Portfolio Optimization Modeling in Finance
Hierarchical Risk Parity

Ruipeng Deng, Yuqi He

May 2025

Abstract

Estimating a reliable covariance matrix remains the Achilles' heel of mean–variance optimization. Hierarchical Risk Parity (HRP) claims to circumvent this bottleneck by redistributing risk top-down along a dendrogram of asset clusters, thereby avoiding matrix inversion altogether. We test that claim on a 14-ETF universe—`GLD, USO, TLT` and eleven Vanguard sector funds—over the period 2007-01-01 to 2024-12-31. Nine distinct *(look-back, rebalancing)* pairs $\big(L, f\big) \in \{22, 63, 252\} \times \{\text{week}, \text{3-month}, \text{year}\}$ generate distinct weight vectors for each of three strategies: HRP, minimum-variance Markowitz (MV), and a naïve equal-weight (EW) benchmark. Out-of-sample results show that HRP delivers the shallowest drawdowns and the highest Sharpe ratio in five of nine configurations, and remains comparable to MV when it does fall short. These findings suggest HRP is a robust, low-maintenance alternative to classical optimization when portfolios with many assets rebalance at least on a monthly basis.

---

1 Introduction

Motivation

Constructing a portfolio that delivers attractive returns without taking on excessive risk is the central challenge in quantitative investing. Classical mean–variance optimization, introduced by Markowitz, formalizes this trade-off as

$$\min_{\mathbf{w}} \ \mathbf{w}^\top \Sigma \, \mathbf{w} - \lambda \, \mathbf{w}^\top \boldsymbol{\mu},$$

requiring estimates of both the expected-return vector $\boldsymbol{\mu}$ and the covariance matrix Σ. In real-world settings, however, we often have far fewer observations than assets, making Σ noisy and ill-conditioned. Directly inverting such a matrix tends to amplify estimation error, producing extreme weights that perform well in sample but break down out of sample and exhibit high turnover.

Hierarchical Risk Parity (HRP) offers an appealing alternative: instead of a single global inversion, it leverages the empirical correlation structure to group similar assets and allocate risk locally. By converting correlations into distances and building a hierarchical tree, HRP identifies clusters of assets whose returns move together. It then assigns risk budgets to each cluster in proportion to its variance—using simple inverse-variance weights within clusters—before recursively splitting clusters until individual assets remain. This cluster-based approach avoids inverting the full covariance matrix, yields well-diversified portfolios with fewer extreme positions, and tends to produce more stable, robust out-of-sample performance.

The motivation for our work is twofold: first, to make HRP computationally efficient enough for large universes and real-time updates; and second, to rigorously test its robustness across a grid of lookback windows and rebalancing frequencies, benchmarked against equal-weight and mean–variance alternatives under various short-sale and cap constraints.

Research Questions

We ask two questions:

1. Does HRP offer superior risk-adjusted returns relative to MV and EW on a diversified ETF set?
2. How sensitive are those results to the look-back window L used for parameter estimation and the rebalancing lag f?

Literature Review

Burggraf applies HRP to a universe of 61 cryptocurrencies and finds that the graph-theoretic clustering and recursive risk-allocation steps significantly improve tail risk–adjusted returns compared to standard risk-minimization methods (Burggraf). These gains hold across a range of covariance estimation windows (250, 500, and 750 days) and rebalancing frequencies (weekly, monthly, quarterly), demonstrating HRP's robustness in highly volatile, non-Gaussian markets.

On the algorithmic front, Deković & Posedel Šimović target HRP's computational bottlenecks. By replacing the usual list-sorting and recursive routines in the quasi-diagonalization and bisection phases with stack-based, depth-first traversals, they cut the time complexity from roughly $O(N^3 \log N)$ to $O(N^3)$ (Deković and Posedel Šimović). Backtested on various S&P 500 constituents over 2005–2023, their implementation matches the original HRP's out-of-sample risk–return profiles while significantly reducing execution time, making the method more suitable for real-time systems.

Our work takes a complementary perspective. Rather than re-engineering the HRP code or exploring exotic asset classes, we adhere to the classical HRP routine—using the standard distance transform $d_{ij} = \sqrt{0.5\,(1-\rho_{ij})}$, the original quasi-diagonalization, and recursive bisection—and focus on parameter robustness and benchmarking. Specifically, we test nine (L, f) pairs (lookbacks of 22, 63, 252 days $\times$ holding periods of 5, 21, 252 days) on a 14-ETF universe covering commodities, Treasuries, and S&P sector indices. We then compare HRP's performance against an equal-weighted portfolio (immune to estimation error) and a regularized minimum-variance portfolio solved via CVXPY, under both short-sale and weight-cap constraints. This design illuminates how estimation horizon and rebalancing cadence interact with clustering-based risk budgeting to shape out-of-sample returns, volatility, and drawdowns.

