setting $\mathbf{1}_N^{\text{MV}} = \text{sign}(\boldsymbol{A}'\mathbf{1}_N)$. As a result, this strategy is less subject to estimation risk since it does not require any estimate for the signs of the Sharpe ratios since only the (observable) sign of the global position in the asset classes is needed to obtain the sign of the factor weights. The risk-adjusted performance of the FRP-MV portfolio is less attractive than that of the FRP-MSR portfolio, but still remains reasonably close to that of the unconstrained GMV portfolio.

4.2 FRP Strategies in the Presence of Shortsale Constraints

In this section, we analyze the same four strategies as in Section 4.1, except that each portfolio allocation is now implemented with short-selling constraints applied at the underlying asset class level. We also consider other strategies that have been introduced in Section 3 that satisfy the no shortsale constraint, namely the EW and ERC portfolio strategies. Finally, we also simulate the performance of the current allocation of the pension fund regarded as a fixed-mix strategy. Table 5 provides all the descriptive statistics for the portfolio optimization strategies introduced in Section 3 together with the ad-hoc pension fund policy allocation described in Section 2.4.2 (denoted by "fund" in Table 5). To generate the FRP portfolios, we use the following procedure. We first solve the maximum FRP objective under shortsale constraints, in order to know the maximal effective number of bets, $\text{ENB}_1^{\max}$, that can be attained in the presence of shortsale constraints. Then, the implementation of the FRP-MSR portfolio consists in numerically maximizing the Sharpe ratio among all the FRP portfolios that achieve an effective number of bets equal to $\text{ENB}_1^{\max}$. In the same spirit, the implementation of the FRP-MV relies on a numerical minimization of the variance among all the FRP portfolios that achieve an effective number of bets equal to $\text{ENB}_1^{\max}$.

We find that the current policy portfolio, as well as the EW strategy, exhibit high average returns together with significantly higher volatilities, as well as extremely large drawdown levels, which result in relatively low Sharpe ratios for both strategies. We also notice that the out-of-sample minimum volatility is equal to 4.13% and is attained by the GMV strategy, suggesting that this strategy has achieved its stated objective despite the presence of estimation risk. We also find that the highest out-of-sample Sharpe ratio is now obtained by the FRP portfolios (0.90 for the FRP-MV portfolio and 0.86 for the FRP-MSR portfolio, versus a Sharpe ratio of 0.83 for the MSR portfolio). While firm general conclusions cannot be drawn on the basis of a single sample and a single universe of assets, this result illustrates the fact that introducing ENB constraints can be an effective approach to generating higher risk-adjusted performance in the presence of estimation errors on risk and return parameters, with expected return parameter estimates being known to be the more noisy estimates (see Merton (1980)). Moreover, we also notice that the ERC portfolio, which does not rely on any expected return estimate, achieves the same out-of-sample Sharpe ratio as the MSR portfolio.

It is interesting to note that adding the shortsale constraints seems to have narrowed down the number of portfolios that have the highest ENB measure. In particular, the FRP-MV and the FRP-MSR portfolios now appear to be very similar (they have a correlation of 99.51%). If we look at the top panel of Figure 4, it seems that they even strictly coincide until 2011, and then they show a slight discrepancy that has been beneficial to the FRP-MV in terms of expected return and Sharpe ratio. Concerning the other statistics that are provided in Table 5, we see that the average ENB is equal to 5.30 for both the FRP-MSR and the FRP-MV portfolios since the addition of short-selling constraints prevents the allocation from achieving its highest ENB equal to $N = 7$. On the other hand, the average ENB is much lower for ad-hoc strategies (1.20 for the fund allocation policy and 1.43 for the EW strategy), which confirms again that deconcentrating a portfolio is not equivalent to generating a well-balanced exposure to underlying risk factors. The other portfolio optimization approaches, namely the GMV, MSR, and ERC portfolios, exhibit average ENB levels that lie between those of the EW and FRP strategies.

When looking at the maximum drawdown and VaR levels, we find that the ERC portfolio is the most exposed to extreme risks amongst all optimized strategies. In particular, the GMV has an extremely low maximum drawdown (less than 6%), while the MSR and FRP strategies have similar extreme risk statistics (maximum drawdown around 12%). However, these levels of extreme risks are substantially lower compared to the maximum drawdown resulting from the use of the fund allocation (close to 50%) or the EW strategy (around 45%). Overall, the FRP strategies display lower extreme risk levels, with reasonably strong performance, resulting in attractive Sharpe ratios. On the other hand, the fund allocation or the EW strategy, which have the highest ENC measures, are by far the most exposed to extreme risk, since their resulting position is essentially loaded on a single factor driving the equity risk. In the top panel of Figure 4, we show the cumulative wealth of each strategy and notice that the wealth generated by the EW portfolio is by much more volatile than the wealth generated by optimized portfolios. As expected, the GMV portfolio leads to the lowest cumulative wealth, but the least volatile one.

The descriptive statistics in Table 5 provide global measures for portfolio allocations over the whole out-of-sample period. In order to better understand the time-series properties of diversification levels, we plot the evolution of the ENB measure through time in the bottom panel of Figure 4. We notice that at each point in time, the two FRP portfolios have the highest ENB measure (the two curves actually coincide since they relate to two FRP portfolios, that is two portfolios with the maximum effective number of bets), while the EW has the lowest one. Other optimized strategies have ENB measures that remain somewhere between these two extremes. In Figures 5 and 6, we analyze for each strategy the evolution over time of the weight in each asset class, the intrinsic contribution of each asset class position to the overall volatility and finally the contribution of each factor exposure to the overall variance. The first row of tables in Figure 5 gives us the results for the EW strategy. In particular, we find that while the maximal diversification is attained in terms of effective number of constituents, this does not translate into a diversification in terms of risk loadings with a portfolio that is mostly exposed to the first factor, which has been interpreted as equity risk. Moving to the second row of figures, showing the results for the ERC strategy, we notice that this portfolio shows (by construction) a perfect diversification in terms of intrinsic risk contribution from each asset, but this does not translate into a well-diversified risk exposure either. In Figure 6, we show the results for the GMV, MSR and both FRP portfolios. In particular, we note that both GMV and MSR portfolios lead to overweighting the Treasury and corporate bond indices. This makes perfect sense since the volatilities of these two classes are the lowest, and their Sharpe ratios are higher than those of the other asset classes. If we look at the second and fourth rows, which relate to the two FRP portfolios, we find very similar patterns as explained above. The most important observation is that the ENB is strictly lower than the number of asset classes due to the presence of long-only constraints.

An analysis of the weight profile for the FRP portfolios shows a relative under-weighting of equity indices and a related over-weighting of bond indices. It also shows a high turnover related in particular to the relative allocation to the two bond indices. Intuitively, this comes from the high correlation between the two asset classes, which leads to a lack of robustness in the extraction of uncorrelated factors via PCA. In this context, it would be desirable to analyze whether more robust results could be obtained with other competing statistical methods that could be used to turn correlated asset returns into uncorrelated factor returns (see the Conclusion for a further discussion of this question).[19]

4.3 Minimum Variance Strategies under ENB and Shortsale Constraints

In this section, we analyze the results of MV portfolios under ENC or ENB constraints. In the presence of estimation errors, adding structure to the optimization problem can indeed lead to improved out-of-sample risk-adjusted performance, as explained in Section 3.3. Table 6 shows the results of GMV portfolios under two types of ENC constraints: $\text{ENC}_2 \geq 3$ and $\text{ENC}_2 \geq 4$; and two types of ENB constraints: $\text{ENB}_1 \geq 3$ and $\text{ENB}_1 \geq 4$. First of all, we notice that higher Sharpe ratios are obtained for ENB-constrained portfolios (0.88 for the MV-ENB3 portfolio and 0.92 for the MV-ENB4 portfolio), compared to what was obtained for the GMV portfolio with long-only constraints, but no constraints on ENB (0.83). Interestingly, we notice that the levels of maximum drawdown are lower for the MV-ENB3 (5.86%) portfolio compared to the GMV itself (5.93%). Overall, this illustration suggests that the combination of ENB constraints and GMV objectives may lead to portfolios with more attractive characteristics (higher Sharpe ratios, lower maximum drawdowns) than the GMV or the FRP-MV portfolios alone. Moreover the level of turnover for ENB-constrained MV portfolios is significantly lower than those obtained in Table 5 for FRP-MV or FRP-MSR portfolios. On the other hand, the addition of ENC constraints to an MV objective results in portfolios with lower Sharpe

[19] In unreported results, we have found that imposing a constant mix allocation within the fixed-income universe would generate a strong reduction in turnover with little impact on risk-adjusted performance and effective number of bets.

ratios and higher maximum drawdown than the GMV portfolio. This is because the ENC constraints shrinks the GMV portfolio towards the EW portfolio, which implies in general a high opportunity cost in terms of Sharpe ratio.

