Risk-Budgeted Mean-Variance Portfolios

Bernardo Freitas Paulo de Costa[1], Raul Riva[2], and Rodrigo Targino[1]

[1]*School of Applied Mathematics, Getulio Vargas Foundation (FGV EMAp)*
[2]*Brazilian School of Economics and Finance, Getulio Vargas Foundation (FGV EPGE)*

This version: April 6, 2025

Abstract

We introduce the Risk-Budgeted Mean-Variance (RBMV) portfolio, a novel framework that connects the classical Markowitz mean-variance problem and the risk budgeting approach. By modifying the risk budgeting optimization problem to include constraints on expected returns and volatility, RBMV offers a disciplined way to manage the trade-off between risk concentration and return maximization. The investor gains a lever to adjust how close the portfolio sits to either framework, depending on her preferences. We show that the optimization problem that defines the RBMV portfolio is convex, efficiently computable, and typically delivers competitive returns with reduced risk concentration in the context of long-only portfolios. We illustrate our methodology using daily equity returns from the U.S. and show that our methodology efficiently controls the volatility of returns while also delivering Sharpe ratios that are consistently higher than the traditional mean-variance approach.

Keywords Portfolio Allocation, Risk Parity, Risk Budgeting, Mean-Variance, Volatility;
JEL Classification G1, G11;

1 Introduction

While several quantitative investment strategies have been proposed the academic literature as well as in industry, one has stood up the test of time and reigns supreme: the Markowitz mean-variance portfolio. This portfolio is the cornerstone of modern portfolio theory and has been the subject of countless academic papers and books. The Markowitz portfolio is the solution to an optimization problem that aims to maximize the expected return of a portfolio while minimizing its volatitly. Although theorectically sound, the Markowitz portfolio has some drawbacks. One of the most important is that it tends to concentrate the portfolio on a few assets with high Sharpe ratios, which makes the portfolio's overall risk heavily dependent on what happens to a few assets. Another drawback is how sensitive the optimal portfolio is to small changes in the inputs parameters, which is especially important as the expected returns and the covariance matrix have to be estimated.

To circumvent the issue of risk concentration, Maillard et al. (2010) and Roncalli (2013) introduced the heuristic idea of forcing the investor to define a budget for the risk of each individual asset in the overall portfolio. These authors define the *risk contribution* of an asset as the fraction of the total portfolio risk assigned to each asset. A *risk budget portfolio* is then constructed by ensuring the risk contribution of each asset matches a predefined proportion of the total portfolio risk. In the seminal paper Maillard et al. (2010), the authors show that the risk budget portfolio is the solution to an optimization problem where the investor minimizes the portfolio risk subject to a seemingly arbitraty constraint: that the sum of the product between the pre-specified risk contributions and the log of the exposures should be non-negative. This optimization problem is convex and can be solved efficiently (see e.g. Freitas Paulo da Costa et al. (2023)). Whilst the risk budget portfolio naturally diversifies the risk of the portfolio among its assets and its less sensitive to estimation error (see discussion in Section 2.3.2, below), it is not necessarily the portfolio with the highest expected return nor the portfolio with the highest Sharpe ratio.

To blend the benefits of both approaches, we propose a novel portfolio optimization framework (named the Risk Budgeted Mean-Variance Portfolio) that interpolates between the Markowitz mean-variance and a risk budgeting portfolio, naturally nesting these two frameworks. Our methodology allows investors to balance the trade-off between maximizing expected returns and diversifying risk contributions across assets. This is done by modifying the risk budget portfolio optimization problem by including constraints on both the expected return and the volatility of the resulting portfolio. As added constraints are also convex, we mantain the convexity of the optimization problem, which ensures that the resulting portfolio is computationally efficient to obtain.

We take our methodology to the data and provide an application focusing on the construction of long-only equity portfolios for the U.S. market. Using more than 30 years of daily data from CRSP, we show that our methodology, in contrast to the traditional Mean-Variance framework, generates portfolios with lower volatility, lower portfolio concentration, and higher Sharpe ratios. In constrast to the risk budgeting framework, our methodology generates portfolios with higher returns. Importantly, we show how tuning the parameters of our methodology enables the investor to have either a portfolio closer to the one generated by Mean-Variance or a portfolio closer to one generated by risk budgeting.

The remainder of this paper is organized as follows. In Section 2 we review the basic Mean-Variance and Risk Budgeting portfolios and thoroughly discuss the empirical differences observed between both, including the consequences of adding estimation error to the population moments. In Section 3, we formally introduce the RBMV framework, discuss some of its properties, and compare it with the portfolio obtained by simply averaging Mean-Variance and Risk Budgeting portfolio weights. Section 4 presents a detailed empirical study on the performance of the RBMV portfolio for the U.S. stock market. Section 5 concludes the paper.

1.1 Literature review

1.2 Extensions to the Mean-Variance Approach

Since the seminal work of Markowitz (1952), the mean-variance portfolio has inspired a vast body of research. Many extensions aim to address key practical challenges, such as sensitivity to estimation error and excessive concentration. One important theoretical extension is the two-fund separation theorem introduced by Tobin (1958), which shows that under certain assumptions, investors can achieve optimal portfolios by combining a risk-free asset with a tangency portfolio of risky assets. Despite its elegance, this result relies on strong assumptions and has limited direct applicability.

In practice, one of the most pressing issues is the instability of the Markowitz solution when inputs are estimated from historical data. Several authors have proposed incorporating robust estimators to alleviate this problem. Jorion (1986) introduced a Bayes-Stein shrinkage approach to improve expected return estimates, while the Black-Litterman model (Black and Litterman, 1991) combined prior market equilibrium returns with investor views to generate more stable and intuitive expected returns. Ledoit and Wolf (2003) proposed a shrinkage estimator for the covariance matrix, now widely used in portfolio construction. DeMiguel et al. (2009) showed that simple heuristics like the equally weighted portfolio often outperform optimized portfolios out-of-sample due to estimation error. In response, Michaud (1989) introduced resampled efficient frontiers, emphasizing the importance of accounting for estimation uncertainty.

To reduce portfolio concentration, Jagannathan and Ma (2003) demonstrated that imposing norm constraints can improve out-of-sample performance by effectively regularizing the solution. Alternative risk measures have also been proposed. Konno and Yamazaki (1991) introduced the mean-absolute deviation model, replacing variance with a linear objective. Cardinality and transaction cost constraints were addressed in Chang et al. (2000), Bienstock (1996), and Lobo et al. (2007). These extensions, while useful, often lead to non-convex optimization problems, raising computational challenges and reducing scalability.

1.3 Extensions to the risk budgeting approach

The idea of allocating portfolio weights based on risk contributions was formalized in Maillard et al. (2010) and extended in Roncalli (2013), who introduced the concept of risk budgeting. These models aim to construct portfolios where each asset contributes a predefined fraction of the portfolio's total risk. This approach naturally leads to better risk diversification, as no single asset dominates the total portfolio risk. Moreover, risk budgeting portfolios are known by practitioners to be less sensitive to estimation error and remain tractable due to their convex optimization formulation. Mathematical properties of this investment strategy, such as existence and uniqueness, have been studied in Freitas Paulo da Costa et al. (2023) and Cetingoz et al. (2024).

A large body of work has explored alternative risk measures within the risk budgeting framework. Bruder et al. (2016) and Jurczenko and Teiletche (2019) derived closed-form expressions for Expected Shortfall (ES) risk contributions under Gaussian assumptions, enabling ES-based budgeting. Ji and Lejeune (2018) incorporated downside risk, Bellini et al. (2021) studied expectile risk measures, and Anis and Kwon (2022) added cardinality constraints to control portfolio sparsity. Freitas Paulo da Costa et al. (2023) proposed a cutting-planes-based stochastic optimization algorithm for computing risk budgeting portfolios under general convex risk measures, capable of handling distributions where closed-form risk contributions are unavailable.

Other researchers focused on structural extensions. Bai et al. (2016) proposed a least-squares formulation for risk parity, while Haugh et al. (2017) explored budgeting risk across overlapping groups of assets in conjunction with return objectives. Risk factor budgeting—where diversification is applied to risk factors rather than individual assets—has been developed by Meucci et al. (2015), Roncalli and Weisang (2016), and Lassance et al. (2022), the latter introducing independent component analysis to improve factor orthogonality. More recently, Cetingoz and Guéant (2025) proposed a framework that balances the risk contributions from both assets and factors. Li et al. (2022) and Pesenti et al. (2024) extended risk parity into multi-period settings, highlighting its potential in dynamic environments.

Finally, algorithmic innovations have enhanced the applicability of these models. Mausser and Romanko (2018) proposed convex optimization methods for ES-based risk parity using discrete loss distributions. Alternatively, Griveau-Billion et al. (2013), Chaves et al. (2012), and Spinu (2013) developed efficient algorithms for variance-based risk budgeting. These formulations—many rooted in or inspired by the original logarithmic formulation of Maillard et al. (2010)—continue to evolve, expanding the toolbox of practitioners seeking robust and diversified portfolio solutions under uncertainty.

The literature reviewed in this section demonstrates that while many extensions of the mean-variance approach offer improvements, they typically address only one aspect of the problem: either estimation error, diversification, or tractability.

Despite these innovations, few approaches offer a unified framework that balances the trade-off between expected return maximization and risk diversification. This gap motivates the development of models that can interpolate between the mean-variance and alternative risk-based allocations in a computationally efficient way.

2 Mean-Variance and Risk Budgeting

This section briefly reviews the mean-variance and the risk budget portfolio problems. We start with basic notation. We use uppercase letters to indicate random variables, and bold symbols to denote a vector of variables with the same name, so that $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$, for example. The investor can allocate resources across $d \in \mathbb{N}$ assets with current prices $\boldsymbol{p} \equiv (p_1, \ldots, p_d)$, and corresponding future (random) prices $\boldsymbol{P} \equiv (P_1, \ldots, P_d)$.

Even though portfolio allocation problems are generally described in terms of portfolio proportions, the risk budgeting problem is more easily described in terms of dollar amounts. That will also be true for our methodology since we connect the traditional mean-variance framework to the risk budgeting problem. Hence, we develop our setup assuming the investor has a total *dollar* amount v_0 to invest across the risky assets. His problem will consist of picking a vector of dollar amounts $\boldsymbol{v} = (v_1, v_2, ..., v_d)$ to invest. We further assume that $v_i \geq 0$, i.e., the investor will build a long-only portfolio. The risk budgeting problem, and as a consequence, our methodology, is not well-defined when short positions are allowed.

After choosing the dollar amount v_i to be allocated on asset i, this investor acquires v_i/p_i shares of this asset, which will have a value of $v_i \cdot \frac{P_i}{p_i}$ in the future. We call v_i the *exposure* of the portfolio to the i^{th} asset, and $w_i \equiv v_i/v_0 \geq 0$ the portfolio *weights*, that is, the proportion of the budget invested into asset i, so that $\sum_{i=1}^{d} w_i = 1$.

The investor's dollar *return* on a portfolio with exposure $\boldsymbol{v}$ is the random variable given by

$$R(\boldsymbol{v}) \equiv \sum_{i=1}^{d} v_i \frac{P_i}{p_i} - v_0 = \sum_{i=1}^{d} v_i \left(\frac{P_i}{p_i} - 1 \right) = v_0 \left[\sum_{i=1}^{d} w_i \left(\frac{P_i}{p_i} - 1 \right) \right] . \tag{1}$$

We further assume that $\mu \in \mathbb{R}^d$ and $\Sigma \in \mathbb{R}^{d \times d}$ are the expected returns and the covariance matrix of asset returns, respectively. We assume that Σ is positive-definite. In our setup, the investor knows these population moments. In our empirical application, they will be estimated using daily returns.

2.1 The Mean-Variance Portfolio Problem

The traditional long-only mean-variance portfolio problem consists in finding the portfolio weights $\boldsymbol{w}$ that maximizes the expected return $\mu(R(\boldsymbol{w})) \equiv \mu^\top \boldsymbol{w}$ while minimizing the portfolio variance

$\sigma^2(R(\boldsymbol{w})) \equiv \boldsymbol{w}^\top \Sigma \boldsymbol{w}$. One mathematical formulation for the portfolio selection problem is:

$$\begin{aligned} \min_{\boldsymbol{w} \geq 0} \quad & \sigma(R(\boldsymbol{w})) \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^{d} w_i = 1 \\ & \mu(R(\boldsymbol{w})) \geq \mu_{\min}^{MV}, \end{aligned} \tag{2}$$

where $\mu_{\min}^{MV}$ is the minimum expected return accepted by investor. The first constraint corresponds to the intuitive idea that the w_i's are indeed portfolio weights, while the second is the *target return* constraint. We also could have cast this problem as a return maximization program under a volatility constraint. The formulation above, however, will be more convenient when we introduce our methodology.