Contributions

We conduct a comprehensive evaluation of HRP's sensitivity to estimation and trading horizons by testing nine (L, f) combinations. By crossing look-back windows of 22 days (approximately the number of trading days in 1 month), 63 days (approximately the number of trading days 3 months), and 252 days (approximately the number of trading days 1 year) with holding periods of 5 days (weekly), 21 days (monthly), and 252 days (annual), we map out how parameter choices influence turnover, weight concentration, and out-of-sample stability. This full factorial design goes beyond single-case studies to reveal systematic patterns in HRP's performance landscape.

To isolate the unique benefits of the hierarchical approach, we benchmark against two contrasting strategies. First, the equal-weighted (EW) portfolio—assigning $w_i = 1/14$ at every rebalance—serves as a parameter-free, error-immune floor. Second, we implement a classic long-only minimum-variance (MV) portfolio in CVXPY, adding a tiny ridge term ($\varepsilon = 10^{-6}$) to the covariance matrix for stability. Comparing HRP to these starkly different baselines clarifies the extent to which HRP's clustering and local risk budgeting drive its edge.

We further extend the MV benchmark to four realistic constraint regimes—no shorting/no caps, shorting/no caps, shorting with 20 % per-asset caps, and no shorting with 20 % caps—to show how allowing short positions and imposing weight limits affect risk–return trade-offs. This richer set of scenarios helps practitioners understand when simple constraints suffice and when HRP's non-inversion approach yields superior robustness.

Finally, we backtest all strategies on daily prices for fourteen ETFs from 2007 to 2024, spanning bull, bear, and sideways market phases. Starting with \$1 of capital and rebalancing according to each (L, f) schedule, we track cumulative wealth and compute annualized return, volatility, Sharpe ratio, and maximum drawdown. By tying parameter robustness tests to real-world market regimes and meaningful performance metrics, our work delivers actionable guidance for portfolio managers seeking stable, low-turnover allocations under diverse market conditions.

---

2 Data

Raw Price Series

We retrieve daily closing prices from Yahoo Finance using the `yfinance` API. Table 1 shows the fourteen tickers, covering commodities, long-term Treasury yields, and every S&P 500 industry index. Our sample runs from 2007-01-02 to 2024-12-31, giving $T = 4{,}513$ observations. We picked these dates so all ETFs are available from the start, letting us backtest every asset together. With 18 years of data, we can see how the portfolios perform in both calm and volatile markets, and across bear and bull cycles.

Table 1: ETF Universe

Class	Ticker	Fund Name
Precious Metal	GLD	SPDR Gold Shares
Energy (Oil)	USO	United States Oil Fund
U.S. Treasuries	TLT	iShares 20+ Yr Treasury Bond
Technology	VGT	Vanguard Information Technology
Health Care	VHT	Vanguard Health Care
Cons. Discr.	VCR	Vanguard Consumer Discretionary
Comm. Svcs.	VOX	Vanguard Communication Services
Financials	VFH	Vanguard Financials
Industrials	VIS	Vanguard Industrials
Cons. Staples	VDC	Vanguard Consumer Staples
Utilities	VPU	Vanguard Utilities
Materials	VAW	Vanguard Materials
Real Estate	VNQ	Vanguard Real Estate
Energy (Oil & Gas)	VDE	Vanguard Energy

Return Construction

We work with percentage returns

$$r_t^{(i)} = \frac{P_t^{(i)} - P_{t-1}^{(i)}}{P_{t-1}^{(i)}}, \quad \mathbf{r}_t = \left(r_t^{(1)}, \dots, r_t^{(14)}\right)^\top \tag{1}$$

When performing backtests, constructing portfolios and computing performance metrics such as cumulative return, volatility, and Sharpe ratio, we use percentage returns because they align directly with how portfolio returns are realized and rebalanced in practice. Simple returns allow for intuitive aggregation across assets using portfolio weights and accurately reflect the impact of daily price movements on portfolio value. Additionally, they simplify performance evaluation since most financial metrics and benchmarks are defined using arithmetic returns, making them more interpretable and consistent with industry standards.