If we now observe the impact of such constraints on the weight evolution in Figure 8, we see that, as expected, imposing ENC constraints leads to positions that are closer to the EW portfolios. In particular we see that the allocations between treasury and corporate bonds are very similar, and the allocation to equities increases compared to the GMV. On the other hand, if we look at the weight profile of ENB-constrained portfolios, they tend to be closer to that of the FRP-GMV portfolio, which is somewhat similar to the weight profile of the GMV portfolio except that the exposure to the uncorrelated factors is more evenly spread out. In particular, we see that the equity risk is more represented towards the end of the backtest for ENB-constrained portfolios than it is for the GMV, but is not dominating as observed for ENC-constrained portfolios.

5 Conclusion

The ability to construct well-diversified portfolios is a challenge of critical importance in the context of designing good proxies for performance-seeking portfolios. A seemingly well-diversified allocation (e.g., EW, GMV, MSR or ERC) to asset classes may well result in a portfolio that is heavily concentrated in terms of factor exposures. In this context, it is of high relevance to measure and manage the effective number of bets in a portfolio. The analysis conducted in this paper suggests that these questions can be addressed within a comprehensive framework by designing portfolio strategies that maximize the effective number of bets, or that minimize risk subject to a target level in terms of minimum effective number of bets.

Our analysis can be extended in a number of ways. In particular, it should be noted that PCA is only one among many possible approaches to extracting factors, and not necessarily the one leading to the most stable results, nor the one leading to the best ease of interpretation. In a companion paper (Meucci et al. (2013)), we explore a competing approach, known as *minimal torsion* approach, for extracting uncorrelated factors from correlated asset returns, which is shown to alleviate the aforementioned concerns raised by the use of PCA. This framework can be applied in contexts where the underlying constituents can be securities (as in the traditional approaches), but also exogenously specified fully general, and possibly correlated, factors. Another useful extension would consider the use of these techniques for the construction of efficient asset class benchmarks, as opposed to asset allocation benchmarks. Finally, that the risk parity approach inevitably leads to a substantial overweighting of bonds versus equities might be a concern in a high-bond-price-low-bond-yield environment. In this context, it would be useful to explore paper more *conditional* risk parity strategies, explicitly designed to optimally respond to changes in market conditions. We leave these questions for further research.

A Proofs of Propositions

A.1 Proof of Proposition 1

We start from the MSR expression given in Equation (3.6) where we replace the vector $\boldsymbol{\mu}$ with the expression given in Proposition 1. Then, the MSR is proportional to:

$$\begin{aligned}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^{j} &= \kappa\,\boldsymbol{\Sigma}^{-1}\left(\boldsymbol{A}'\right)^{-1}\boldsymbol{\Sigma}_F^{\frac{1}{2}}\mathbf{1}_N^{j} \\ &= \kappa\,\left(\boldsymbol{A}'\boldsymbol{\Sigma}\right)^{-1}\boldsymbol{\Sigma}_F^{\frac{1}{2}}\mathbf{1}_N^{j}\end{aligned} \tag{A.1}$$

From the expression of $\boldsymbol{\Sigma}_F = \boldsymbol{A}'\boldsymbol{\Sigma}\boldsymbol{A}$, it is easy to see that $\boldsymbol{A}'\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_F\boldsymbol{A}^{-1}$, leading to:

$$\left(\boldsymbol{A}'\boldsymbol{\Sigma}\right)^{-1} = \boldsymbol{A}\boldsymbol{\Sigma}_F^{-1}.$$

Then, we use the above identity in Equation (A.1) and obtain:

$$\begin{aligned}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}^{j} &= \kappa\,\boldsymbol{A}\boldsymbol{\Sigma}_F^{-1}\boldsymbol{\Sigma}_F^{\frac{1}{2}}\mathbf{1}_N^{j} \\ &= \kappa\,\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{j},\end{aligned}$$

which concludes the proof since it is proportional to the FRP expression given in Equation (3.9).

A.2 Proof of Proposition 2

Let us consider any FRP portfolio $\boldsymbol{w}^{\mathrm{FRP},j}$ that is solution to Problem (3.10). Its Sharpe ratio is given by:

$$\lambda\left(\boldsymbol{w}^{\mathrm{FRP},j}\right) = \frac{\boldsymbol{\mu}'\boldsymbol{w}^{\mathrm{FRP},j}}{\sqrt{\left(\boldsymbol{w}^{\mathrm{FRP},j}\right)'\boldsymbol{\Sigma}\boldsymbol{w}^{\mathrm{FRP},j}}}.$$

We first start with the computation of the variance. Using the following identity $\boldsymbol{\Sigma} = \left(\boldsymbol{A}'\right)^{-1}\boldsymbol{\Sigma}_F\boldsymbol{A}^{-1}$, we can show that:

$$\begin{aligned}\left(\boldsymbol{w}^{\mathrm{FRP},j}\right)'\boldsymbol{\Sigma}\boldsymbol{w}^{\mathrm{FRP},j} &= \frac{1}{\left(\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{j}\right)^2}\times\left(\mathbf{1}_N^{j}\right)'\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\boldsymbol{A}'\boldsymbol{\Sigma}\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{j} \\ &= \frac{1}{\left(\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{j}\right)^2}\times\left(\mathbf{1}_N^{j}\right)'\mathbf{1}_N^{j} \\ &= \frac{N}{\left(\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{j}\right)^2}.\end{aligned}$$

Second, we compute the expected excess return of the FRP portfolio:

$$\begin{aligned}
\boldsymbol{\mu}'\boldsymbol{w}^{\mathrm{FRP},j} &= \frac{1}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j} \times \boldsymbol{\mu}'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j \\
&= \frac{1}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j} \times \left(\boldsymbol{A}'\boldsymbol{\mu}\right)'\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j \\
&= \frac{1}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j} \times \left(\boldsymbol{\mu}_F\right)'\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j \\
&= \frac{1}{\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j} \times \left(\boldsymbol{\lambda}_F\right)'\mathbf{1}_N^j.
\end{aligned}$$

Therefore, we finally obtain a Sharpe ratio equal to:

$$\lambda\left(\boldsymbol{w}^{\mathrm{FRP},j}\right) = \mathrm{sign}\left(\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^j\right)\frac{\left(\boldsymbol{\lambda}_F\right)'\mathbf{1}_N^j}{\sqrt{N}} \leq \frac{\sum_{k=1}^N |\lambda_{F_k}|}{\sqrt{N}}. \tag{A.2}$$

From the Sharpe ratio expression given in Equation (A.2), we see that if we choose the signs of vector $\mathbf{1}_N^j$ such that they coincide with the signs of vector $\boldsymbol{\lambda}_F$, then we have:

$$\lambda\left(\boldsymbol{w}^{\mathrm{FRP\text{-}MSR}}\right) = \mathrm{sign}\left(\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{\mathrm{MSR}}\right)\frac{\sum_{k=1}^N |\lambda_{F_k}|}{\sqrt{N}}. \tag{A.3}$$

Equation (A.3) and inequality (A.2) show that if $\mathbf{1}_N'\boldsymbol{A}\boldsymbol{\Sigma}_F^{-\frac{1}{2}}\mathbf{1}_N^{\mathrm{MSR}} > 0$, then $\boldsymbol{w}^{\mathrm{FRP\text{-}MSR}}$ achieves the highest Sharpe ratio. It is then equal to: $\frac{\sum_{k=1}^N |\lambda_{F_k}|}{\sqrt{N}}$.