If, instead, the investor is interested in the portfolio exposures $\boldsymbol{v} = v_0 \boldsymbol{w}$, substituting $\boldsymbol{w} := \boldsymbol{v}/v_0$ in the Markowitz portfolio problem yields

$$\begin{aligned} \min_{\boldsymbol{v} \geq 0} \quad & \tfrac{1}{v_0} \sigma(R(\boldsymbol{v})) \\ \text{s.t.} \quad & \tfrac{1}{v_0} \textstyle\sum_{i=1}^{d} v_i = 1 \\ & \tfrac{1}{v_0} \mu(R(\boldsymbol{v})) \geq \mu_{\min}^{MV}. \end{aligned} \tag{3}$$

A positive constant multiplying the objective function will not change the optimal solution. Moreover, we can multiply both sides of the constraints by v_0 and have the equivalent optimization problem:

$$\begin{aligned} \min_{\boldsymbol{v} \geq 0} \quad & \sigma(R(\boldsymbol{v})) \\ \text{s.t.} \quad & \textstyle\sum_{i=1}^{d} v_i = v_0 \\ & \mu(R(\boldsymbol{v})) \geq \mu_{\min}^{MV} \cdot v_0. \end{aligned} \tag{4}$$

The first constraint implies that we have a *fully invested* portfolio, from where we easily interpret v_0 on its right-hand side as the total wealth of the investor. Note that the expected return constraint, when formulated as a function of the exposures $\boldsymbol{v}$, must also include a v_0 factor, since the left-hand side is homogeneous of degree one in $\boldsymbol{v}$. Furthermore, if we rescale the total wealth by a factor $\gamma > 0$, the corresponding portfolio exposures will be scaled by γ as well.

We denote by $\boldsymbol{v}_{MV}^*$ the solution to problem (4). The return rate corresponding to this allocation is denoted $\mu_{MV} = \frac{1}{v_0} \cdot \mu(R(\boldsymbol{v}_{MV}^*))$. Usually, this is equal to the parameter $\mu_{\min}^{MV}$ in problem (4) and will differ only when $\mu_{\min}^{MV}$ is chosen as something smaller than the return rate of the minimum variance portfolio. Likewise, we define $\sigma_{MV} = \frac{1}{v_0} \cdot \sigma(R(\boldsymbol{v}_{MV}^*))$.

2.2 The Risk Budgeting Portfolio Problem

One important drawback of the mean-variance approach is concentrating the portfolio on a few assets with high Sharpe ratios (the ratio between expected returns and volatility). This is an undesirable feature from the risk management perspective since it will make the portfolio's overall risk heavily dependent on what happens to a few assets. Another drawback is how sensitive the optimal portfolio is to small changes in μ and Σ (see, for instance, Li (2015) for a modern review). Small changes in these parameters might induce sharp changes to optimal weights.

To circumvent the issue related to risk concentration, Maillard et al. (2010) and Roncalli (2013) introduced the heuristic idea of forcing the investor to define a budget for the maximum risk each individual asset can expose the overall portfolio to. This budget is a safeguard against overly concentrating overall risk. Before defining the risk budgeting portfolio problem, we define the risk contribution of individual assets.

Definition 1. *The* risk contribution *of asset i to the total portfolio risk $\sigma(R(\boldsymbol{v}))$, which we denote by $\mathcal{RC}_i(\boldsymbol{v})$, is given by:*

$$\mathcal{RC}_i(\boldsymbol{v}) \equiv v_i \cdot \frac{\partial \sigma(R(\boldsymbol{v}))}{\partial v_i} \tag{5}$$

Using this definition, we see that the risk contribution of a certain asset to the overall risk can be high for either of two reasons: 1) there is a large exposure to asset i, i.e., v_i is high; 2) an increment in the exposure v_i will cause a large increase in $\sigma(R(\boldsymbol{v}))$, which obviously depends on the whole correlation structure of asset returns.[1]

Since $\sigma(R(\boldsymbol{v}))$ is a homogeneous function of degree one when seen as a function of $\boldsymbol{v}$, Euler's theorem for homogeneous functions implies that we can write:

$$\sigma(R(\boldsymbol{v})) = \sum_{i=1}^{d} \mathcal{RC}_i(\boldsymbol{v}). \tag{6}$$

In this sense, it is natural to interpret $\mathcal{RC}_i(\boldsymbol{v})$ as how much asset i contributes to the portfolio's total risk. We also note that $\mathcal{RC}_i(\boldsymbol{v})$ is a homogeneous function of degree one. To define the risk budgeting problem, we assume that the investor chooses a set of d proportions $0 < b_i < 1$ that represent how much of the overall risk of the portfolio each asset can command. We call $\boldsymbol{b} = (b_1, \ldots, b_d)$ the *risk budget*. We also assume $\sum_i^d b_i = 1$. We are now ready to define the risk budgeting portfolio.

Definition 2 (The Risk Budgeting Portfolio)**.** *For a risk budget $\boldsymbol{b} = (b_1, \ldots, b_d)$ and a total endowment v_0, the* risk budgeting *(RB) portfolio is defined by a vector of exposures $\boldsymbol{v}$ that satisfies*

$$b_i \cdot \sigma(R(\boldsymbol{v})) = \mathcal{RC}_i(\boldsymbol{v})\,, \qquad \textit{for all} \quad i = 1, \ldots, d\,, \tag{7}$$

and ensures that $v_0 = \sum_{i=1}^{d} v_i$.

We highlight that if some $\boldsymbol{v}$ satisfies (7), any other vector $\gamma \cdot \boldsymbol{v}$ with $\gamma > 0$ will also satisfy the same condition. It is the requirement of full investment that identifies the scale of the risk budgeting portfolio.

[1] Even though our discussion uses the standard deviation of (dollar) returns as the chosen risk measure, this definition can be applied for any coherent risk measure (in the sense of Artzner et al. (1999)) that is homogeneous of degree one as a function of the exposures $\boldsymbol{v}$. In fact, our whole methodology can be extended to these more general risk measures. We chose not to pursue this more general problem because the connection with the traditional mean-variance setup would not be so clear.

One example of a risk budgeting portfolio is the *risk parity portfolio*, in which every asset has equal risk contribution, that is $\boldsymbol{b} \equiv (\frac{1}{d}, \ldots, \frac{1}{d})$. In this case, even though assets are allowed to have different volatilities, each asset will contribute equally to the overall risk. As another example, assume $d = 3$ the risk budget $\boldsymbol{b} = (0.5, 0.3, 0.2)$. This means that the investor targets a portfolio in which the assets have risk contributions of 50%, 30%, and 20%, respectively.

Now, we further characterize the risk budgeting portfolio. The next proposition, whose proof we delegate to the appendix, is crucial for our methodology since it hints at how to compute $\boldsymbol{v}$ in practice.

Lemma 1. *Given a budget $\boldsymbol{b}$ and a total endowment v_0, the risk budgeting portfolio is the portfolio with exposure $\boldsymbol{v}$ that satisfies*

$$b_i \, \mathcal{RC}_j(\boldsymbol{v}) = b_j \, \mathcal{RC}_i(\boldsymbol{v}) \,, \qquad \text{for all} \quad i, j = 1, \ldots, d \,, \tag{8}$$

$$v_0 = \sum_{i=1}^{d} v_i \,. \tag{9}$$

This lemma shows that the correct $\boldsymbol{v}$ that ensures risk budgeting is the solution of a non-linear system of equations. However, for large d, obtaining $\mathcal{RC}_i(\boldsymbol{v})$ in closed form is cumbersome and, in fact, could be impossible if we had picked a more involved risk measure other than the standard deviation. Following an insight from Maillard et al. (2010), we propose computing $\boldsymbol{v}$, up to a scaling factor, as the solution of an optimization problem that will generate the conditions in Lemma 1 as first-order conditions of a convex optimization problem. This is introduced by the next proposition, which is proven below.

Proposition 1. *Given a risk budget $\boldsymbol{b}$, any optimal solution $\boldsymbol{v}^*$ to*

$$\min_{\boldsymbol{v} \in \mathbb{R}^d_+} \sigma(R(\boldsymbol{v})) \,, \qquad \text{subject to} \quad \sum_{i=1}^{d} b_i \log(v_i) \geq 0 \tag{10}$$

is proportional to the exposure $\boldsymbol{v}$ of the RB portfolio for risk budget $\boldsymbol{b}$.

Proof. We define the Lagrangian function

$$J(\boldsymbol{v}, \lambda) \equiv \sigma(R(\boldsymbol{v})) - \lambda \sum_{i=1}^{d} b_i \log(v_i) \,, \tag{11}$$

where the parameter $\lambda \geq 0$ corresponds to the Lagrange multiplier. Taking derivatives with respect to v_i, $i = 1, \ldots, d$, we obtain

$$\partial_{v_i} J(\boldsymbol{v}^*, \lambda) = \frac{\partial \sigma(R(\boldsymbol{v}^*))}{\partial v_i} - \lambda \frac{b_i}{v_i^*} \,. \tag{12}$$

Imposing the first order condition yields, for all $i = 1, \ldots, d$,

$$\lambda = \frac{v_i^*}{b_i} \cdot \frac{\partial \sigma(R(\boldsymbol{v}^*))}{\partial v_i} = \frac{1}{b_i} \mathcal{RC}_i(\boldsymbol{v}^*), \qquad \text{for all} \quad i = 1, \ldots, d \,. \tag{13}$$

Hence, the optimal $\boldsymbol{v}^*$ fulfills (8). By homogeneity, we may define $\boldsymbol{v} \equiv C\boldsymbol{v}^*$ so that the exposure $\boldsymbol{v}$ also satisfies $v_0 = \sum_{i=1}^{d} v_i$. By Lemma 1, this yields the desired risk budgeting portfolio.

Conversely, if we are given strictly positive exposures v_i such that the conditions in (8) hold, then all ratios $\frac{\mathcal{RC}_i(v)}{B_i}$ are equal. We can scale the v_i's using a factor C such that $x_i \equiv Cv_i$ and $\sum_i b_i \log(x_i) = 0$, yielding a vector $\boldsymbol{x}$ at the boundary of the feasible region. By homogeneity of the risk contribution, we still have that all ratios $\frac{\mathcal{RC}_i(x)}{b_i}$ are equal (scaled up by C), so we can define a positive λ via Equation (13). Given such λ, the x_i's are feasible and satisfy the first order conditions. Since the optimization problem is convex and has a strictly feasible interior point, the first-order conditions are both necessary and sufficient, implying that $\boldsymbol{x}$ is a solution of (10). $\square$

Proposition 1 provides a powerful numerical result for the characterization of $\boldsymbol{v}$. It identifies the ray along $\boldsymbol{v}^*$ that will contain the correct exposure $\boldsymbol{v}$ the risk budgeting problem looks for. The exact risk budgeting portfolio will be a rescaled version of $\boldsymbol{v}^*$ such that the portfolio is fully invested. Formally, the risk budgeting portfolio is $\boldsymbol{v} = \frac{v_0}{\sum_i^d v_i^*} \cdot \boldsymbol{v}^*$.

From the optimization point of view, the objective function is strictly convex, and the constraint is the upper-contour set of a strictly concave function. This will imply that numerical algorithms can efficiently solve this problem and that the solution $\boldsymbol{v}^*$, if it exists, it will be unique. Hence, given a budget $\boldsymbol{b}$ and endowment v_0, we have also proven that the risk budget portfolio $\boldsymbol{v}$, if it exists, it is unique.[2]

Similarly to the mean-variance optimization problem, we denote by $\boldsymbol{v}_{RB}^*$ the solution to problem (10). The return rate and volatility corresponding to this allocation are denoted, respectively, by $\mu_{RB} \equiv \frac{1}{\sum_{i=1}^d v_{RB,i}^*}\mu(R(\boldsymbol{v}_{RB}^*))$ and $\sigma_{RB} \equiv \frac{1}{\sum_{i=1}^d v_{RB,i}^*}\sigma(R(\boldsymbol{v}_{RB}^*))$.