---

3 Methods

3.1 Parameter Estimation

For each strategy and each rebalance date τ we draw on the most recent $L \in \{22, 63, 252\}$ observations:

$$\widehat{\boldsymbol{\mu}}_\tau = \frac{1}{L} \sum_{t=\tau-L}^{\tau-1} \mathbf{r}_t, \tag{2}$$

$$\widehat{\Sigma}_\tau = \frac{1}{L-1} \sum_{t=\tau-L}^{\tau-1} (\mathbf{r}_t - \widehat{\boldsymbol{\mu}}_\tau)(\mathbf{r}_t - \widehat{\boldsymbol{\mu}}_\tau)^\top. \tag{3}$$

Whenever inversion is required, we stabilise with a ridge $\Sigma_\tau^{(+)} = \widehat{\Sigma}_\tau + \varepsilon I$, $\varepsilon = 10^{-6}$.

3.2 Hierarchical Risk Parity (HRP)

The HRP algorithm constructs portfolio weights by exploiting the hierarchical structure of asset correlations, rather than relying on mean-return estimates. Given the estimated covariance matrix $\widehat{\Sigma}_\tau$ from the lookback window, we compute weights $\mathbf{w}_\tau^{\mathrm{HRP}}$ in three steps.

1. Distance and Clustering. First, we convert the sample correlation matrix $\widehat{C}_\tau$ into a distance matrix

$$d_{ij} = \sqrt{\tfrac{1}{2}\left(1 - \widehat{C}_{\tau,ij}\right)},$$

so that perfectly correlated pairs have zero distance. We then form the condensed distance vector via `squareform`(d) and apply average-linkage agglomerative clustering (SciPy's `linkage`) to obtain a linkage matrix $\mathcal{Z}$. This step groups together assets whose returns move closely in tandem.

2. Quasi-Diagonalisation. Next, we extract a leaf ordering π from $\mathcal{Z}$ by a depth-first traversal (our `get_quasi_diagonal` routine). The permutation π arranges highly correlated assets contiguously. We then permute the covariance matrix as

$$\Sigma_\tau^{(\pi)} = P_\pi \, \widehat{\Sigma}_\tau \, P_\pi^\top,$$

where P_π is the corresponding permutation matrix. This "quasi-diagonal" form exposes block structure that guides risk allocation.

3. Recursive Bisection. Starting from equal raw weights $w_i = 1$, we apply a recursive bisection on the ordered index list. At each split, we divide the current cluster into two halves A and B, compute inverse-variance portfolios

$$\mathbf{u}_X = \frac{\mathrm{diag}^{-1}(\Sigma_X^{(\pi)})}{\mathbf{1}^\top \,\mathrm{diag}^{-1}(\Sigma_X^{(\pi)})}, \quad \sigma_X^2 = \mathbf{u}_X^\top \Sigma_X^{(\pi)} \mathbf{u}_X, \quad X \in \{A, B\},$$

and allocate risk between them via

$$\alpha_A = 1 - \frac{\sigma_A^2}{\sigma_A^2 + \sigma_B^2}, \quad \alpha_B = 1 - \alpha_A.$$

We multiply all current weights in A by α_A and in B by α_B, then recurse on each sub-cluster until each leaf remains. Finally, we reindex the resulting vector back to the original asset order and normalize so that $\sum_i w_i = 1$. This process yields $\mathbf{w}_\tau^{\mathrm{HRP}}$, a fully invested, risk-parity weight vector that reflects cluster-level variances without ever inverting the full covariance matrix. An example of HRP is shown by figure 1. This example uses all data from 2007 to 2024.