A.3 Proof of Proposition 4

First, we rewrite the ENB$_2$-constrained GMV problem (3.13) with respect to an optimization over the weights in the factors:

$$\min_{\boldsymbol{w}_F} \boldsymbol{w}_F'\boldsymbol{\Sigma}_F\boldsymbol{w}_F \quad \text{subject to } \mathbf{1}_N'\boldsymbol{A}\boldsymbol{w}_F = 1, \text{ and } \mathrm{ENB}_2(\boldsymbol{w}) \geq \mathrm{ENB}^*, \tag{A.4}$$

where the ENB-constraint can be made explicit in the factors' weights as:

$$
\begin{aligned}
\mathrm{ENB}_2(\boldsymbol{w}) \geq \mathrm{ENB}^* \quad &\Leftrightarrow \quad \boldsymbol{p}'\boldsymbol{p} \leq \frac{1}{\mathrm{ENB}^*} \Leftrightarrow \sum_{k=1}^{N} \left(w_{F_k}\sigma_{F_k}\right)^4 \leq \frac{\left(\sum_{k=1}^{N} \left(w_{F_k}\sigma_{F_k}\right)^2\right)^2}{\mathrm{ENB}^*} \\
&\Leftrightarrow \quad \frac{\mathrm{ENB}^* - 1}{\mathrm{ENB}^*} \sum_{k=1}^{N} \left(w_{F_k}\sigma_{F_k}\right)^4 - \frac{2}{\mathrm{ENB}^*} \sum_{i \neq j} w_{F_i}^2 w_{F_j}^2 \sigma_{F_i}^2 \sigma_{F_j}^2 \leq 0 \\
&\Leftrightarrow \quad \left(\mathrm{ENB}^* - 1\right) \sum_{k=1}^{N} \left(w_{F_k}\sigma_{F_k}\right)^4 - 2 \sum_{i \neq j} w_{F_i}^2 w_{F_j}^2 \sigma_{F_i}^2 \sigma_{F_j}^2 \leq 0,
\end{aligned}
$$

where the last equivalence holds since $\mathrm{ENB}^* \geq 1 > 0$. Now, if we define matrix $\boldsymbol{M}$ as:

$$
\begin{aligned}
m_{ij} &= -w_{F_i} w_{F_j} \sigma_{F_i}^2 \sigma_{F_j}^2 && \text{if } i \neq j \\
m_{ii} &= \left(\mathrm{ENB}^* - 1\right) w_{F_i}^2 \sigma_{F_i}^4 && \text{otherwise}
\end{aligned}
$$

then we see that

$$
\mathrm{ENB}_2(\boldsymbol{w}) \geq \mathrm{ENB}^* \Leftrightarrow \boldsymbol{w}_F' \boldsymbol{M} \boldsymbol{w}_F \leq 0,
$$

To conclude, it is enough to notice that, at the solution to problem (A.4), a Lagrange multiplier $\nu \geq 0$ exists such that the solution to this problem coincides with the solution to:

$$
\min_{\boldsymbol{w}_F} \boldsymbol{w}_F' \left(\boldsymbol{\Sigma}_F + \nu \boldsymbol{M}\right) \boldsymbol{w}_F \quad \text{subject to } \boldsymbol{1}_N' \boldsymbol{A} \boldsymbol{w}_F = 1.
$$

Using the following identity $\boldsymbol{\Sigma} = \left(\boldsymbol{A}'\right)^{-1} \boldsymbol{\Sigma}_F \boldsymbol{A}^{-1}$, the above optimization problem can be written as an optimization over the weights in the asset classes:

$$
\min_{\boldsymbol{w}} \boldsymbol{w}' \left(\boldsymbol{A}'\right)^{-1} \left(\boldsymbol{\Sigma}_F + \nu \boldsymbol{M}\right) \boldsymbol{A}^{-1} \boldsymbol{w} \quad \text{subject to } \boldsymbol{1}_N' \boldsymbol{w} = 1,
$$

which completes the proof.

B Tables and Figures

Table 1: Descriptive Statistics of Asset Classes

(a) Risk and Performance Statistics

	Tr. Bonds	Co. Bonds	US Eq.	Ex-US Eq.	Private Eq.	Real Est.	Commo.
Av. Ret. (%)	6.33	7.11	8.04	5.58	9.31	10.29	3.37
Volatility (%)	4.63	5.20	17.36	18.25	20.39	22.27	21.55
Sharpe Ratio	0.69	0.76	0.28	0.13	0.30	0.32	0.01
$VaR_{5\%}$ (%)	0.98	1.05	3.78	3.77	4.53	4.17	4.64
$VaR_{1\%}$ (%)	1.60	1.70	6.74	7.54	8.48	9.27	9.02

(b) Correlation in %

	Tr. Bonds	Co. Bonds	US Eq.	Ex-US Eq.	Private Eq.	Real Est.	Commo.
Tr. Bonds	100.00	84.05	-17.87	-18.69	-24.54	-15.88	-9.01
Co. Bonds	84.05	100.00	2.68	5.26	-4.33	1.54	0.83
US Eq.	-17.87	2.68	100.00	73.27	86.49	64.19	22.72
Ex-US Eq.	-18.69	5.26	73.27	100.00	69.40	52.80	34.45
Private Eq.	-24.54	-4.33	86.49	69.40	100.00	73.21	26.06
Real Est.	-15.88	1.54	64.19	52.80	73.21	100.00	18.19
Commo.	-9.01	0.83	22.72	34.45	26.06	18.19	100.00

The top table displays various statistics on the asset classes that are used in the empirical analysis. The average return and volatility are annualized, while both VaR computations are done with weekly returns. The risk-free return for the Sharpe ratio computation is the US T-bill with a 3-month maturity. The bottom table displays correlations between each asset class. All these statistics are computed over a data set that uses 1,069 weekly returns from January 1992 to the end of June 2012.

Figure 1: Volatilities of Asset Classes

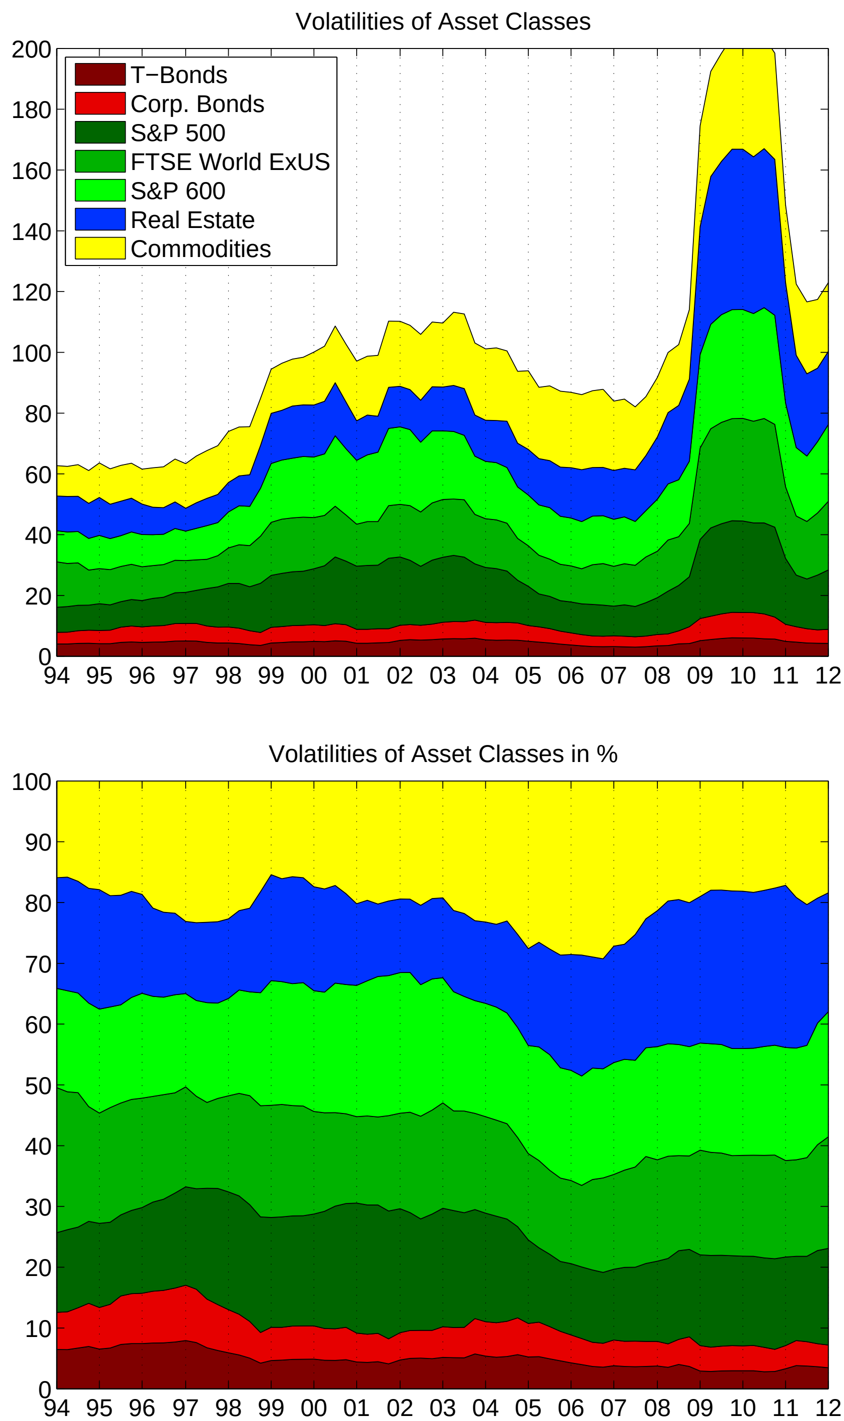

This figure displays the volatilities of the different asset classes estimated over rolling windows of 104 weekly returns (i.e. two years of data) from January 1994 to the end of June 2012. The top panel gives volatilities contribution in values (%) whereas the bottom panel gives the volatilities contribution as a percentage of the total volatility.