2.3 Comparing Both Approaches

We now contrast both approaches. The mean-variance portfolio delivers the portfolio with the lowest overall risk among the portfolios with high enough expected returns. The price to pay for such optimality is typically the concentration of the portfolio on a few assets and a portfolio that is highly sensitive to small changes in inputs parameters, which is especially important in a context where μ and Σ have to be estimated.

On the other hand, the risk budgeting framework allows the investor to have full control over the distribution of overall risk across different assets. The price to pay for such control is, in general, a portfolio with lower returns when compared with a mean-variance portfolio with the same volatility. To make this trade-off clear, we provide a numerical example with calibrated parameters. Then, we introduce moment estimation errors via a simulation exercise.

We assume $d = 5$ and use the following population moments for the mean returns μ, the individ-

[2] Freitas Paulo da Costa et al. (2023) shows that, in fact, the constraint will be binding at the optimal solution.

ual standard deviations s and correlation matrix C:

$$\mu = \begin{bmatrix} 0.5 \\ 0.12 \\ 0.09 \\ 0.05 \\ 0.15 \end{bmatrix}, \quad s = \begin{bmatrix} 0.10 \\ 0.20 \\ 0.15 \\ 0.08 \\ 0.13 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0.20 & 0.40 & 0.25 & 0.50 \\ 0.20 & 1 & -0.20 & 0.40 & 0.6 \\ 0.40 & -0.20 & 1 & -0.10 & 0.30 \\ 0.25 & 0.40 & -0.10 & 1 & 0.30 \\ 0.50 & 0.60 & 0.30 & 0.30 & 1 \end{bmatrix}. \tag{14}$$

We further define $\Sigma \equiv s^\intercal C s$.

2.3.1 No Estimation Error

With these population moments at hand, we compute the Markowitz efficient frontier, defined as the points in the (σ, μ)-plane that denote the expected returns and volatility of mean-variance portfolios. This is represented by the dashed line in Figure 1. Given the moments above, it is not feasible to have any portfolios above this frontier, by definition.

Figure 1: We show the different portfolios in the (σ, μ)-plane, computed using the calibrated parameters from (14). Both expected returns and volatility are measured in percentages. We only consider long-only portfolios.

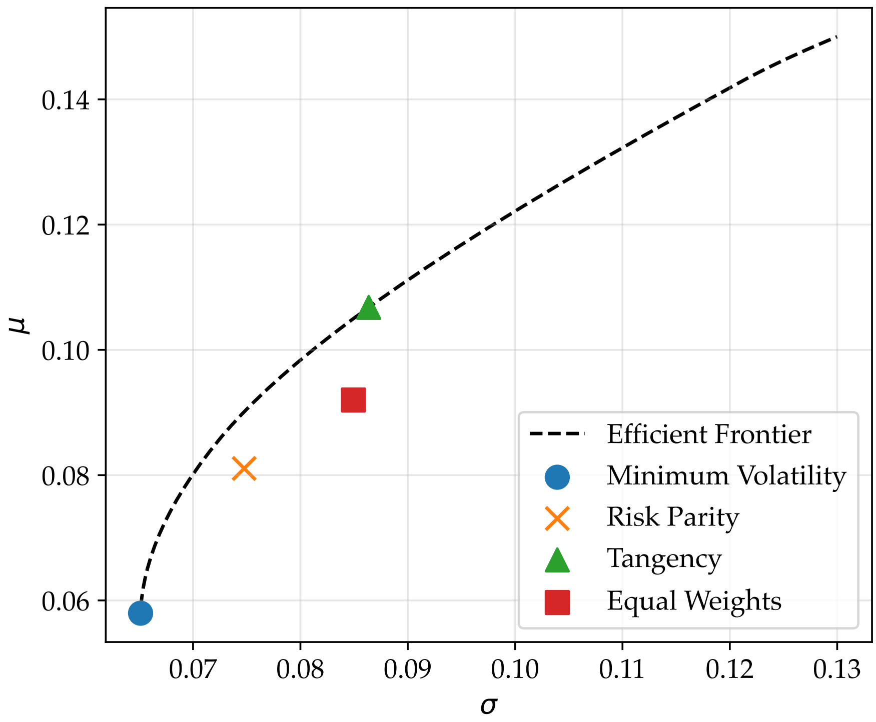

Figure 1 also shows the location of four portfolios. The blue dot represents the minimum volatility portfolio, defined as the portfolio with the lowest possible volatility across all possible portfolios. The diagonal cross represents the risk parity portfolio, which is the risk budgeting portfolio where all assets contribute equally to overall risk, i.e., $\boldsymbol{b} = (0.2, 0.2, 0.2, 0.2, 0.2)$. The triangle locates the tangency portfolio, which is defined as the point along the frontier with the highest ratio between expected return and volatility (the Sharpe ratio). Finally, the square represents the portfolio with equal weights on all assets, identified by $\boldsymbol{w} = (0.2, 0.2, 0.2, 0.2, 0.2)$.

Table 1: For each of the portfolios from Figure 1, we report their expected return, volatility, Sharpe ratio (measured as the ratio between expected return and volatility), the Gini index of associated weights, and the Gini index of risk contributions.

Portfolio	Return	Volatility	Sharpe Ratio	Gini Index (w_i)	Gini Index ($\mathcal{RC}_i$)
Equal weights	0.092	0.085	1.083	0.000	0.238
Minimum volatility	0.058	0.065	0.890	0.568	0.568
Tangency	0.107	0.086	1.236	0.534	0.625
Risk parity	0.081	0.075	1.084	0.184	0.000

We highlight two important features in Figure 1. First, we note that the risk parity portfolio is not identical to the equal weights one. In the risk budgeting framework, even if we require risk parity, weights are not necessarily equal across assets exactly because different assets will have different volatilities, and these returns are not assumed to be independent. Second, the risk parity portfolio is strictly *inside* the frontier. This means that there exist portfolios that dominate it in the mean-variance sense. However, these dominating portfolios will not feature any other particularly desirable characteristics, such as risk parity. The cost of risk budgeting is exactly the loss of mean-variance optimality. In other words, risk budgeting, in general, will push portfolios toward the interior of the frontier.

Table 1 illustrates these points by computing different portfolio summary statistics. For instance, by design, the tangency portfolio has the highest Sharpe ratio, even though the risk parity portfolio is not that far behind. This table also shows one of the important drawbacks of the mean-variance framework, which is the concentration of weights on a few assets. In the fourth column, we report the Gini index for the associated weights of these portfolios.[3] This is a concentration measure that ranges from zero to one. When it is zero, it means that all weights are equal to $1/d$, so the portfolio is evenly distributed. Alternatively, it measures one when 100% of a portfolio is concentrated in a single asset. Table 1 makes it clear that the risk parity portfolio has a much lower Gini index than the tangency portfolio. Going inside the frontier implies a necessarily lower Sharpe ratio, but it enables much lower concentration in terms of portfolio weights.

The last column in Table 1 reports the Gini index of $\mathcal{RC}_i$ for the different portfolios. The tangency portfolio has the highest concentration of risk using that measure, while the risk parity portfolio has the lowest by design. The last column also helps illustrate that a portfolio with evenly distributed weights will not have, in general, evenly distributed risk contributions.

Figure 2 further documents this tendency of the mean-variance approach to concentrate risk. The panel on the left computes the Gini index for all portfolios along the efficient frontier and compares this measure with the Gini index for the risk parity portfolio. Along the frontier, the concentration measure is almost always in the $[0.5, 0.8]$ range, uniformly higher than the concentration measure for the risk parity portfolio.[4]

[3] Given a vector $\boldsymbol{x} = (x_1, ..., x_n)$, we define the Gini index of $\boldsymbol{x}$ as $\frac{\sum_{i=1}^{n}\sum_{j=1}^{n}|x_i - x_j|}{2n\sum_{i=1}^{n}x_i}$.

[4] The Gini index is also a commonly used measure of wealth inequality. As a reference of absolute magnitudes, the Gini index for South Africa, ranked as the most unequal country in the World Bank's database, has a Gini index of 0.63.

Figure 2: In panel (a), we show the Gini index of portfolio weights for all the portfolios along the efficient frontier. We also show, as a reference, the Gini index of weights for the Risk Parity portfolio as a dashed line. In panel (b), we show the risk contributions $\mathcal{RC}_i$ for each of the d assets and each of the portfolios.

(a) Portfolio weight Gini indices

(b) Risk Contribution Gini indices

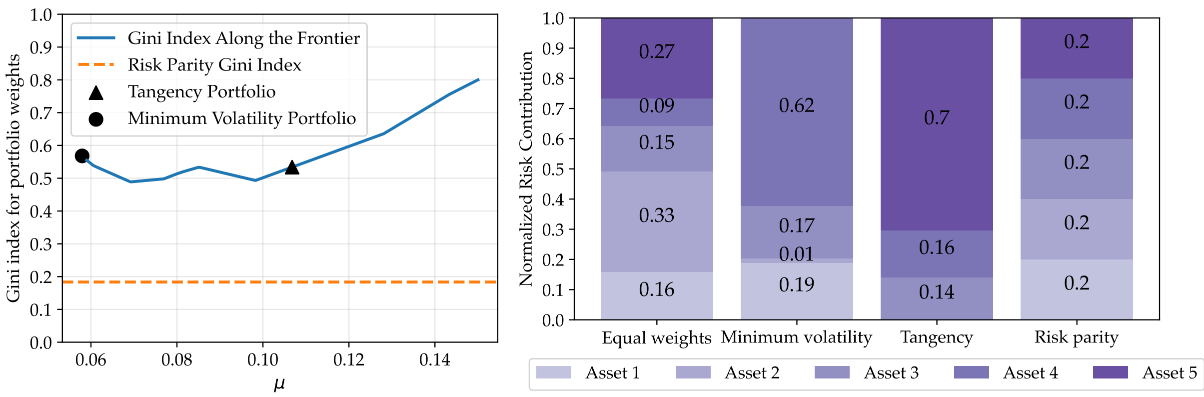

The panel on the right shows the risk contributions $RC_i(v)$ for each asset and portfolio. We normalize these contributions as a fraction of the overall risk. As expected, the risk parity portfolio has $RC_i(v)/\sigma(R(v)) = \frac{1}{5}$ across the five assets. On the other hand, the risk contributions for the tangency portfolio are very different across assets. Asset 5 contributes roughly 70% of the overall risk of that portfolio, indicating a strong level of concentration of the overall risk. In fact, even the equally weighted portfolio displays some level of concentration. In that case, Asset 2 makes up a third of the overall risk, while Asset 4 accounts for less than 10% of the overall standard deviation.

Mean-variance and risk budgeting are different allocation frameworks with different strengths and weaknesses. This first exercise without estimation error shows the tendency mean-variance portfolios have in concentrating weights and the overall risk in a few assets. Risk budgeting is an explicit safeguard against this issue. The cost of such a safeguarding mechanism is the lack of mean-variance optimality.

2.3.2 Adding Estimation Error

Our discussion so far assumed that the portfolio manager knows μ and Σ. Now, we study the role of estimation error and how these two frameworks will behave in a more challenging scenario. We conduct the following procedure 500 times:

1. We simulate 252 returns from $N(\mu, \Sigma)$ – which is the equivalent of a trading year;
2. With these returns, we compute the sample average and the sample covariance matrix, denoted by $\widehat{\mu}$ and $\widehat{\Sigma}$, respectively;

Brazil has a Gini index of 0.52, while the U.S. Gini index is around 0.41, and Sweden is around 0.29. See the data at https://data.worldbank.org/indicator/SI.POV.GINI.

Figure 3: In panel (a), each of the dots represents the volatility and expected return of a portfolio constructed using either the Mean-Variance (blue) or the Risk Parity (orange) methodologies based on population moments estimated using simulated data from a Gaussian distribution parametrized as (14). The black diagonal crosses represent the theoretical locations of the Mean-Variance portfolio with $\sigma_{max} = 0.1$ and of the Risk Parity portfolio. In Panel (b), we report the histogram of the realized Gini index for portfolio weights.