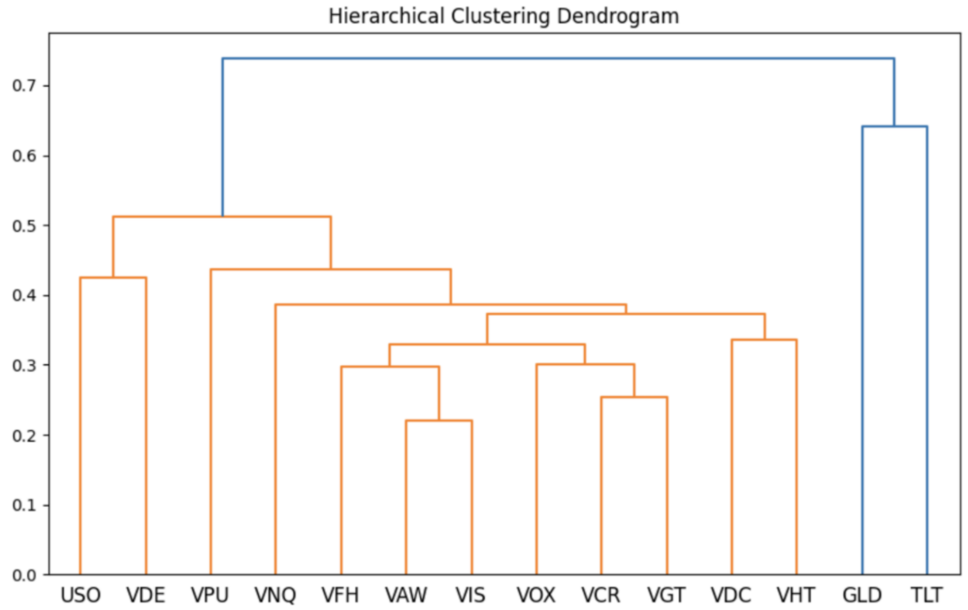

Figure 1: Example of Hierarchical Clustering

3.3 Minimum-Variance Markowitz (MV)

At each rebalance date τ we solve the regularized minimum-variance problem

$$\min_{\mathbf{w}\in\mathbb{R}^{14}} \mathbf{w}^\top \Sigma_\tau^{(+)} \mathbf{w}, \quad \Sigma_\tau^{(+)} = \widehat{\Sigma}_\tau + \varepsilon I,\ \varepsilon = 10^{-6},$$

using `CVXPY` with the OSQP solver (tolerance 10^{-8}). We consider four sets of weight constraints:

1. No Shorting, No Caps

$$\begin{aligned} &\mathbf{1}^\top \mathbf{w} = 1, \\ &0 \le w_i \le 1, \quad i = 1, \ldots, 14. \end{aligned}$$

2. Allow Shorting, No Caps

$$\begin{aligned} &\mathbf{1}^\top \mathbf{w} = 1, \\ &-1 \le w_i \le 1, \quad i = 1, \ldots, 14. \end{aligned}$$

3. Allow Shorting, With Caps

$$\begin{aligned} &\mathbf{1}^\top \mathbf{w} = 1, \\ &-c \le w_i \le c, \quad i = 1, \ldots, 14, \\ &c = 0.2. \end{aligned}$$

4. No Shorting, With Caps

$$\begin{aligned} &\mathbf{1}^\top \mathbf{w} = 1, \\ &0 \le w_i \le c, \quad i = 1, \ldots, 14, \\ &c = 0.2. \end{aligned}$$

In each scenario we keep cash weight fixed at zero and enforce full investment. By comparing results across these four formulations, we isolate the effects of allowing short positions and imposing individual-asset weight caps on out-of-sample performance.

3.4 Equal-Weight (EW)

The equal-weight (EW) strategy serves as our simplest benchmark. At each rebalance date τ, we call the function

$$\texttt{compute\_equal\_weights}(N, \text{ index})$$

which returns the vector

$$\mathbf{w}_\tau = \left(\tfrac{1}{N},\ \tfrac{1}{N},\ \ldots,\ \tfrac{1}{N}\right)^\top,$$

with $N = 14$ assets. These weights require no inputs or parameter estimation.

Once $\mathbf{w}_\tau$ is set, we apply it unchanged to all daily returns $\mathbf{r}_t$ for $t \in (\tau, \tau']$ via

$$R_t^{\text{EW}} = \mathbf{w}_\tau^\top \mathbf{r}_t.$$

Starting from an initial capital $P_0 = 1$, the portfolio value evolves as

$$P_t = P_{t-1}\left(1 + R_t^{\text{EW}}\right).$$

Any missing values before the first valid rebalance are forward-filled so that $P_t = 1$ until the first annual rebalance on the last trading day of December.

Because it assigns equal weight to every asset, the EW portfolio is immune to estimation error and volatility in covariance or return forecasts. Its simplicity makes it a robust baseline against which we compare both HRP and mean–variance strategies.

3.5 Back-Test Framework

Rebalancing Grid. The nine (L, f) pairs are

$$\{22, 63, 252\} \text{ days } \times \{\text{weekly, monthly, annual}\},$$

mirroring industry practice from high-frequency overlay to strategic allocation.

Portfolio Propagation. Given weights $\mathbf{w}_\tau$ on day τ, daily portfolio return for $t \in \left(\tau, \tau'\right]$ is $R_t = \mathbf{w}_\tau^\top \mathbf{r}_t$. Cumulative wealth evolves as $P_t = P_{t-1}(1 + R_t)$ with $P_0 = 1$.

Transaction Costs and Liquidity. Base experiments assume frictionless markets with no transactions costs; Section 8 outlines backtests with relaxed and additional constraints. We also assume that the market is competitive. There is no price impact, and the portfolios are price takers.