Table 2: Factors' Exposures in %

	F_1	F_2	F_3	F_4	F_5	F_6	F_7
Tr. Bonds	-2.94	-0.02	1.30	4.59	48.41	-42.89	76.06
Co. Bonds	0.11	0.28	-0.66	8.77	56.13	-50.69	-64.83
US Eq.	43.61	-11.95	-32.73	-33.27	53.74	53.75	0.36
Ex-US Eq.	42.20	6.67	-53.73	70.82	-14.43	-7.23	3.39
Private Eq.	53.53	-14.13	-13.14	-53.25	-36.96	-50.59	0.51
Real Est.	52.78	-26.89	74.24	29.41	5.42	9.17	0.71
Commo.	25.62	94.29	18.86	-8.84	3.68	2.50	0.29
Total Position (%)	214.91	48.23	-5.86	18.23	115.99	-85.98	16.49
	-	-	-	-	-	-	-
Variance (%)	0.24	0.08	0.04	0.02	0.01	0.01	0.00
Percent Explained (%)	60.84	20.46	9.43	4.94	2.37	1.81	0.15
Cumulative (%)	60.84	81.30	90.73	95.67	98.04	99.85	100.00
Sharpe Ratio	0.29	-0.15	0.12	-0.06	0.60	-0.55	0.02

This table displays the exposures (in %) with respect to asset classes of each factor obtained from running a PCA on the covariance matrix estimated with 1,069 weekly returns from January 1992 to the end of June 2012.

Figure 2: Volatilities of Factors

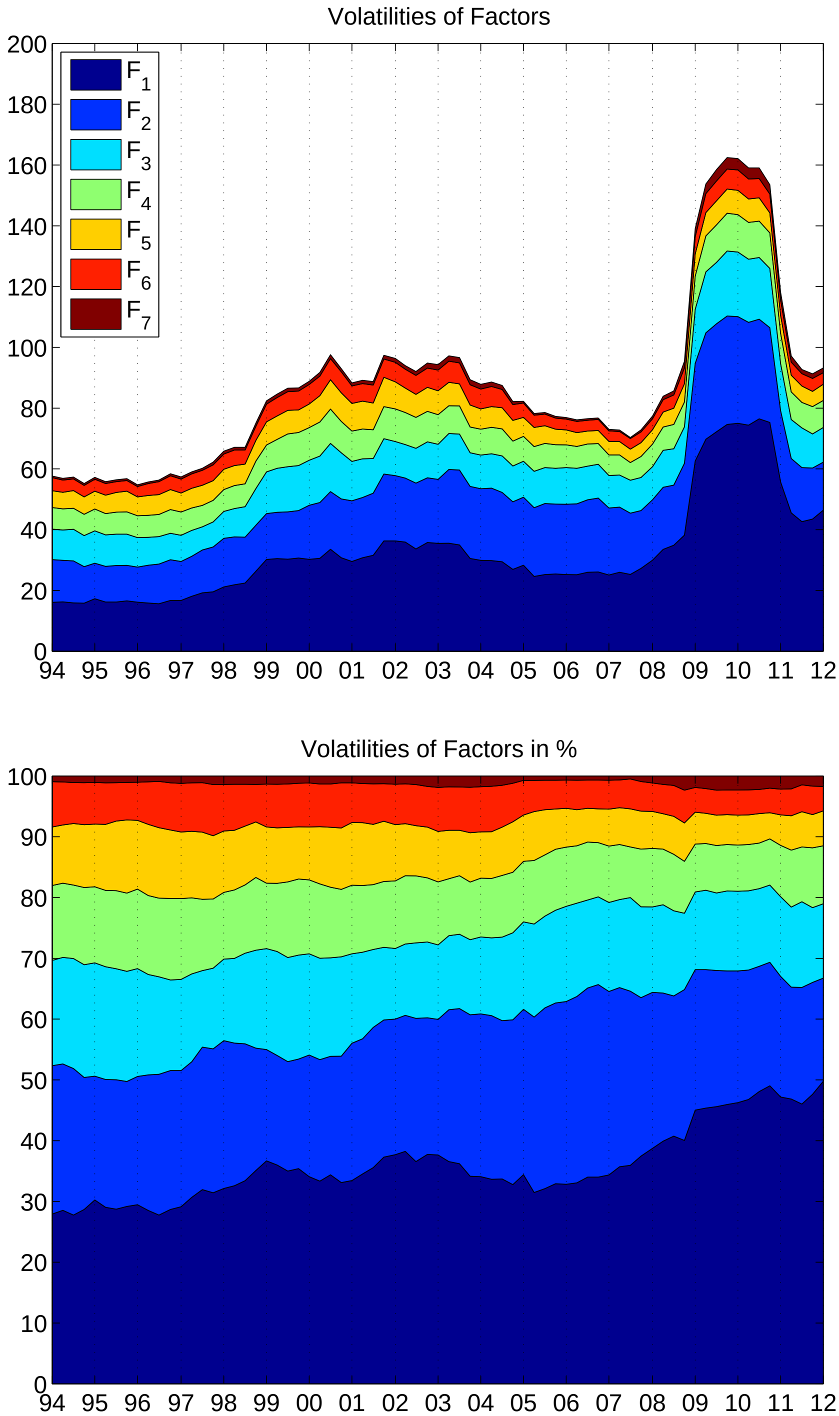

This figure displays the volatilities of the different factors obtained from running a PCA on the covariance matrix estimated over rolling windows of 104 weekly returns (i.e. two years of data) from January 1994 to the end of June 2012. The top panel gives volatilities contribution in values (%) whereas the bottom panel gives the volatilities contribution as a percentage of the total volatility.

Table 3: Allocation and Exposure of the Policy Portfolio

(a) Allocation

	Tr. Bonds	Co. Bonds	US Eq.	Ex-US Eq.	Private Eq.	Real Est.	Commo.
Weight	4%	16%	25%	25%	13%	13%	4%

(b) Exposure and Contribution

	F_1	F_2	F_3	F_4	F_5	F_6	F_7
Exposure: $\boldsymbol{w}_F$	36.20%	-2.84%	-12.97%	7.52%	16.79%	-3.48%	-6.22%
Variance Contribution: $\boldsymbol{p}$	96.69%	0.20%	1.92%	0.34%	0.81%	0.03%	0.01%
Implied Risk Premiums	100%	-4.55%	-14.11%	5.92%	9.16%	-1.66%	-0.85%

The top table displays the allocation of a portfolio representing the policy of a large US state pension fund. The bottom table displays the exposure of the policy portfolio to each factor, together with the variance contribution of each factor. The last line of the table corresponds to implicit views regarding the Sharpe ratios for the factors so as to rationalize the allocation of the policy portfolio displayed in the top table.