(a) Simulated portfolios in the (σ, μ)-plane

(b) Distribution of realized Gini indices for $\boldsymbol{w}_i$

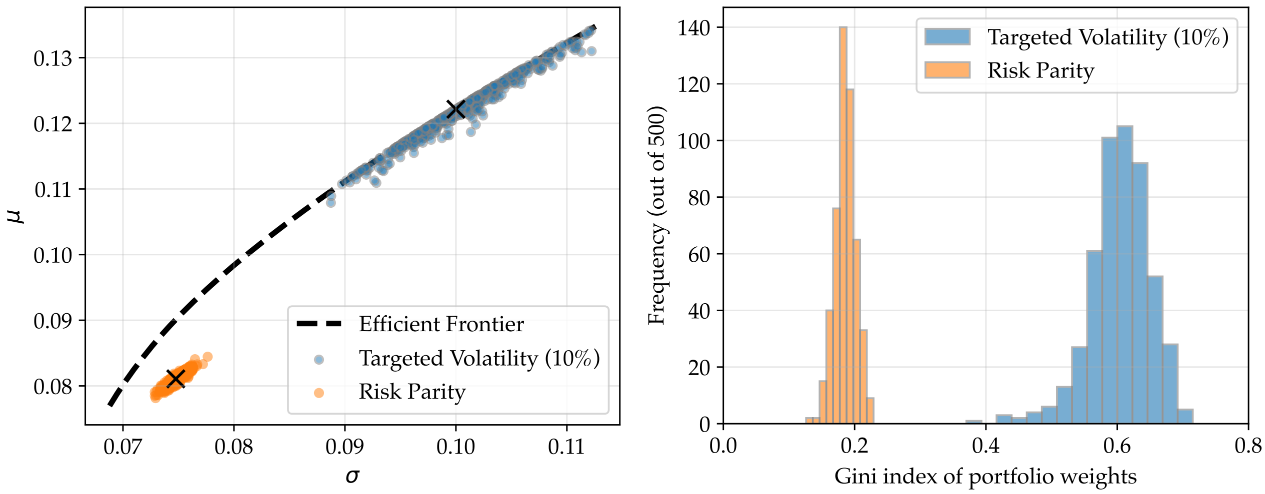

3. Given $\widehat{\mu}$ and $\widehat{\Sigma}$, we compute the weights of the risk parity portfolio and the weights for a mean-variance portfolio with maximum volatility of 10%, which we denote by $\boldsymbol{w}_{RP}$ and $\boldsymbol{w}_{10\%}$;

4. We compute these portfolios' true expected returns and volatility, which depend on μ and Σ, and not on estimated moments.

Each of these simulated samples generates two dots on the left panel of Figure 3: one around the true risk parity portfolio and another one around the mean-variance portfolio whose true volatility is 10%. Hence, each of these clouds is composed of 500 dots, and the efficient frontier is constructed with the population parameters. We use black diagonal crosses to identify the true risk parity portfolio and the true mean-variance portfolio with 10% volatility.

The most striking fact from these simulations is that the cloud around the risk parity portfolio is much more concentrated around its center than the cloud around the mean-variance portfolio with a volatility target of 10%. Since the mean-variance portfolios typically concentrate weights (and risk) in assets with the highest *estimated* Sharpe ratios, estimation error will be particularly detrimental for this framework. Conversely, the risk parity framework handles estimation errors with more success. Once again, the cost of such robustness is generally lower returns.

The panel on the right from Figure 3 shows the associated histograms for the realized Gini indices of $\boldsymbol{w}_{RP}$ and $\boldsymbol{w}_{10\%}$, computed after each of the 500 simulations. We see that the risk parity approach systematically generates less concentrated weights (lower Gini) than the mean-variance portfolio, leading to better robustness against estimation error.

Overall, this section showed that mean-variance optimality implies the cost of concentrating risk, leading to portfolios sensitive to estimation error. Spreading that risk through risk budgeting is one possible solution, but one with its own costs: no control over expected returns and typically a lower Sharpe ratio. The risk budgeting framework, specialized here to risk parity for the sake of an example, was also more robust to estimation error of population parameters.

3 Our Methodology

Now, we connect the two frameworks we described and present our methodology. Our motivation is to give the investor the ability to have, at the same time, some control over both individual risk contributions and the minimum expected return a portfolio can achieve.

There are two ways to view our methodology. First, it is a regularized version of the mean-variance approach in which the optimal exposures will be tilted to incorporate the idea that the investor might want to spread the overall risk across assets. Another way of understanding it is a tilt in the risk budgeting portfolio that will give up perfect budgeting in favor of higher expected returns. Our portfolio construction framework is defined by the following optimization problem below.

Definition 3 (The Risk Budgeted Mean-Variance Portfolio)**.** *Given a risk budget $\boldsymbol{b}$, an endowment v_0, a minimum required expected return μ_{min}, and a maximum volatility bound σ_{max}, the Risk Budgeted Mean-Variance Portfolio (RBMV) is given by $\boldsymbol{v} = \frac{v_0}{\sum_{i=1}^d v_i^*} \cdot \boldsymbol{v}^*$, where $\boldsymbol{v}^*$ is the solution to the following optimization program if it exists:*

$$\begin{aligned} \min_{\boldsymbol{v}\in\mathbb{R}^d_+} \quad & \sigma(R(\boldsymbol{v})) && && (15)\\ s.t. \quad & \textstyle\sum_{i=1}^d b_i \log(v_i) \geq 0 && [\lambda_v] \\ & \mu(R(\boldsymbol{v})) \geq \mu_{\min} \textstyle\sum_{i=1}^d v_i && [\lambda_\mu] \\ & \sigma(R(\boldsymbol{v})) \leq \sigma_{\max} \textstyle\sum_{i=1}^d v_i\,, && [\lambda_\sigma] \end{aligned}$$

where we have indicated Lagrange multipliers for each constraint in brackets. The corresponding portfolio weights are given by $\boldsymbol{w} = \frac{1}{v_0}\boldsymbol{v}$.

Similarly to the pure risk budgeting problem, the solution $\boldsymbol{v}^*$ to the program above does not need to satisfy $\sum_{i=1}^d v_i^* = v_0$. But we can normalize the solution for the problem above such that $\boldsymbol{v} = \frac{v_0}{\sum_{i=1}^d v_i^*} \cdot \boldsymbol{v}^*$ is a fully invested factor of exposures. This will also scale the risk contributions, but the *ratio* of these risk contributions across assets will remain the same.

The first constraint in this problem forces the individual risk contributions to be as close as possible to the desired budget $\boldsymbol{b}$. The second constraint ensures that the portfolio delivers an expected return of at least $\mu_{\min}$. The last constraint further ensures that seeking higher returns while also maintaining some proximity to the desired risk budget does not lead to excessive volatility.

Analyzing the Lagrangian function of this problem is instructive. If we let $\mathcal{L}(\boldsymbol{v}; \lambda_v, \lambda_\mu, \lambda_\sigma)$ be the

Figure 4: On the left, we show the path created by our methodology (diagonal crosses) as we vary μ_{min}. The triangles represent the expected returns and volatilities of portfolios created by linearly interpolating between the weights provided by the Mean-Variance approach and the Risk Parity framework. On the right, we compute the L_2-distance between the realized risk contributions along the RBMV path and the desired risk budget $\boldsymbol{b}$ as a function of μ_{min}.

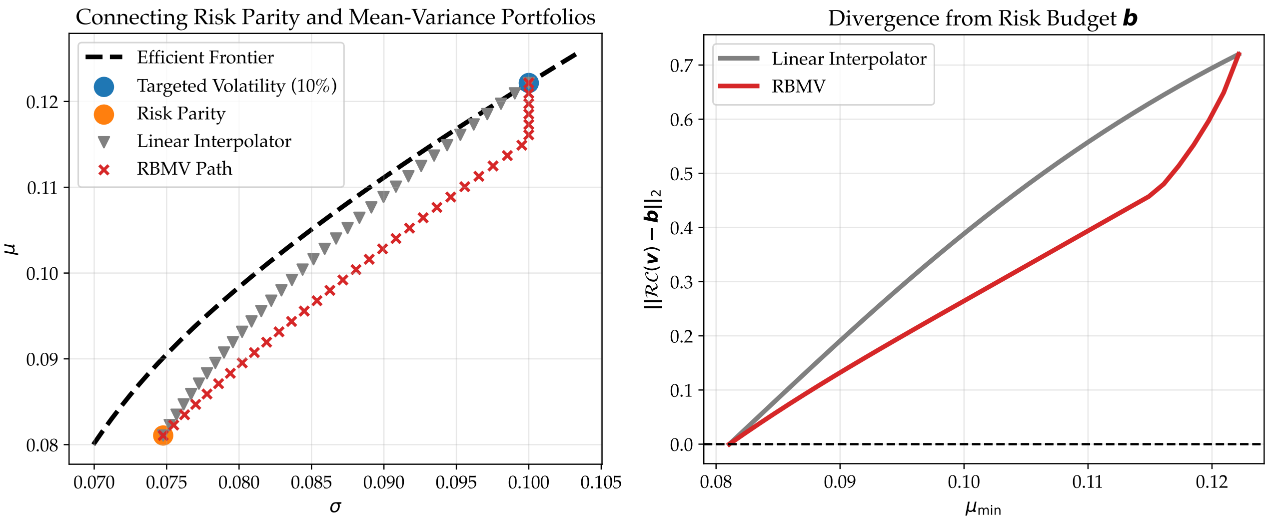

associated Lagragian, we have:

$$\frac{\partial \mathcal{L}(\boldsymbol{v};\lambda_v,\lambda_\mu,\lambda_\sigma)}{\partial v_i} = \frac{\partial \sigma(R(\boldsymbol{v}))}{\partial v_i} - \lambda_v \frac{b_i}{v_i} - \lambda_\mu(\mu_i - \mu_{\min}) - \lambda_\sigma\left(\sigma_{\max} - \frac{\partial \sigma(R(\boldsymbol{v}))}{\partial v_i}\right). \tag{16}$$

The results from Freitas Paulo da Costa et al. (2023) show that the first constraint will always bind, which implies $\lambda_v > 0$. In case the Lagrange multipliers satisfy $\lambda_\mu = \lambda_\sigma = 0$, the second and third restrictions are not binding, and we are back at the pure risk budgeting problem. The first-order conditions of the Lagrangian will imply perfect risk budgeting in this case. Hence, if $\mu_{\min}$ is low enough and $\sigma_{\max}$ is high enough, our methodology will deliver the risk budgeting portfolio from the previous section, perfectly nesting that methodology.

The more interesting case, however, happens when $\lambda_\mu > 0$ or $\lambda_\sigma > 0$. This will happen either when $\mu_{\min}$ is large enough or when $\sigma_{\max}$ is low enough. In these cases, the first-order conditions will not be equivalent to (8), implying that we are giving up perfect risk budgeting in favor of higher expected returns or lower volatility. These Lagrange multipliers' magnitude controls exactly how much the final risk contributions will deviate from the desired risk budget $\boldsymbol{b}$.

It might also happen that the solution to the problem above does not exist. That will happen if the pair $(\sigma_{\max}, \mu_{\min})$ lies outside of the mean-variance frontier traced by the mean-variance problem. By definition, the frontier traces the smallest possible volatility a portfolio can have if it ensures a certain minimum expected return. Choosing a point that lies outside this frontier for program (15) would imply that we are optimizing over an empty set.

We conduct a simulation exercise using the same calibrated population moments from the previous section to show the RBMV portfolio in action and how it compares to the two other frameworks.

Figure 4, on the panel on the left, shows the efficient frontier as a black dashed line. The blue and orange dots represent a mean-variance portfolio with a targeted volatility of 10% and the risk parity portfolio, respectively. The red diagonal crosses mark the expected return and volatility of the RBMV portfolio as we gradually increase $\mu_{\min}$, starting from $\boldsymbol{w}_{RP}^{\intercal}\mu$ and going up to $\boldsymbol{w}_{10\%}^{\intercal}\mu$. We solve the optimization problem (15) with $\sigma_{\max} = 10\%$ and $\boldsymbol{b} = (0.2, 0.2, 0.2, 0.2, 0.2)$.

The path created by the RBMV portfolio naturally connects the two colored dots. When we have $\mu_{\min} = \boldsymbol{w}_{RP}^{\intercal}\mu$, we can enjoy perfect risk budgeting since the risk parity portfolio satisfies all of the constraints. As the investor starts requiring higher expected returns, we have to give up perfect risk budgeting, i.e., the solution for the optimization will imply $\lambda_{\mu} \neq 0$, leading to a distortion of the final risk contributions.