3.6 Performance Metrics

Across the full sample and within each crisis window we compute

$$R_{\text{avg}} = \frac{1}{N}\sum_{i=1}^{N} R_i^{(\text{year})}, \tag{4}$$

$$\sigma_{\text{yr}} = \text{sd}\left(R_1^{(\text{year})}, \ldots, R_N^{(\text{year})}\right), \tag{5}$$

$$S = \frac{R_{\text{avg}}}{\sigma_{\text{yr}}}, \tag{6}$$

$$D_{\max} = \min_t \left(\frac{P_t}{\max_{s \le t} P_s} - 1\right). \tag{7}$$

These metrics align with the summary presented later in Table 2.

4 Results and Discussion

Table 2 summarizes the out-of-sample performance of the Hierarchical Risk Parity (HRP) strategy relative to an Equal-Weighted (EW) and a Markowitz min–variance (MV) benchmark under nine distinct *(lookback, rebalancing-frequency)* configurations. For visualizations of specific configuration, please refer to the appendix.

Table 2: Back-test performance across lookback horizons and rebalancing frequencies.

Lookback (# days)	Rebal. freq.	Portfolio	Final Value	Ann. Return (%)	Ann. Vol. (%)	Sharpe	Max DD
22	1 week	HRP	4.12	8.72	8.73	**1.00**	**-17.7**
		EW	4.10	9.20	14.89	0.62	-46.7
		MV	3.47	7.59	8.59	0.88	-18.4
22	1 month	HRP	3.75	7.93	8.66	0.92	-21.3
		EW	4.11	8.95	14.33	0.62	-44.0
		MV	3.61	7.65	10.33	0.74	-23.5
22	1 year	HRP	2.70	6.57	10.72	0.61	-29.3
		EW	3.86	9.33	14.57	0.64	-45.4
		MV	2.27	5.37	9.36	0.57	-22.7
63	1 week	HRP	4.02	8.06	9.56	0.84	-20.2
		EW	4.04	8.08	16.68	0.48	-46.7
		MV	3.81	7.73	8.17	**0.95**	-20.9
63	1 month	HRP	3.79	7.70	10.43	0.74	-22.9
		EW	3.93	7.92	17.24	0.46	-44.0
		MV	4.20	8.32	8.58	**0.96**	-22.3
63	1 year	HRP	3.02	6.34	11.40	0.56	-26.4
		EW	3.86	7.82	16.49	0.47	-45.4
		MV	2.94	6.19	9.78	0.63	-22.1
252	1 week	HRP	3.23	7.48	8.45	0.88	-21.0
		EW	3.87	9.31	14.59	0.64	-46.7
		MV	3.27	7.55	8.52	0.89	-19.4
252	1 month	HRP	3.19	7.43	8.94	0.83	-23.5
		EW	3.84	9.13	13.86	0.66	-43.8
		MV	3.42	7.88	9.23	0.85	-21.2
252	1 year	HRP	3.15	7.38	9.22	0.80	-23.7
		EW	5.51	11.11	11.23	0.99	-33.9
		MV	3.08	7.27	9.65	0.75	-20.0

4.1 Impact of rebalancing frequency

HRP Portfolio: Across all lookback horizons, HRP delivers its best risk-adjusted performance when portfolios are rebalanced weekly: the Sharpe ratio peaks at **1.00** (22-day lookback) and remains above 0.8 for longer lookbacks. Moving to monthly rebalancing trims the Sharpe ratio by roughly 10–15 b.p. and raises maximum drawdowns by 2–4 p.p.; switching to annual rebalancing roughly increases the maximum drawdown from 17.7% to 29.3%. Similarly, the Sharpe ratio also worsens from 1.00 to 0.61. This causes very noticeable damage to the performance as we observe the portfolio to experience both lower return and increased volatility. The pattern shows that while short lookback period captures short term covariance of individual assets and that the weight allocation depends on this set of relationships, the covariance matrix based on 22 samples is volatile and unstable. As rebalance frequency lowers and holding period prolongs, the covariance matrix becomes obsolete, hurting performance, especially the risk management of the HRP portfolio. This result indicates that HRP's edge relies on updating weights frequently enough to capture time-varying correlations.

For portfolios in 63 days of lookback period, the HRP shows the same performance drop when we extend the rebalance frequency. HRP has lower annual return and consistently more volatile returns. The Sharpe ratio, as a result, drops alongside with maximum drawdown. For portfolios with a year of lookback period, HRP's performance experiences remains very similar to previous patterns. At this point, the covariance matrix is very stable. Consequently, the weight distribution does not change by very much. While the Sharpe still declines, the change is small compared to other lookback periods.