Figure 3: Performance of Strategies without Short-Selling Constraints

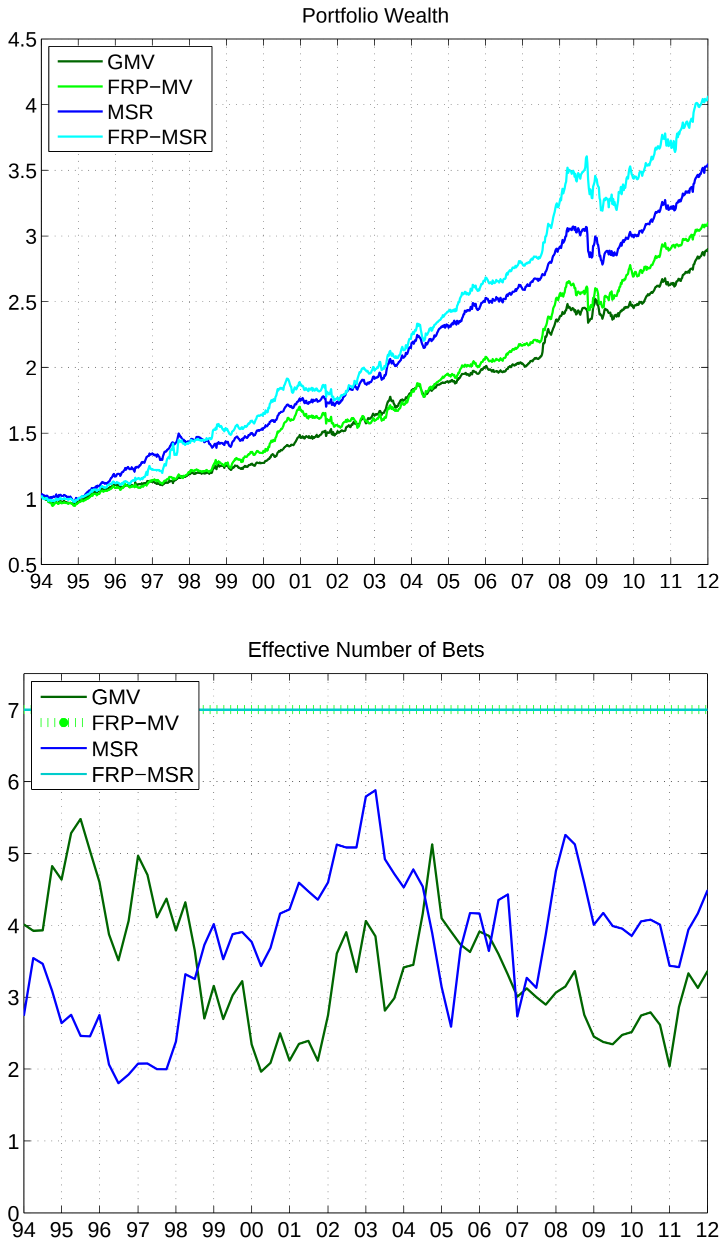

This figure displays the performances of the following strategies: global minimum variance (GMV), factor risk parity with minimum variance (FRP-MV) defined in Proposition 3, maximum Sharpe ratio (MSR) with expected excess-return proportional to the assets volatilities, and factor risk parity with maximum Sharpe ratio (FRP-MSR) defined in Propositions 2. The top panel represents the evolution of wealth between January 1994 to the end of June 2012, while the bottom panel displays the effective number of bets over the same period.

Table 4: Descriptive Statistics of Strategies Without Short-Selling Constraints

(a) Risk and Performance Statistics

	GMV	FRP-MV	MSR	FRP-MSR
Average Return (%)	6.03	6.36	7.14	7.92
Volatility (%)	4.22	5.33	4.83	5.88
Sharpe Ratio	0.68	0.60	0.83	0.81
Max Drawdown (%)	6.37	9.52	9.44	11.48
$\text{VaR}_{5\%}$ (%)	0.78	0.93	0.91	1.07
$\text{VaR}_{1\%}$ (%)	1.47	1.74	1.67	2.11
Average ENC_2	0.86	0.47	1.62	0.43
Average ENB_1	3.40	7.00	3.77	7.00

(b) Correlation in %

	GMV	FRP-MV	MSR	FRP-MSR
GMV	100.00	78.25	75.22	62.16
FRP-MV	78.25	100.00	74.92	67.90
MSR	75.22	74.92	100.00	75.74
FRP-MSR	62.16	67.90	75.74	100.00

The top table displays various statistics on the following strategies: global minimum variance (GMV), factor risk parity with minimum variance (FRP-MV) defined in Proposition 3, maximum Sharpe ratio (MSR) with expected excess-return proportional to the assets volatilities, and factor risk parity with maximum Sharpe ratio (FRP-MSR) defined in Propositions 2. The average return and volatility are annualized, while both VaR computations are done with weekly returns. The risk-free return for the Sharpe ratio computation is the US T-bill with a 3-month maturity. Both the ENC and ENB are averaged over the period starting in January 1994 until the end of June 2012. The bottom panel represents the correlations estimated over the same period of time.

Figure 4: Performance of Strategies with Short-Selling Constraints

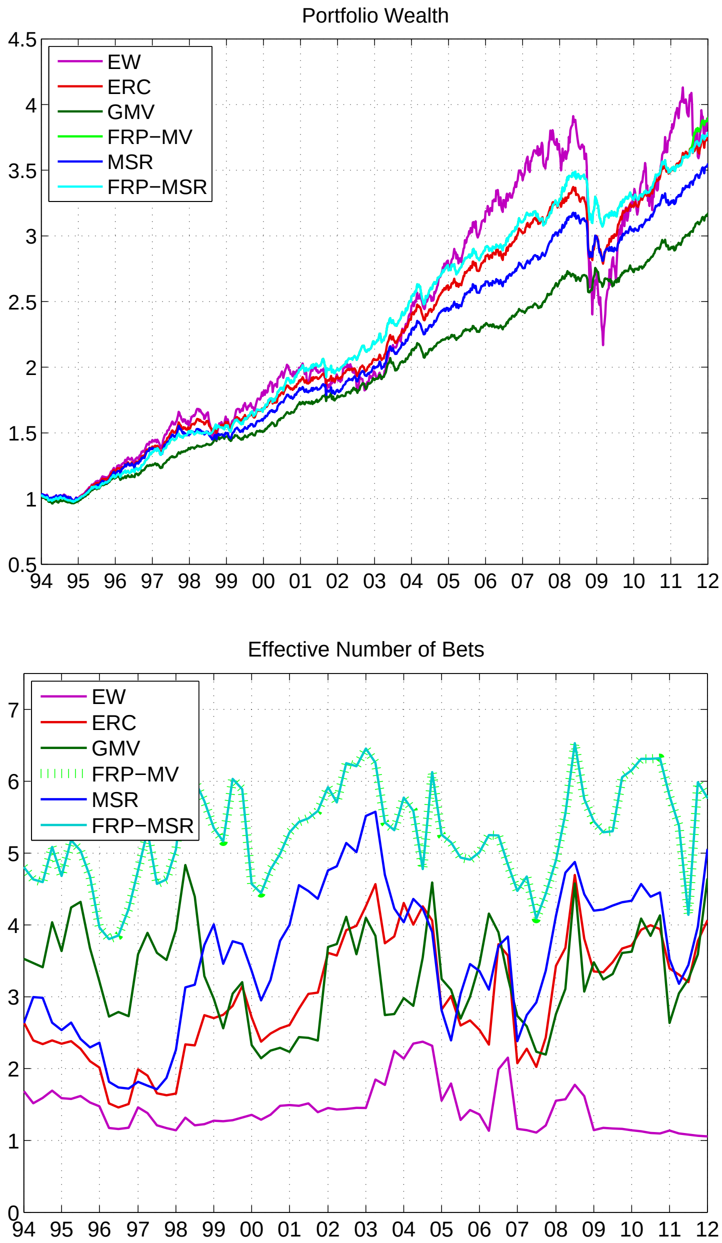

This figure displays the performances of the following strategies: equally weighted (EW), equal risk contribution (ERC), global minimum variance (GMV), factor risk parity with minimum variance (FRP-MV), maximum Sharpe ratio (MSR) with expected excess-return proportional to the assets volatilities, and factor risk parity with maximum Sharpe ratio (FRP-MSR). To obtain portfolios FRP-MV and FRP-MSR, we compute the maximal $\text{ENB}_1^{\max}$, and then numerically solve both the MV and MSR objectives under the constraint that $\text{ENB}_1(\boldsymbol{w}) = \text{ENB}_1^{\max}$. The top panel represents the evolution of wealth between January 1994 to the end of June 2012, while the bottom panel displays the effective number of bets over the same period.