Increasing $\mu_{\min}$ implies that the RBMV portfolio will also feature higher volatility. As soon as we hit $\sigma_{\max}$, there is a kink in the RBMV path. Since the optimization problem will not allow a final volatility higher than $\sigma_{\max}$, the RBMV path goes upwards. In other words, from that point on, we have $\lambda_{\sigma} \neq 0$ as well, further distorting the risk contributions.

The nature of these distortions is evidenced by the panel on the right in Figure 4. For each value of $\mu_{\min}$, we can solve the optimization problem and recover a $d \times 1$ vector of risk contributions $RC(\boldsymbol{v})$. We plot in red the L_2-distance between $RC(\boldsymbol{v})$, which is the realized risk contribution, and $\boldsymbol{b}$, the desired risk budget. The higher this measure is, the further away the final risk contributions are from $\boldsymbol{b}$. The red curve naturally starts at zero since the lowest $\mu_{\min}$ we consider is $\boldsymbol{w}_{RP}^{\intercal}\mu$, implying that perfect risk budgeting was available and was indeed attained.

The red curve increases as the investor becomes greedier, showing that the final risk contributions start deviating from $\boldsymbol{b}$. The slope of this red curve is a measure of the trade-off between two objectives: achieving higher expected returns and still managing risk through a desired risk budget. This curve is at the heart of our contribution because it informs the investor how much risk budgeting she needs to give up in order to achieve higher returns. Upon seeing this curve, an investor who is very concerned with risk management might say it is too steep and, as a result, she should prefer portfolios that are closer to achieving perfect risk budgeting. Alternatively, an investor who is not so concerned about risk concentration might think this curve is not too steep and that it actually represents "a good deal" in terms of the risk budgeting/higher returns trade-off. Such a conclusion would draw the investor to pick weights closer to the mean-variance frontier.

The red curve also displays a kink. It suddenly becomes steeper when λ_{σ} becomes positive because both λ_{μ} and λ_{σ} are distorting the first-order conditions from (15) to what they should have been to achieve perfect risk budgeting (see (8)). In that case, the final portfolio balances out three objectives: risk budgeting, a minimum expected return requirement, and a maximum tolerated level of volatility.

We also compare the RBMV portfolio to what the investor could get if she used a simple convex combination of risk parity and mean-variance weights. That would obviously provide another path connecting the two colored dots by construction. This path is marked by grey triangles in Figure

4. Such an alternative path is entirely above the RBMV path, implying that points along the linear interpolation provide better Sharpe ratios than the corresponding RBMV portfolio. We do not see this as a weakness of our methodology since there is no reason to think about risk budgeting if all one cares about is mean-variance optimality in the first place. In that case, one should stick with the traditional mean-variance approach – and bear the costs associated with it.

The linear interpolation procedure has two important downsides, nonetheless. First, linear interpolation will not provide *simultaneous* control over *both* minimum expected returns and volatility, as our methodology provides. Second, and perhaps more importantly, linear interpolation does not provide an explicit way to limit risk contributions. The panel on the right in Figure 4 shows the distance between the risk contributions implied by the linear interpolation method and $\boldsymbol{b}$ as we progressively get closer to the portfolio on the efficient frontier, plotted as a grey line. Such a curve is much steeper than the corresponding curve generated by the RBMV path, which means a worse trade-off between spreading risk across assets and achieving higher returns. The RBMV portfolios guard against the concentration of risk as much as they can, leading to a red curve entirely below the grey one. This means that risk contributions are closer to the desired goal of risk parity across the entire paths (with the exceptions of the endpoints) when we use our methodology.

In summary, the RBMV provides a disciplined optimization problem for the investor. Through it, she can tilt the optimal portfolio to better control risk contributions and is assured that the final portfolio achieves a minimum level of expected return while also displaying a level of volatility that is below some maximum allowed threshold. We make clear the trade-off between higher returns and risk budgeting, providing a tool for the investor to pick where she wants to be.

4 Empirical Application

We showcase how our methodology can be deployed focusing on an application of long-only portfolio formation in the U.S. equity market. Our data comes from the Center of Research in Security Prices (CRSP), ranging from 1990 to 2022 at the daily frequency.

One important choice for the application of any methodology for portfolio formation is the universe of assets available for investment. On the one hand, we would like to include as many assets as possible, aiming at a comprehensive study. On the other hand, increasing the number of assets, or increasing d in our notation, implies estimating more parameters with the same amount of data. Our exercise, trying to take this trade-off seriously, uses $d = 50$. We build portfolios using three different frameworks: Mean-Variance, Risk Parity (as a special case of risk budgeting), and the Risk-Budgeted Mean-Variance (RBMV). We detail how these portfolios are created below and then compare their evolution over time.

4.1 Portfolio Formation Procedure

Starting in January 1990, at the beginning of each month, we deploy the following procedure:

1. We select the top $d = 50$ firms with the largest market capitalization in the U.S. on the first trading day of the month;

2. We consider the last two years of daily returns to estimate a $d \times 1$ vector of expected returns $\widehat{\boldsymbol{\mu}}$ and a $d \times d$ covariance matrix $\widehat{\Sigma}$. We use the sample average of returns for $\widehat{\boldsymbol{\mu}}$ and the sample covariance matrix for $\widehat{\Sigma}$;

3. We solve for the long-only minimum variance portfolio weights, whose volatility we denote by σ_{MinVol}. Then, we set $\sigma_{max} = \min\{\sigma_{MinVol} + 0.02, 0.2\}$;

4. We solve the traditional Mean-Variance problem, finding the non-negative weights that will deliver the highest possible expected return with a maximum volatility of σ_{max}. The expected return of this portfolio is denoted by μ_{MV};

5. Given these estimated parameters, we set the risk budget $\boldsymbol{b} = \left(\frac{1}{50}, ..., \frac{1}{50}\right)$ and solve the Risk Parity problem, as in (10);

6. We finally compute two RBMV portfolios, fixing $\boldsymbol{b}$ and σ_{max} as above, but imposing two different values for μ_{min}. We define

$$\mu_{min,\,\text{conservative}} \equiv \min\{\mu_{MV} - 0.05, 0.1\} \tag{17}$$

$$\mu_{min,\,\text{greedy}} \equiv \min\{\mu_{MV} - 0.05, 0.2\} \tag{18}$$

   One of the RBMV portfolios will be more conservative, while the second will typically require higher expected returns;

7. We hold these four portfolios for twelve months, keeping track of their daily realized returns. We call the date in which these portfolios were computed their "portfolio formation date"';

8. We repeat the same procedure, including the estimation of sample moments, for the subsequent months;

As seen above, to ensure feasibility, we cannot simply define $\sigma_{max} = 0.2$ and $\mu_{min} = 0.1$, for example. The reason is that the RBMV problem might look for a portfolio that is outside the efficient frontier. That can happen if the volatility of the minimum volatility portfolio is below 20% or if, given a maximum volatility value, $\mu_{MV} < 0.1$. Then, in that case, we would be looking for a portfolio that is too good to exist, given the estimated sample moments. To accommodate this issue, we adjust the optimization parameters to deliver a feasible estimate.

In Appendix A, Figures A.1 and A.2 show the evolution of σ_{max} and μ_{MV}, respectively. The volatility threshold of 20% was only an issue for the period immediately after the start of the Global Financial Crisis in 2008. Even during that period, however, the volatility of the minimum variance portfolio was around 26%, implying that our adaptive procedure will not distort estimates in a significant way. Similarly, it was only during the aftermath of the dot-com bubble and the Global Financial Crisis that μ_{MV} dropped below 20%, dropping below 10% even less frequently.

We will have a time series of daily realized returns for each portfolio and each formation date. We can also compute the Gini index of these weights as a measure of portfolio concentration. We analyze the characteristics of these portfolios below.

4.2 Results

4.2.1 Portfolio Concentration

It is instructive to first analyze the Gini index of the portfolio weights over time. Unlike realized returns, these are not subjective to uncertainty *after* the portfolios are formed. They are only potentially affected by estimation error of population moments. Figure 5 displays the time-series evolution of the Gini indexes.

The dashed lines refer to the traditional Mean-Variance and Risk Parity frameworks. Building on intuition from Figure 4, they are extreme poles in terms of portfolio concentration. The solid lines in Figure 5 represent the Gini index of the RBMV portfolios.

Figure 5: For each of the portfolios, we show the Gini index of portfolio weights over time. The dates refer to the dates when these portfolios were formed. The universe of investable assets is the $d = 50$ largest stocks in the U.S. market at the moment of portfolio formation.

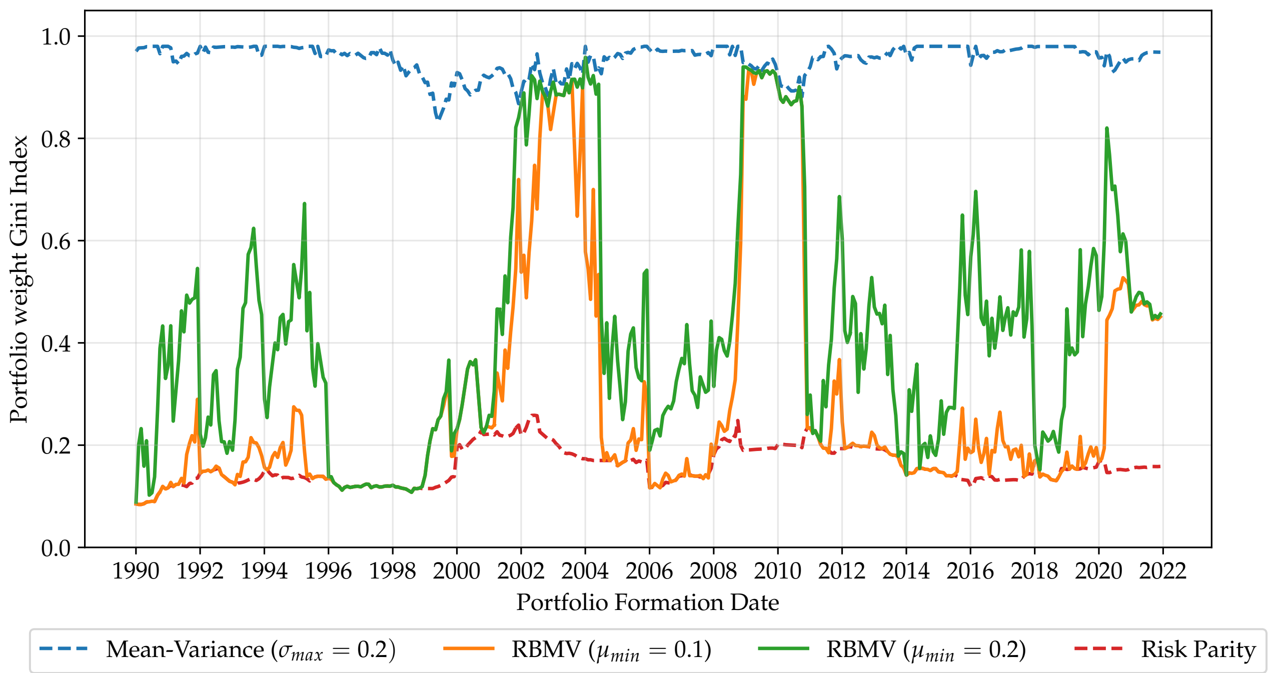

The Gini index for the Mean-Variance portfolio is consistently close to 1 and always above 0.8. Throughout our exercise, this framework never chose more than five assets to invest in, leading to high Gini index values in general. On the other extreme in terms of portfolio concentration, the Gini index values for the Risk Parity portfolio are consistently below 0.2.

On several occasions, the RBMV portfolios coincide with Risk Parity, leading to the same Gini index. In light of our discussion from Section 3, these are cases in which perfect risk budgeting is available, given estimated parameters. Using the notation from program (15), λ_v is the only non-zero Lagrange multiplier.

However, given the other constraints, perfect risk budgeting will not be feasible in other moments. Hence, the RBMV portfolios will have to concentrate more risk on some assets, leading to typically higher values for the Gini index. Given our discussion about the simulation results presented in panel (b) from Figure 4, this is the expected behavior. Indeed, especially during moments of higher volatility and lower expected returns, like after the dot-com bubble burst and the Global Financial Crisis, the RBMV methodology had to deliver portfolios close to the Mean-Variance one. In those cases, the constraints on volatility and expected returns were relatively stringent given the market conditions, which implied that the RBMV methodology had to deliver a portfolio far away from Risk Parity to satisfy the other constraints.