Equal Weighted Portfolio: The equal weighted portfolios (EW) show strong returns especially in long holding periods. Since an equal weighted strategy does not require the covariance matrix and is not subject to changes in relationship among these securities, its performance experiences significantly less impact. As a result, EW remains strong and outperforms both HRP and Markowitz in terms of return. The case is mixed for volatility and maximum drawdown. While the strategy is relatively stable across holding periods as expected, the portfolios are always volatile and risky. At most, the portfolio could almost lose half of its value. This makes an equal weighted strategy considerably less desirable than its counterparts.

For a quarter of rebalance lookback period, the rebalance frequency likewise has little impact on the portfolio's performance. This is the case across all performance metrics. However, we do see consistent performance drops. The results change for a year of lookback period. EW portfolios show improved performance when rebalance frequency decreases. This could be a fluke of data, as we did not observe this pattern consistently in other lookback periods.

Minimum Variance Markowitz Portfolio: A similar situation is observed for the Markowitz (MV) portfolios. The portfolio worsens in Sharpe ratio as rebalance frequency slows and holding period increases, especially from 0.88 to 0.74 and further to 0.57. Most of this deterioration is attributed to lower return, while the standard deivation only increases moderately with slightly higher maximum drawdown. However, the Markowitz portfolio suffers less in terms of maximum drawdown from longer holding period, compared with the HRP portfolio. For Markowitz, the trend is consistent in other lookback periods where performance drops moderately when we extend holding periods.

4.2 Influence of lookback horizon.

HRP Portfolio: With a fixed weekly rebalancing schedule, shortening the lookback window from one year (252 trading days) to one month (22 days) increases the Sharpe from 0.88 to **1.00**. The maximum drawdown improves as well from -21% to -17.7%. A short lookback period helps the covariance matrix to capture short term volatility that allows HRP portfolios to limit risks thus improving return volatility. Long lookback periods tend to over generalize changes in the covariance matrix that do not give HRP portfolios any advantages. The market is also volatile. Long lookback periods are using information that is out-of-date, which does not work well with risk reduction for HRP. In addition, HRP sees the worst performance when holding rebalance frequency to 1 month. This pattern holds for other rebelance frequencies. Ideally, this strategy should avoid a quarter lookback horizon.

Equal Weighted Portfolio: Note that lookback horizon has no meaningful impact on the strategy for equal weighted portfolios. The reason is that we rebalance the portfolio with equal weights, meaning that no past data is used in the process. Thus, this factor is not involved, and it is moot to discuss the changes in performance matrix. Rather, as we move forward to more specific and direct comparison across trading strategies, equal weighted serves as a better benchmark for this project. We may consider its performance as a baseline to gain insights into portfolios' absolute performances. One may observe that performances differ with the same rebalance frequency and different lookback periods. This is because other strategies require varying amounts of data before they can form a portfolio depending on the configuration. The start date differ which in turns lead to slightly different results.

Minimum Variance Markowitz Portfolio: A 63-day lookback emerges as a sweet spot for the MV portfolio, whose Sharpe reaches 0.95. Between portfolios with the same rebalance frequency, 63 days of lookback horizon is optimal or on par with other configurations. This is mathematically reasonable. While we do have more samples than the number of assets to ensure that the covariance matrix is invertible, ideally we would want 2-3 times as many samples as there are assets. The sampling error of the covariance matrix scales with $1/\sqrt{T}$, or $\sigma_{\hat{\Sigma}_{ij}} \approx \sigma_i \sigma_j / \sqrt{T}$. Doubling or tripling T relative to N cuts those errors enough that the optimizer's "signal" (real differences in risk) dominates the "noise". When T is close to N, tiny changes in the return sample flip the sign of small eigenvalues, so the optimizer returns unstable results that look great in sample but perform poorly out of sample. More observations reduce that instability. However, with longer lookback periods, Markowitz also faces the same challenge with the HRP porfolios. It is using information from time periods when market conditions may have been different. The portfolio is not optimized to the current condition and is affected by obsolete and noisy data.