Table 5: Descriptive Statistics of Strategies With Short-Selling Constraints

(a) Risk and Performance Statistics

	Fund	EW	ERC	GMV	FRP-MV	MSR	FRP-MSR
Average Return (%)	7.81	7.87	7.56	6.57	7.66	7.21	7.49
Volatility (%)	13.32	11.20	5.31	4.13	5.04	4.91	5.03
Sharpe Ratio	0.35	0.42	0.83	0.83	0.90	0.83	0.86
Max Drawdown (%)	49.53	44.59	17.32	5.93	12.00	11.58	12.00
$\text{VaR}_{5\%}$ (%)	2.71	2.29	1.01	0.78	0.96	0.91	0.96
$\text{VaR}_{1\%}$ (%)	5.76	4.48	1.74	1.38	1.69	1.57	1.69
Turnover (%)	7.98	10.00	11.57	25.60	88.37	33.92	87.71
Average ENC_2	5.90	7.00	4.39	1.56	1.82	2.64	1.83
Average ENB_1	1.20	1.43	3.02	3.36	5.30	3.60	5.30

(b) Correlation in %

	EW	ERC	GMV	FRP-MV	MSR	FRP-MSR
EW	100.00	80.60	38.61	56.93	69.27	58.88
ERC	80.60	100.00	79.84	84.62	95.59	85.32
GMV	38.61	79.84	100.00	82.47	83.67	82.70
FRP-MV	56.93	84.62	82.47	100.00	86.65	99.51
MSR	69.27	95.59	83.67	86.65	100.00	87.13
FRP-MSR	58.88	85.32	82.70	99.51	87.13	100.00

The top table displays various statistics on the following strategies: equally weighted (EW), equal risk contribution (ERC), global minimum variance (GMV), factor risk parity with minimum variance (FRP-MV), maximum Sharpe ratio (MSR) with expected excess-return proportional to the assets volatilities, and factor risk parity with maximum Sharpe ratio (FRP-MSR). To obtain portfolios FRP-MV and FRP-MSR, we compute the maximal $\text{ENB}_1^{\max}$, and then numerically solve both the MV and MSR objectives under the constraint that $\text{ENB}_1(\boldsymbol{w}) = \text{ENB}_1^{\max}$. The average return and volatility are annualized, while both VaR computations are done with weekly returns. The risk-free return for the Sharpe ratio computation is the US T-bill with a 3-month maturity. The turnover is an annual one-way measure, and both the ENC and ENB are averaged over the period starting in January 1994 until the end of June 2012. The bottom panel represents the correlations estimated over the same period of time.

Figure 5: Weights and Risk Decompositions of EW and ERC Strategies with Short-Selling Constraints

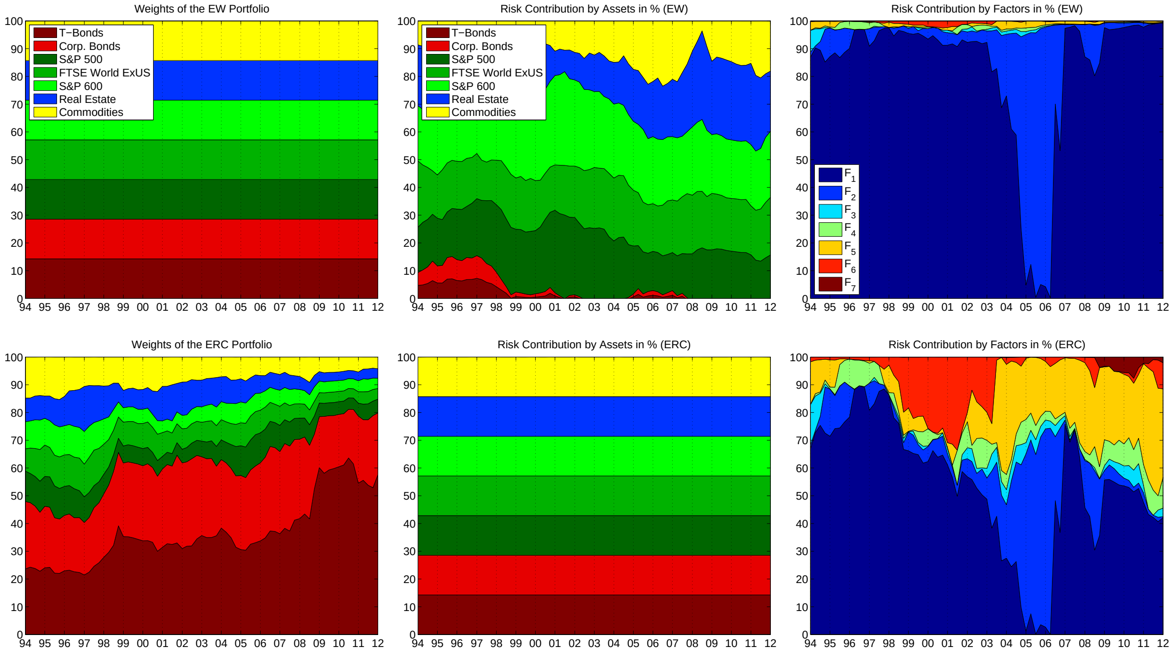

This figure displays the weights in each asset class (left panels) together with the risk contribution of each asset class (middle panels), and the risk contribution of each factor (right panels), both expressed in percentage of the total portfolio variance. The portfolio strategies considered are: equally weighted (EW) and equal risk contribution (ERC).

Figure 6: Weights and Risk Decompositions of GMV, MSR and FRP Strategies with Short-Selling Constraints