We also note that the Gini index of the greedier RBMV portfolio is generally higher than the values for the other RBMV portfolio. This is also in line with our previous discussion: as we require a higher minimum expected return, our methodology has to deviate more from the risk budget we imposed and will allow for more concentration as a way to meet a higher expected return requirement. We view Figure 5 as confirmation that our methodology presents a way to bridge the gap between the Mean-Variance and Risk Parity (or risk budgeting more generally) frameworks. Parameters such as μ_{min} act as a lever that allows the investor to choose how close she wants to be to either of these poles.

4.2.2 Realized Returns and Volatility

We now turn to the analysis of the returns generated by these portfolios. Figure 6 reports the evolution of the 12-month returns generated by these methodologies, while Figure 7 shows the annualized realized volatility of the daily returns. The randomness from the data will affect results in two ways, in comparison with the ideal setup from Section 3. First, estimation error of population moments distorts optimal weights. Second, given portfolio weights, we can only, at best, approximate the expected returns these portfolios have and associated volatilities with sample moments.

The returns earned by the RBMV portfolios are generally positive but become sharply negative when these portfolios are formed just before large and long-lasting market downturns. We see this behavior for portfolios formed during 2000 and 2001, which would have had low returns one year later due to the dot-com bubble burst. A similar phenomenon happens with portfolios formed on the eve of the Global Financial Crisis. Interestingly, this was not the case at the beginning of 2020, when the Covid pandemic hit the U.S. stock market. The reason for that different behavior is that the recovery from COVID-19, from the stock market point of view, was faster than during the Global Financial Crisis, for example. We do see, nonetheless, negative returns being realized for portfolios formed after 2021 since they suffered from a market downturn caused by tighter monetary policy from the Federal Reserve in the U.S. We also note that portfolios formed around 2009 earned large positive returns as all methodologies could buy almost any of the $d = 50$ assets at a discount.

Figure 6: We hold each portfolio for 12 months and report the return earned by each of them over time. The dates refer to when the portfolios were formed. The universe of investable assets is the $d = 50$ largest stocks in the U.S. market at the moment of portfolio formation.

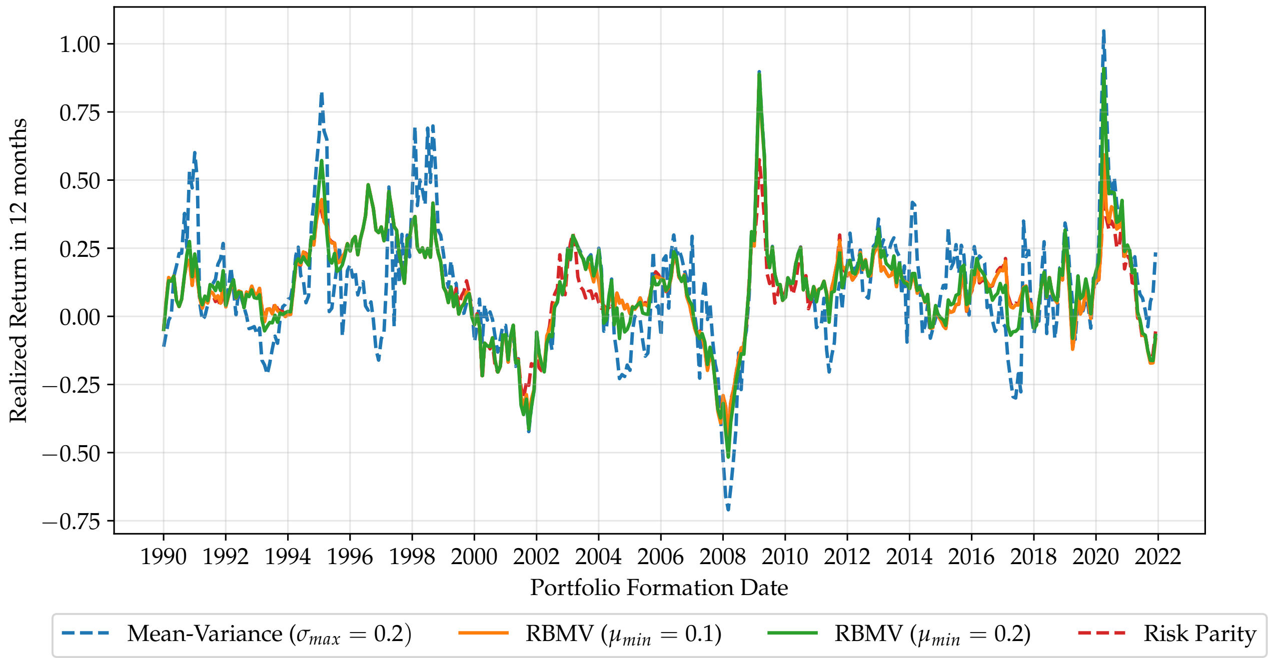

Figure 7: We hold each portfolio for 12 months and report the annualized standard deviation of daily returns. The dates refer to when the portfolios were formed. The universe of investable assets is the $d = 50$ largest stocks in the U.S. market at the moment of portfolio formation.

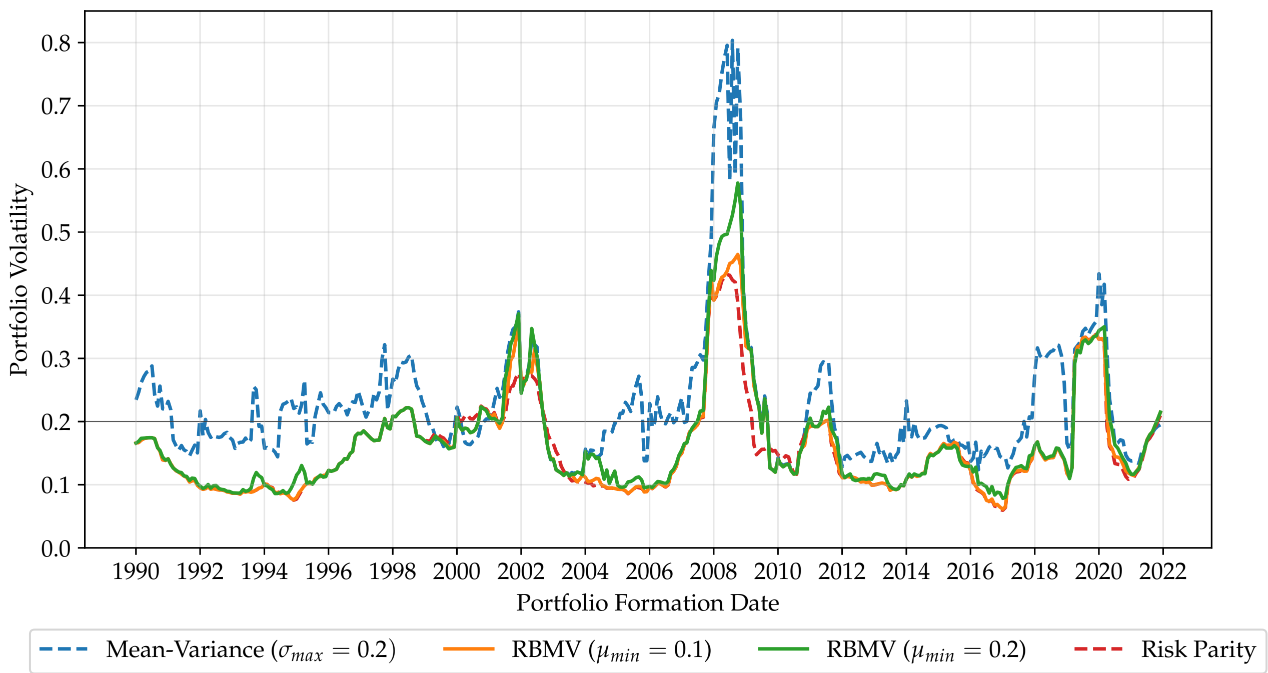

From Figure 7, we see how the portfolio concentration shown by Figure 5 is translated to higher

Figure 8: For each portfolio held for 12 months, we compute its Sharpe ratio as the average between its mean daily return and the standard deviation of these returns. The dates refer to when these portfolios were formed. The universe of investable assets is the set of $d = 50$ largest stocks in the U.S. market at the time of portfolio formation.

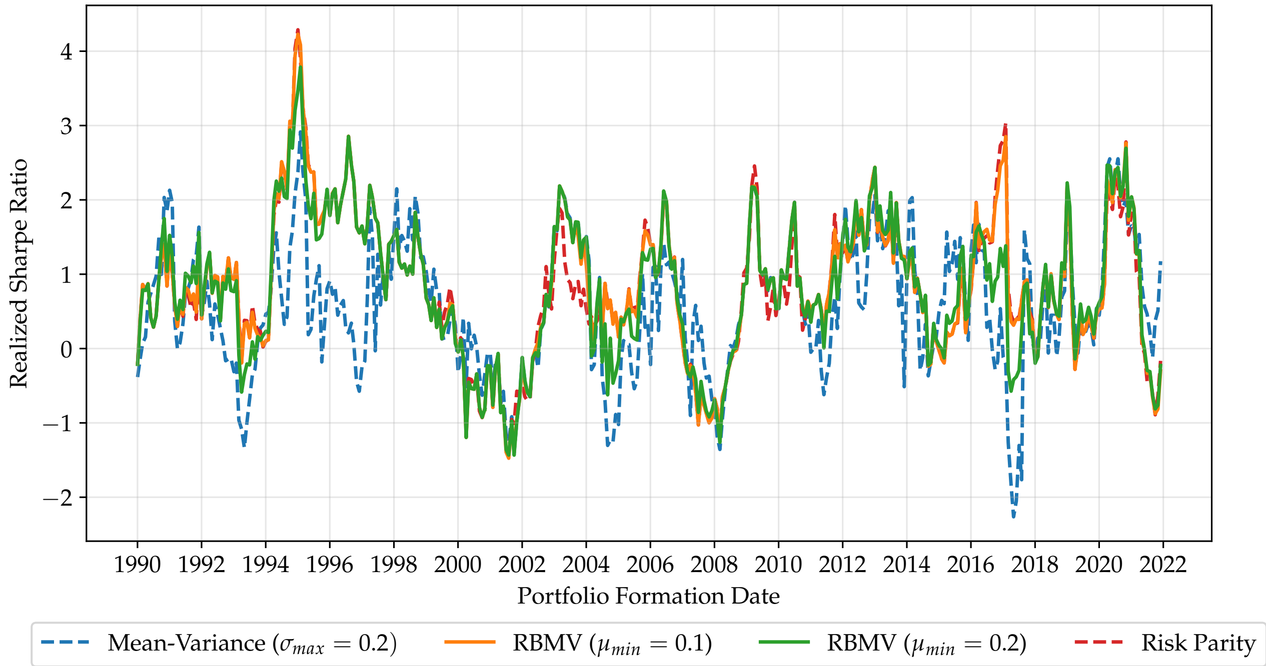

volatility for the Mean-Variance portfolio. The dashed blue line is consistently above the others, showing how the Mean-Variance portfolio is more volatile than any other of the alternatives. It is also frequently above the 20% maximum volatility target. The largest realizations of the volatility measure happened during the Global Financial Crisis and during 2020. We note how the RBMV portfolios do much better in terms of keeping volatility below the 20% threshold.

Given daily returns, we can compute their realized Sharpe ratio as their mean divided by their standard deviation. We report the realized Sharpe ratios for the four portfolios in Figure 8. These four series tend to comove, even though the blue dashed line (Mean-Variance) attains extreme negative realizations more frequently. We also note the slight and expected tendency of the greedy RBMV to follow the Mean-Variance portfolio more closely, while the conservative RBMV stays closer to Risk Parity. This can be seen in the period between 2004 and 2006, and then between 2016 and 2018.

Table 2 reports the time-series averages of the quantities presented in the figures above for different subsamples. Panel A focuses on our full sample (1990-2022). As expected from our previous discussions, the Risk Parity portfolio is the one with the lowest average annual returns. It is also the less volatile portfolio. The RBMV portfolios have higher returns, with moderately more volatility. The Mean-Variance portfolio has the highest volatility, averaging a realized volatility measure of more than 22%.