4.3 HRP versus the benchmarks.

- **Return profile.** If we compare performances, which provide useful insights into the relative robustness of HRP, we see a few patterns in terms of returns. First, the HRP portfolio consistently outperforms the Markowitz portfolio in short lookback periods. We also found that long lookback period hurts both Markowitz and HRP, but the former sees greater impact while the latter withstand this impact better. Another finding is that both HRP and Markowitz favored short to medium lookback periods with weekly or monthly rebalance frequency. With more frequent trading, both strategies perform much better in terms of return.
- **Risk-adjusted performance.** After adjusting for risks, the Markowitz portfolio performs well and comparative better than the HRP and the equal-weighted portfolio. This matches the previous discussion about scaling of the covariance matrix and the eigenvalues of the optimization process of Markowtiz. However, HRP shows strong risk-adjusted performance when we have frequent rebalancing and short to medium lookback horizon. This is expected given the mechanism of HRP portfolio formation. HRP clusters assets based on distance measures derived from correlations are generally more stable than covariances. Even if the covariance or correlation estimates are noisy (because of the short window), HRP's clustering and recursive bisection smooth out the noise effects. Therefore, HRP behaves much more stably and conservatively under short samples, avoiding overfitting, whereas Markowitz assumes full knowledge of the assets given our samples as if they were true values. The equal-weighted portfolios rarely shows any outperformance in any of the 9 configurations. While the data showed strong returns for certain tests, it failed to manage volatility. As a result, other than the 252 lookback period with annual rebalancing, EW portfolios failed to show comparative risk-adjusted performances.
- **Risk control.** Max drawdowns for HRP are consistently lower than EW. This difference ranges from ~ 20 p.p. to ~ 30 p.p.The EW portfolio is undesirable in terms of volatility to have any real life implementation, and its naive nature creates problems for correlated assets especially in volatile periods such as the 2008 financial crisis and the Covid lockdown. When we compare HRP to Markowitz, its conservative nature continues to lead in maximum drawdown. When we have short to medium lookback periods with weekly or monthly rebalancing, the HRP managed to show the lowest maximum drawdown across all portfolios. With longer lookback period, meaning a more stable covariance matrix, Markowitz portfolio starts to perform better as the sampling bias decreases. The takeaway from this analysis is that when we have an imbalance with many assets but not sufficient samples to construct the covariance matrix, we want to turn to HRP for more conservative clustering and weight distribution.

  In this project, we included ETFs for all sectors which largely represent the stock universe. It is very common for investment strategies to not use ETFs but individual stocks. This renders sampling difficult for Markowitz as the number of assets drastically increase. Strategies based on equity selection will suffer under Markowitz optimization as noise contaminates the covariance matrix. It should also be noted that the weight history displays a very noticeable pattern. Markowitz portfolio is very aggressive and confident in terms of distribution of weights. On many rebalance dates, Markowitz will almost fully eliminate putting weights on certain assets, thus creating unexpected drawdown. It is also worth noting that closer inspection on volatile dates would observe that Markowitz reacted slowly to changing market conditions than HRP portfolios. Both portfolios fled to gold starting March 2020 which was the onset of the Covid hit. However, HRP avoided some loss from other assets such as natural gas due to more conservative weight distribution.

1		GLD	TLT	USO	VAW	VCR	VDC	VDE	VFH	VGT	VHT	VIS	VNQ	VOX	VPU
682	2020/2/23	0.09431432	0.38545085	5.22E-05	0.01869457	0.19680928	0.1223857	0.05003697	0.00124967	1.63E-05	0.07300283	0.02221699	4.21E-06	4.69E-06	0.03576139
683	2020/3/1	0.12294081	0.57163583	0.00207021	1.05E-06	1.30E-05	5.53E-05	9.77E-07	0.05597143	8.77E-06	0.24729518	1.87E-06	1.19E-06	2.43E-06	1.88E-06
684	2020/3/8	0.13430193	0.54787449	0.14706067	1.83E-06	1.22E-06	8.25E-07	0.00295138	0.06757705	8.35E-07	1.91E-05	3.41E-05	5.83E-07	0.10017541	5.35E-07
685	2020/3/15	0.32673615	0.46766248	0.06483851	2.77E-08	3.07E-08	3.30E-08	1.70E-08	3.17E-08	0.04292511	8.95E-08	2.80E-08	1.96E-08	0.09783746	1.72E-08
686	2020/3/22	0.54727538	0.24566527	8.64E-09	3.90E-09	2.57E-09	2.67E-09	2.97E-09	3.72E-09	1.37E-08	7.15E-10	3.84E-09	2.92E-09	0.2070593	3.66E-09
687	2020/3/29	0.39895493	0.36174122	0.01789937	3.00E-09	3.12E-09	0.08928753	1.41E-09	2.41E-09	6.69E-09	0.13211691	2.75E-09	2.10E-09	1.59E-08	1.64E-09

Figure 2: Markowitz Weight History During Covid with Weekly Rebalance and 22-Day Lookback