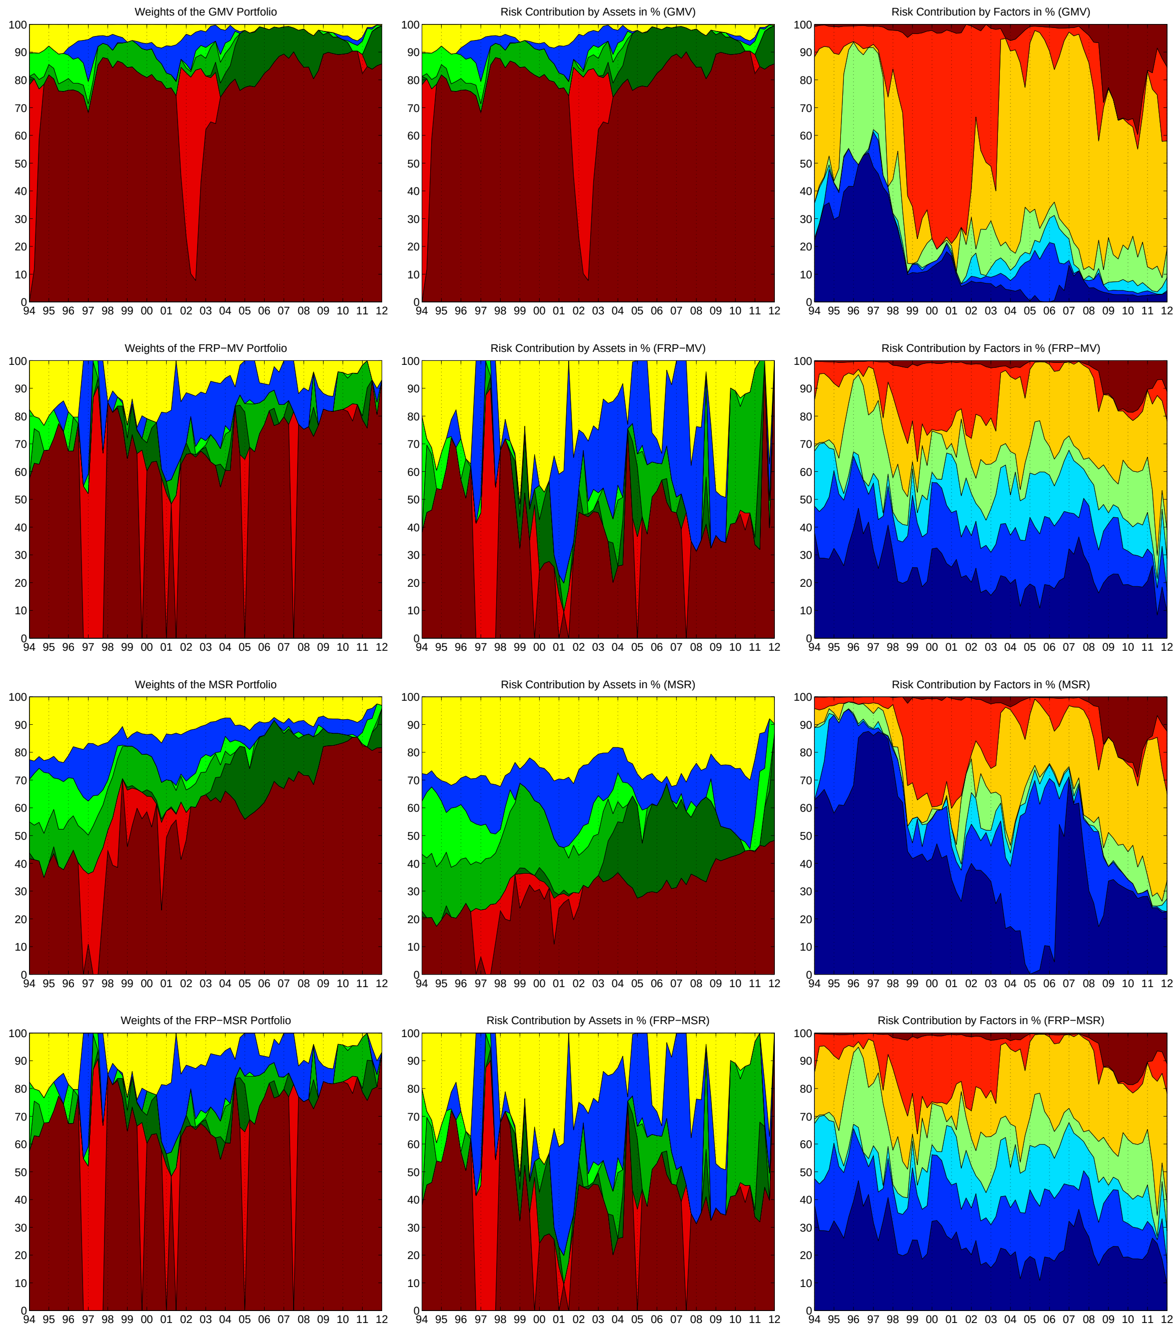

This figure displays the weights in each asset class (left panels) together with the risk contribution of each asset class (middle panels), and the risk contribution of factor (right panels), both expressed in percentage of the total portfolio variance. The portfolio strategies considered are: global minimum variance (GMV), factor risk parity with minimum variance (FRP-MV), maximum Sharpe ratio (MSR) with expected excess-return proportional to the assets volatilities, and factor risk parity with maximum Sharpe ratio (FRP-MSR). To obtain portfolios FRP-MV and FRP-MSR, we compute the maximal $\mathrm{ENB}_1^{\max}$, and then numerically solve both the MV and MSR objectives under the constraint that $\mathrm{ENB}_1(\boldsymbol{w}) = \mathrm{ENB}_1^{\max}$.

Figure 7: Performance of MV Strategies with Short-Selling Constraints and Either ENC or ENB Constraints

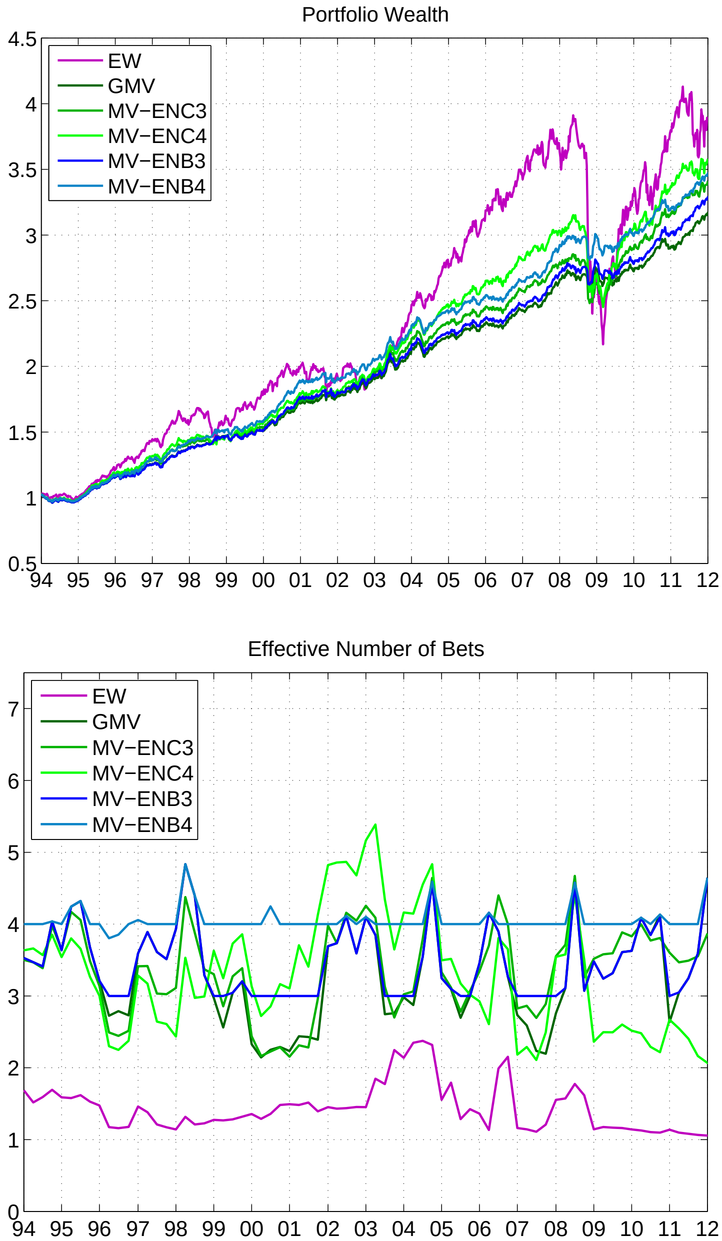

This figure displays the performances of the following strategies: equally weighted (EW), global minimum variance (GMV), minimum variance with $\text{ENC}_2 \geq 3$ and $\text{ENC}_2 \geq 4$ constraints (denoted respectively MV-ENC3 and MV-ENC4), and minimum variance with $\text{ENB}_1 \geq 3$ and $\text{ENB}_1 \geq 4$ constraints (denoted respectively MV-ENB3 and MV-ENB4). The top panel represents the evolution of wealth between January 1994 to the end of June 2012, while the bottom panel displays the effective number of bets over the same period.

Table 6: Descriptive Statistics of MV Strategies with Short-Selling Constraints and Either ENC or ENB Constraints

(a) Risk and Performance Statistics

	EW	GMV	MV-ENC3	MV-ENC4	MV-ENB3	MV-ENB4
Average Return (%)	7.87	6.57	7.02	7.32	6.78	7.09
Volatility (%)	11.20	4.13	4.52	5.50	4.13	4.28
Sharpe Ratio	0.42	0.83	0.86	0.76	0.88	0.92
Max Drawdown (%)	44.59	5.93	13.54	22.22	5.86	6.05
$\text{VaR}_{5\%}$ (%)	2.29	0.78	0.84	1.08	0.77	0.80
$\text{VaR}_{1\%}$ (%)	4.48	1.38	1.39	1.88	1.38	1.46
Turnover (%)	10.00	25.60	17.62	13.59	30.41	44.76
Average ENC_2	7.00	1.56	3.00	4.00	1.60	1.67
Average ENB_1	1.43	3.36	3.40	3.26	3.50	4.08

(b) Correlation in %

	EW	GMV	MV-ENC3	MV-ENC4	MV-ENB3	MV-ENB4
EW	100.00	38.61	63.90	85.38	40.07	45.46
GMV	38.61	100.00	92.22	76.23	99.67	97.06
MV-ENC3	63.90	92.22	100.00	93.67	92.46	92.23
MV-ENC4	85.38	76.23	93.67	100.00	77.19	79.91
MV-ENB3	40.07	99.67	92.46	77.19	100.00	97.73
MV-ENB4	45.46	97.06	92.23	79.91	97.73	100.00

The top table displays various statistics on the following strategies: equally weighted (EW), global minimum variance (GMV), minimum variance with $\text{ENC}_2 \geq 3$ and $\text{ENC}_2 \geq 4$ constraints (denoted respectively MV-ENC3 and MV-ENC4), and minimum variance with $\text{ENB}_1 \geq 3$ and $\text{ENB}_1 \geq 4$ constraints (denoted respectively MV-ENB3 and MV-ENB4). The average return and volatility are annualized, while both VaR computations are done with weekly returns. The risk-free return for the Sharpe ratio computation is the US T-bill with a 3-month maturity. The turnover is an annual one-way measure, and both the ENC and ENB are averaged over the period starting in January 1994 until the end of June 2012. The bottom panel represents the correlations estimated over the same period of time.