The results from Panel A also show that the RBMV portfolios had, in fact, higher Sharpe ratios than the Mean-Variance portfolio, which paid a large price for exhibiting high volatility. This is somewhat surprising because the Mean-Variance portfolio is designed to provide the best risk-return

Table 2: We present averages of different moments of portfolio returns. The first column displays the time-series average of 12-month realized returns across subsamples (see Figure 6). The second column presents the average realized volatility across subsamples (see Figure 7). The third column computes the average realized Sharpe ratio, defined as the mean daily return divided by its standard deviation. The corresponding time series is reported in Figure 8 in Appendix A. The last column reports the average Gini index of portfolio weights across subsamples (see Figure 5).

Panel A: Full Sample (1990-2022)

Portfolio	Return (%)	Volatility (%)	Sharpe Ratio	Gini Index (w_i)
Risk Parity	9.68	15.69	0.84	0.16
RBMV ($\mu_{min} = 0.1$, $\sigma_{max} = 0.2$)	10.29	15.98	0.86	0.29
RBMV ($\mu_{min} = 0.2$, $\sigma_{max} = 0.2$)	10.50	16.70	0.81	0.43
Mean-Variance ($\sigma_{max} = 0.2$)	10.20	22.28	0.58	0.96

Panel B: Before the Great Financial Crisis (1990-2006)

Portfolio	Return (%)	Volatility (%)	Sharpe Ratio	Gini Index (w_i)
Risk Parity	10.79	14.42	0.89	0.15
RBMV ($\mu_{min} = 0.1$, $\sigma_{max} = 0.2$)	10.90	14.45	0.91	0.26
RBMV ($\mu_{min} = 0.2$, $\sigma_{max} = 0.2$)	10.74	14.99	0.85	0.39
Mean-Variance ($\sigma_{max} = 0.2$)	9.92	20.84	0.50	0.95

Panel C: After 2020 (2021-2022)

Portfolio	Return (%)	Volatility (%)	Sharpe Ratio	Gini Index (w_i)
Risk Parity	0.07	15.65	0.21	0.16
RBMV ($\mu_{min} = 0.1$, $\sigma_{max} = 0.2$)	0.93	16.00	0.27	0.46
RBMV ($\mu_{min} = 0.2$, $\sigma_{max} = 0.2$)	1.52	16.01	0.30	0.47
Mean-Variance ($\sigma_{max} = 0.2$)	12.62	16.31	0.84	0.96

trade-off among all feasible portfolios whose volatility is less than 0.2. Yet, from this point of view, the RBMV portfolios did better, providing similar returns with lower volatility.

In the last column, we also report the average Gini index for portfolio weights. In line with Figure 5, the average Gini index is the lowest for the Risk Parity approach, followed by the two RBMV portfolios and then Mean-Variance, which averages a Gini index of 0.96.

Panel B reports the same averages, but focuses on portfolios that were not affected by the Global Financial crisis, with portfolio formation dates between 1990 and 2006. These portfolios are affected by data only up to the end of 2007. We see a pattern similar to Panel A, with the exception that the Mean-Variance framework delivered the lowest expected returns across the board. All portfolios displayed lower volatilities during this period in comparison to the whole sample (Panel A). In this case, the more conservative RBMV portfolio attained the highest Sharpe ratio. We see the same pattern in terms of portfolio concentration from Panel A as well.

Panel C, alternatively, shows that our methodology will not always outperform the other frameworks. We concentrate on portfolios created between 2021 and 2022 for the third panel. These portfolios were formed during the recovery from the major downturn in March 2020 and the subsequent tightening of monetary policy. In this case, only the Mean-Variance framework performed well. The RBMV portfolios and Risk Parity had returns of less than 1%, while the Mean-Variance framework delivered more than 12% of annualized returns. The realized volatility for the Mean-Variance approach was the lowest in Panel C compared to the other panels. This good performance of the Mean-Variance framework is due to its bet, during this period, on big companies from the technology sector, which did well over this subsample, delivering higher returns with moderate volatility.

In summary, our empirical exercise confirms the tendency that Mean-Variance portfolios have to be concentrated on a few assets, which ends up eroding the benefits of diversification. The Risk Parity framework did well, in general, but with lower returns. Our methodology sits at the middle of these two frameworks, providing portfolio weights that aim, at the same time, to have some explicit control regarding concentration and to have higher returns.

4.2.3 The (σ, μ)-plane

Finally, in a similar fashion to Figures 3 and 4, we display these portfolios in the (σ, μ)-plane. Figure 9 displays a scatter plot of realized moments for the four portfolios. There is one element in the plot for each methodology and formation date. Most of the realizations in the far right are blue dots, representing the Mean-Variance framework. These portfolios displayed high volatility and negative returns. We notice that the other methodologies do not visit this part of the plot because of the explicit risk control that both Risk Parity and the RBMV methodology feature.

In a consistent manner with Table 2, the blue dots do not visit the far left of the plot so frequently, which is a low volatility area. Most of the portfolios in this area are squares, diagonal crosses, and stars (conservative RBMV, greedy RBMV, and Risk Parity, respectively). In fact, many of the squares, diagonal crosses, and stars are close to each other.[5] This happens when the RBMV methodology delivers the same portfolio as Risk Parity, i.e., when perfect risk budgeting is available given the estimated parameters.

We further notice how most of the RBMV portfolios are located to the left of the vertical line at 0.2. This implies that these portfolios, designed to have volatility lower than 20%, indeed delivered, very often, realized volatilities below this threshold. On the other hand, it is visually clear that there are many blue dots to the right of the threshold, showing that the Mean-Variance framework, even though designing a portfolio to have a maximum volatility of 20% as well, cannot deliver on its promise. Finally, we also note that there are essentially no stars with more than 50% of realized returns. This is a manifestation of one of the most important drawbacks of the Risk Parity framework, which is delivering typically lower returns. This was confirmed by Table 2.

[5]We offer a “zoomed-in” version of this picture in Figure A.4, in Appendix A, where we can see the elements close to each other but not exactly on top of each other. We “zoom-in” by cropping part of the picture, inevitably leaving out some of the elements with more extreme realized volatility.

Figure 9: For each portfolio created and held for 12 months, we report the its realized volatility and realized return in the (σ, μ)-plane. The dashed vertical line is a reference at $\sigma_{\max} = 0.2$, which was imposed for the Mean-Variance and RBMV approches.

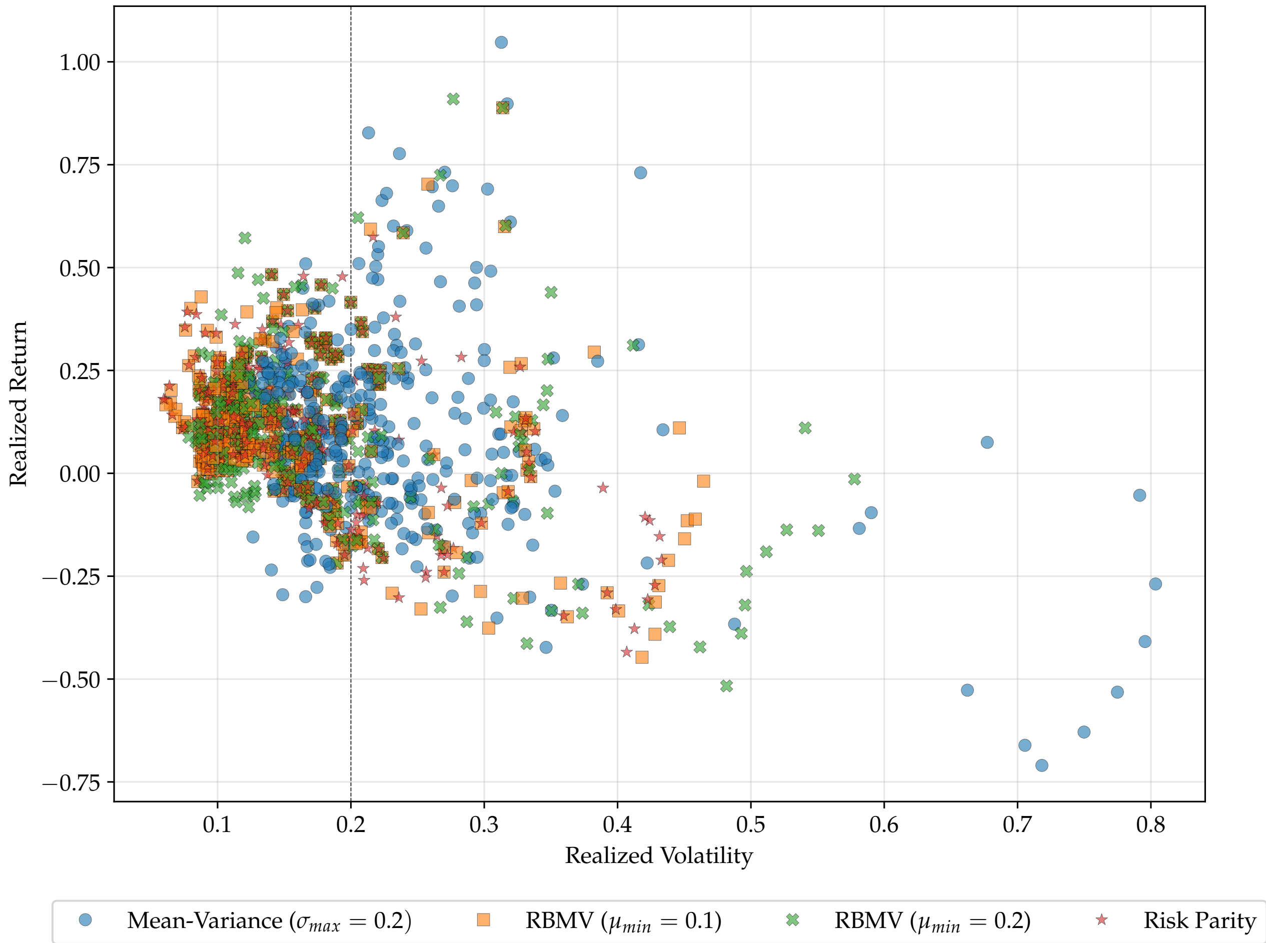

5 Conclusion

The Mean-Variance portfolio formation framework is one or the cornerstones of modern portfolio theory. More recently, the Risk Budgeting framework was introduced to circumvent some of its caveats, such as portfolio risk concentration and its high sensitivity to estimated parameters. However, pure risk budgeting enables no control over expected returns. Our methodology, the Risk-Budgeted Mean-Variance portfolio, introduces a disciplined way to balance the trade-off between risk concentration and the desire for higher returns. We also show how to efficiently compute this portfolio, given a set of population moments.

Our simulation results show our methodology naturally nests the two other frameworks and, crucially, gives the investor a way to choose how close she wants to be to either of them. Our methodology makes explicit the trade-off between risk concentration and expected returns and enables the investor to use the parameters of the optimization problem we propose to assess this trade-off given estimated parameters.

We provide an empirical illustration of our methodology using data from the U.S. stock market. We contrast the behavior of the Risk-Budgeted Mean-Variance portfolios to traditional Mean-Variance and Risk Parity. Our methodology typically delivered effective volatility with realized returns that were comparable to the Mean-Variance approach without stark portfolio concentration. Portfolios constructed using our methodology and held for one full year averaged Sharpe ratios of more than 0.8.

All cases studied in the empirical example require, nonetheless, a previous step in which the investor must estimate expected returns and a covariance matrix for these returns. Even though our application was not made in an extremely high-dimensional setting (fifty assets), estimating these moments efficiently is a crucial step for portfolio formation, and it is a complicated estimation problem due to the number of estimated parameters. One could couple more sophisticated estimators with the subsequent use of our methodology. Since we were focused on comparing our approach to traditional portfolio formation approaches, we leave this possibility for future research.

References

Anis, H. T. and Kwon, R. H. (2022). Cardinality-constrained risk parity portfolios. *European Journal of Operational Research*, 302(1):392–402.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.

Bai, X., Scheinberg, K., and Tutuncu, R. (2016). Least-squares approach to risk parity in portfolio selection. *Quantitative finance*, 16(3):357–376.

Bellini, F., Cesarone, F., Colombo, C., and Tardella, F. (2021). Risk parity with expectiles. *European journal of operational research*, 291(3):1149–1163.

Bienstock, D. (1996). Computational study of a family of mixed-integer quadratic programming problems. *Mathematical Programming*, 74(2):121–140.

Black, F. and Litterman, R. B. (1991). Asset allocation: Combining investor views with market equilibrium. *The Journal of Fixed Income*, 1(2):7–18.

Bruder, B., Kostyuchyk, N., and Roncalli, T. (2016). Risk parity portfolios with skewness risk: An application to factor investing and alternative risk premia. *Available at SSRN 2813384*.