Figure 8: Weights and Risk Decompositions of MV Strategies with Short-Selling Constraints and Either ENC or ENB Constraints

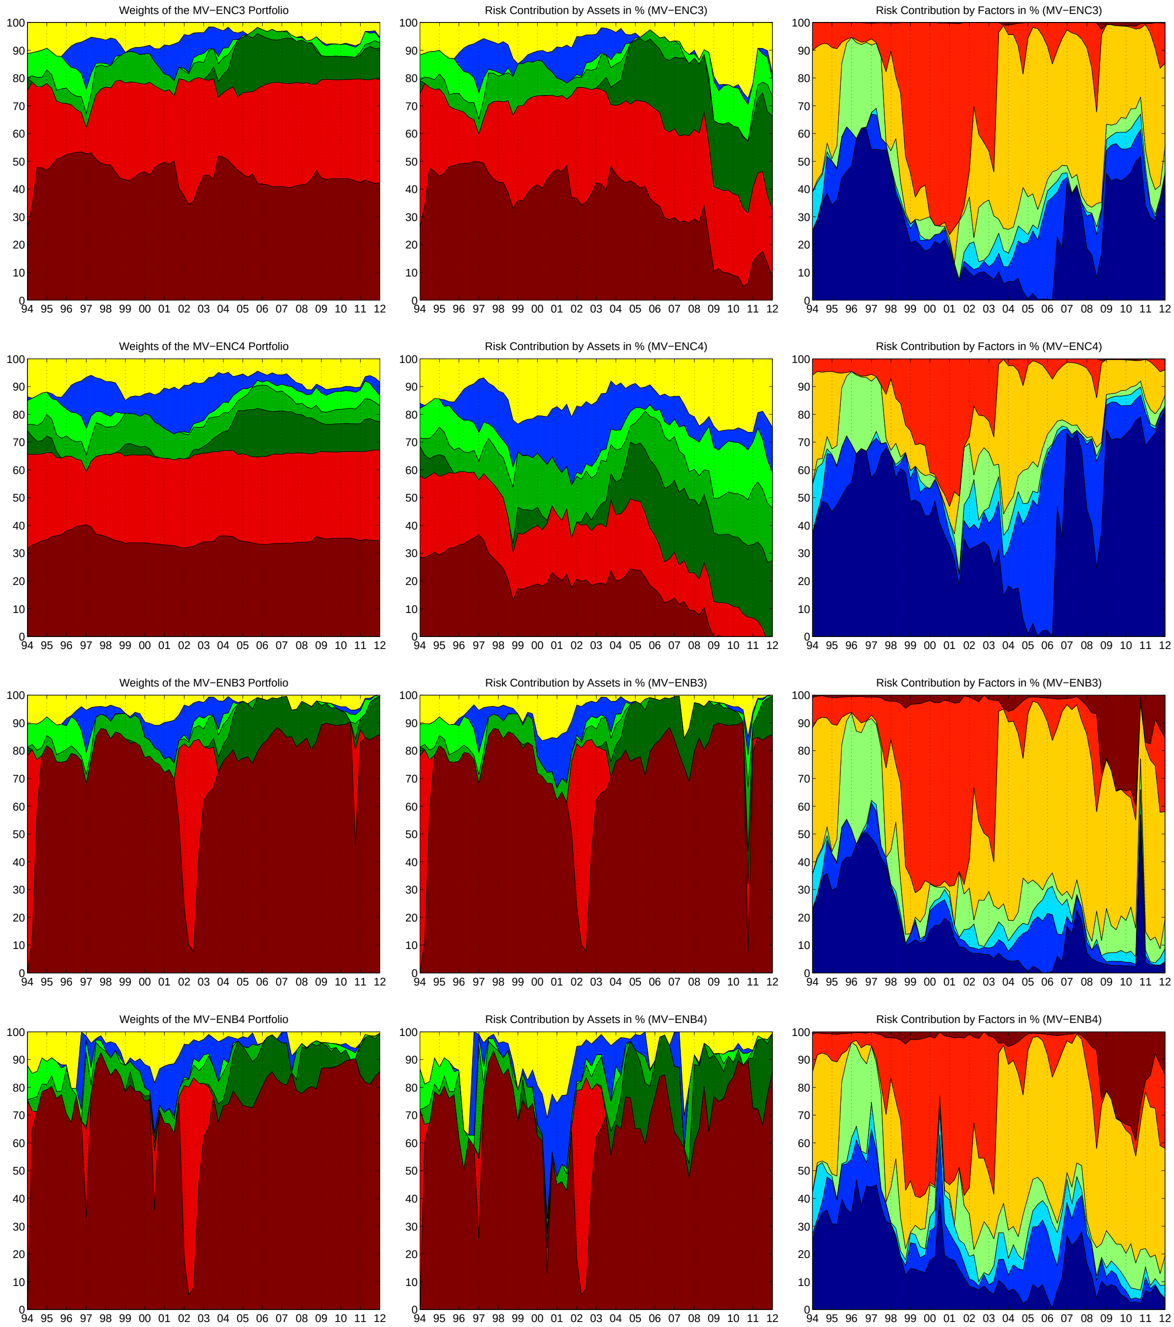

This figure displays the weights in each asset class (left panels) together with the risk contribution of each asset class (middle panels), and the risk contribution of factor (right panels), both expressed in percentage of the total portfolio variance. The portfolio strategies considered are: minimum variance with $\text{ENC}_2 \geq 3$ and $\text{ENC}_2 \geq 4$ constraints (denoted respectively MV-ENC3 and MV-ENC4), and minimum variance with $\text{ENB}_1 \geq 3$ and $\text{ENB}_1 \geq 4$ constraints (denoted respectively MV-ENB3 and MV-ENB4).

References

Amenc, N., L. Martellini, J. Meyfredi, and V. Ziemann (2010). Passive hedge fund replication–beyond the linear case. *European Financial Management 16*(2), 191–210.

Ang, A., W. Goetzmann, and S. Schaefer (2009). Evaluation of active management of the norwegian government pension fund–global. *Report to the Norwegian Ministry of Finance*.

Black, F. and R. Litterman (1992). Global portfolio optimization. *Financial Analysts Journal*, 28–43.

Black, F. and R. B. Litterman (1991). Asset allocation. *The Journal of Fixed Income 1*(2), 7–18.

Bruder, B. and T. Roncalli (2012). Managing risk exposures using the risk budgeting approach. *Available at SSRN 2009778*.

DeMiguel, V., L. Garlappi, F. Nogales, and R. Uppal (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science 55*(5), 798–812.

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.

Evans, J. L. and S. H. Archer (1968). Diversification and the reduction of dispersion: An empirical analysis. *Journal of Finance 23*(5), 761–767.

Goetzmann, W., L. Li, and K. Rouwenhorst (2005). Long-term global market correlations. *The Journal of Business 78*(1), 1–38.

Hasanhodzic, J. and A. Lo (2007). Can Hedge-Fund Returns Be Replicated?: The Linear Case. *Journal of Investment Management 5*(2).

Jagannathan, R. and T. Ma (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance 58*(4), 1651–1683.

Jurczenko, E., T. Michel, and J. Teiletche (2013). Generalized risk-based investing. *Available at SSRN 2205979*.

Ledoit, O. and M. Wolf (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance 10*(5), 603–621.

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

Lee, W. (2011). Risk-based asset allocation: A new answer to an old question? *Journal of Portfolio Management 37*(4), 11–28.

Lohre, H., U. Neugebauer, and C. Zimmer (2012). Diversified risk parity strategies for equity portfolio selection. *The Journal of Investing 21*(3), 111–128.

Lohre, H., H. Opfer, and G. Orszag (2011). Diversifying risk parity. *Available at SSRN 1974446*.

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.

Merton, R. (1971). Optimal portfolio and consumption rules in a continuous-time model. *Journal of Economic Theory 3*(4), 373–413.

Merton, R. (1973). An intertemporal capital asset pricing model. *Econometrica 41*(5), 867–887.

Merton, R. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics 8*(4), 323–361.

Meucci, A. (2009). Managing diversification. *Risk 22*(5), 74–79.

Meucci, A., A. Santangelo, and R. Deguest (2013). Measuring portfolio diversification based on optimized uncorrelated factors. SSRN: http://ssrn.com/abstract=2276632.

Poddig, T. and A. Unger (2012). On the robustness of risk-based asset allocations. *Financial Markets and Portfolio Management 26*(3), 369–401.

Roncalli, T. (2013). *Introduction to Risk Parity and Budgeting.* Chapman & Hall, CRC Financial Mathematics Series.

Roncalli, T. and G. Weisang (2012). Risk parity portfolios with risk factors. *Available at SSRN 2155159*.

Ross, S. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory 13*(3), 341–360.

Statman, M. (1987). How many stocks make a diversified portfolio? *Journal of Financial and Quantitative Analysis*, 353–363.