Cetingoz, A. R., Fermanian, J.-D., and Guéant, O. (2024). Risk budgeting portfolios: Existence and computation. *Mathematical Finance*, 34(3):896–924.

Cetingoz, A. R. and Guéant, O. (2025). Asset and factor risk budgeting: a balanced approach. *Quantitative Finance*, 25(2):181–195.

Chang, T.-M., Meade, N., Beasley, J. E., and Sharaiha, Y. M. (2000). Heuristics for cardinality constrained portfolio optimisation. *Computers & Operations Research*, 27(13):1271–1302.

Chaves, D., Hsu, J., Li, F., and Shakernia, O. (2012). Efficient algorithms for computing riskparity portfolio weights. *The Journal of Investing*, 21(3):150–163.

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953.

Freitas Paulo da Costa, B., Pesenti, S. M., and Targino, R. S. (2023). Risk budgeting portfolios from simulations. *European Journal of Operational Research*, 311(3):1040–1056.

Griveau-Billion, T., Richard, J.-C., and Roncalli, T. (2013). A fast algorithm for computing high-dimensional risk parity portfolios. *Available at SSRN 2325255*.

Haugh, M., Iyengar, G., and Song, I. (2017). A generalized risk budgeting approach to portfolio construction. *Journal of Computational Finance*, 21(2):29–60.

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

Ji, R. and Lejeune, M. A. (2018). Risk-budgeting multi-portfolio optimization with portfolio and marginal risk constraints. *Annals of Operations Research*, 262(2):547–578.

Jorion, P. (1986). Bayes-stein estimation for portfolio analysis. *Journal of Financial and Quantitative Analysis*, 21(3):279–292.

Jurczenko, E. and Teiletche, J. (2019). Expected shortfall asset allocation: A multi-dimensional risk-budgeting framework. *The Journal of Alternative Investments*, 22(2):7–22.

Konno, H. and Yamazaki, H. (1991). Mean-absolute deviation portfolio optimization model and its applications to tokyo stock market. *Management Science*, 37(5):519–531.

Lassance, N., DeMiguel, V., and Vrins, F. (2022). Optimal portfolio diversification via independent component analysis. *Operations Research*, 70(1):55–72.

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

Li, J. (2015). Sparse and stable portfolio selection with parameter uncertainty. *Journal of Business & Economic Statistics*, 33(3):381–392.

Li, X., Uysal, A. S., and Mulvey, J. M. (2022). Multi-period portfolio optimization using model predictive control with mean-variance and risk parity frameworks. *European Journal of Operational Research*, 299(3):1158–1176.

Lobo, M. S., Fazel, M., and Boyd, S. (2007). Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, 152(1):341–365.

Maillard, S., Roncalli, T., and Teïletche, J. (2010). The properties of equally-weighted risk contributions portfolios. *The Journal of Portfolio Management*.

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

Mausser, H. and Romanko, O. (2018). Long-only equal risk contribution portfolios for CVaR under discrete distributions. *Quantitative Finance*, 18(11):1927–1945.

Meucci, A., Santangelo, A., and Deguest, R. (2015). Risk budgeting and diversification based on optimized uncorrelated factors. *Risks*, 11(29):70–75.

Michaud, R. O. (1989). The markowitz optimization enigma: Is "optimized" optimal? *Financial Analysts Journal*, 45(1):31–42.

Pesenti, S. M., Jaimungal, S., Saporito, Y. F., and Targino, R. S. (2024). Risk budgeting allocation for dynamic risk measures. *Operations Research*.

Roncalli, T. (2013). *Introduction to risk parity and budgeting*. CRC Press.

Roncalli, T. and Weisang, G. (2016). Risk parity portfolios with risk factors. *Quantitative Finance*, 16(3):377–388.

Spinu, F. (2013). An algorithm for computing risk parity weights. *Available at SSRN 2297383*.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2):65–86.

A Extra Figures

Figure A.1: We report the σ_{max} value used in our empirical application. It is defined as $\sigma_{max} \equiv \{\sigma_{MinVol} + 0.02, 0.2\}$, where σ_{MinVol} is the volatility of the minimum variance portfolio given the estimated population parameters in each given portfolio formation date.

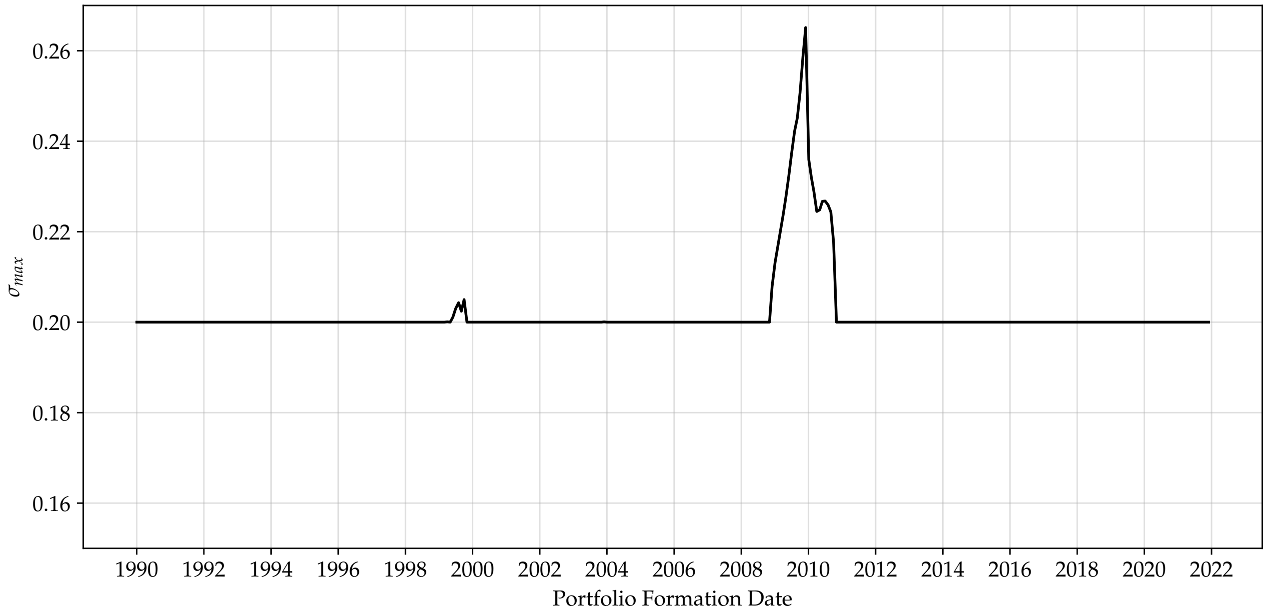

Figure A.2: We report the evolution of μ_{MV}, which is the expected return of the Mean-Variance portfolio when σ_{max} was used as a volatility bound. We recompute the population parameters for each portfolio formation date with the last two years of data.

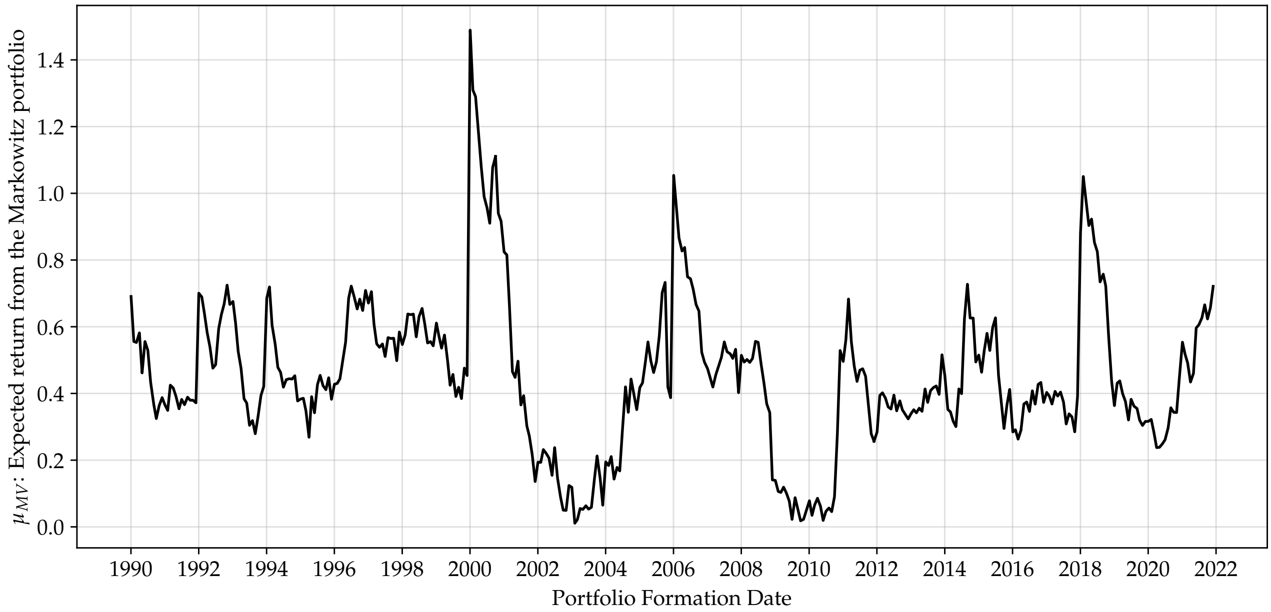

Figure A.3: For each portfolio held for 12 months, we keep track of the daily returns and report the annualized average of these returns. The dates refer to when these portfolios were formed. The universe of investable assets is the set of $d = 50$ largest stocks in the U.S. market at the time of portfolio formation.

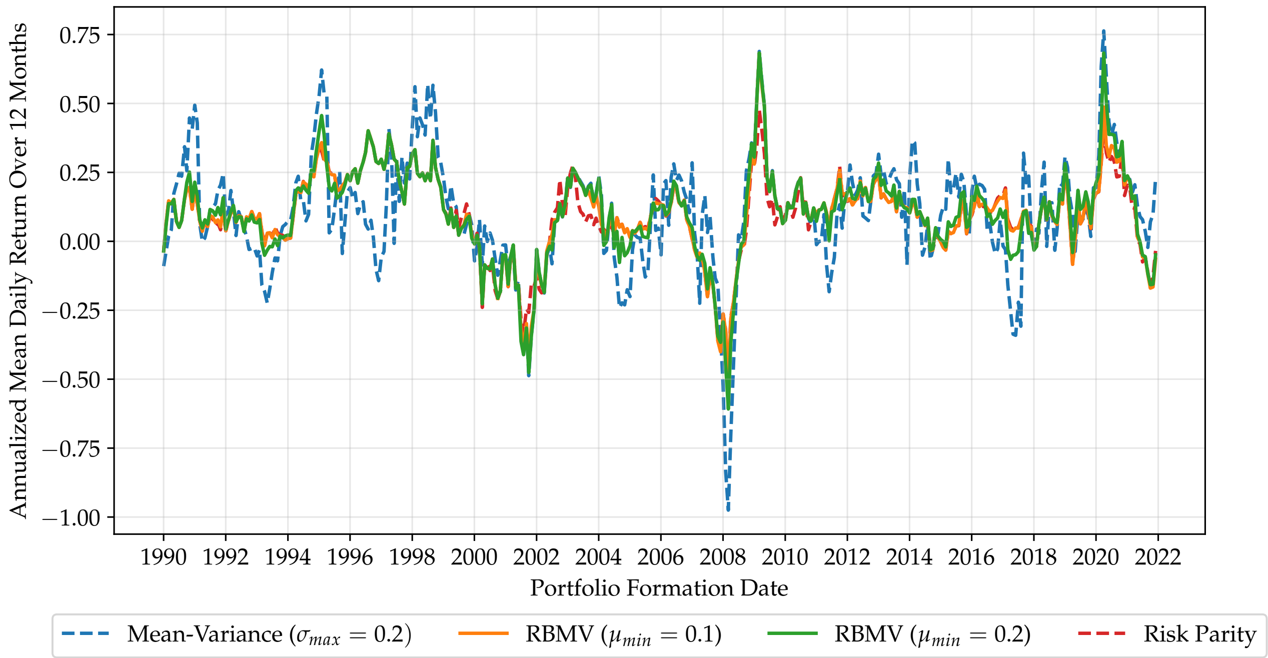

Figure A.4: This is a cropped version of Figure 9. See the details in the main text.