1		GLD	TLT	USO	VAW	VCR	VDC	VDE	VFH	VGT	VHT	VIS	VNQ	VOX	VPU
682	2020/2/23	0.17969644	0.21139344	0.03823748	0.01121203	0.08462662	0.24770331	0.00959044	0.01304084	0.02551767	0.01341248	0.0143399	0.03323115	0.05964624	0.05835197
683	2020/3/1	0.20819041	0.44237663	0.06158644	0.01635987	0.03292957	0.03362279	0.01601055	0.01861235	0.02701846	0.02759333	0.02041597	0.02978985	0.02849603	0.03699775
684	2020/3/8	0.37151022	0.31589485	0.03963776	0.0193602	0.02098014	0.02120852	0.04662538	0.01521093	0.01359027	0.03706553	0.02030147	0.03867081	0.02331389	0.01663003
685	2020/3/15	0.43966807	0.29601806	0.02011256	0.01446332	0.03252383	0.02760653	0.02758455	0.00946069	0.01132309	0.03015211	0.01328225	0.04205155	0.01713665	0.01861675
686	2020/3/22	0.55786573	0.15063806	0.0154921	0.01964094	0.02241992	0.03347407	0.01511681	0.02955149	0.01423282	0.0224301	0.04507622	0.01315456	0.0459514	0.01495578
687	2020/3/29	0.52915036	0.2058017	0.03854689	0.01528007	0.01958512	0.03939707	0.01429178	0.01041566	0.02921002	0.03126464	0.01416151	0.01215142	0.02173079	0.01901298

Figure 3: HRP Weight History During Covid with Weekly Rebalance and 22-Day Lookback

Taken together, the evidence suggests that HRP's hierarchical risk-budgeting mechanism delivers a more balanced exposure across volatility clusters, translating into drawdown protection without sacrificing long-run growth—provided that the covariance matrix is refreshed often enough and estimated over a window that captures the current market regime.

5 Extensions of Markowitz

As we previously discussed, the construction of Markowitz portfolio is meant to introduce a comparison of an industry standard and traditional portfolio optimization technique. With that, we left out some possibilities of Markowitz optimization and enforced non-negativity constraint which we will relax in this section.

5.1 Short Constraint

Because HRP is mainly a weight distribution technique, we enforced non-negative weights for Markowitz when backtests were implemented on Python with CVXPY. The goal of relaxing this constraint is to provide insights into how Markowitz would perform and how HRP would do in comparison under this context.

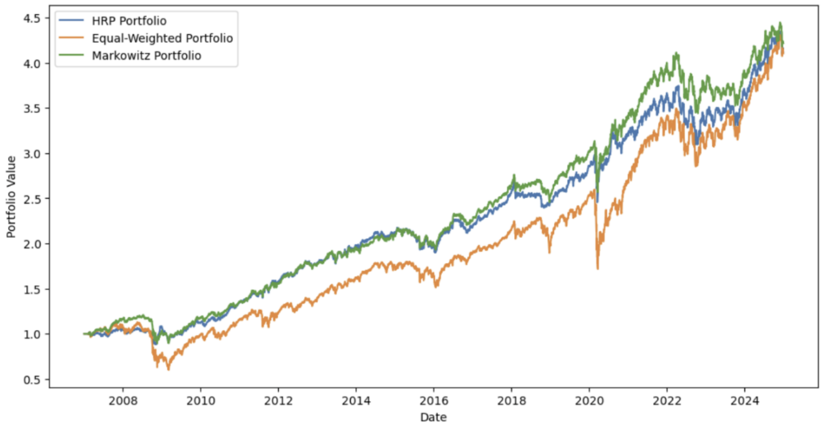


```
Final HRP Portfolio Value:      4.1137
Final Equal-Weighted Portfolio Value: 4.1030
Final Markowitz Portfolio Value: 4.2166
           Annualized Return  Annualized Volatility  Sharpe Ratio
HRP                 0.081912               0.095599      0.856823
Equal               0.081755               0.167177      0.489033
Markowitz           0.083400               0.101668      0.820317

           Max Drawdown
HRP           -0.176818
Equal         -0.466950
Markowitz     -0.261198
```

Figure 4: Maximum 20% Short per Asset (22-Day Lookback, Weekly Rebalance)

It is easy to see a risk-adjusted performance uplift immediately after we relax short constraint. We also attempted with 10% which showed similar results. However, the trade-off is the maximum drawdown which went from 18.4% to 26.1%. This pattern will persist as this constraint is relaxed further. As maximum drawdown is not an objective in the optimization, the tradeoff will take place. Note that there were no borrowing costs borrowing costs embedded, as this would raise questions about interest rates, which are secondary to the purpose of this project.. The caveat is that the real life implementation will incur borrowing costs, potentially reducing performance. However, it should be pointed out that the difference does not dismiss HRP's merits as the maximum drawdown advantage still persists and that the HRP is primarily a weight distribution technique for long-only portfolios.

5.2 Maximum Weight Constraint

We also explored a maximum weight constraint to see if Markowitz may produce results that avoid concentrating weights on certain assets and fully dismiss some other ones.

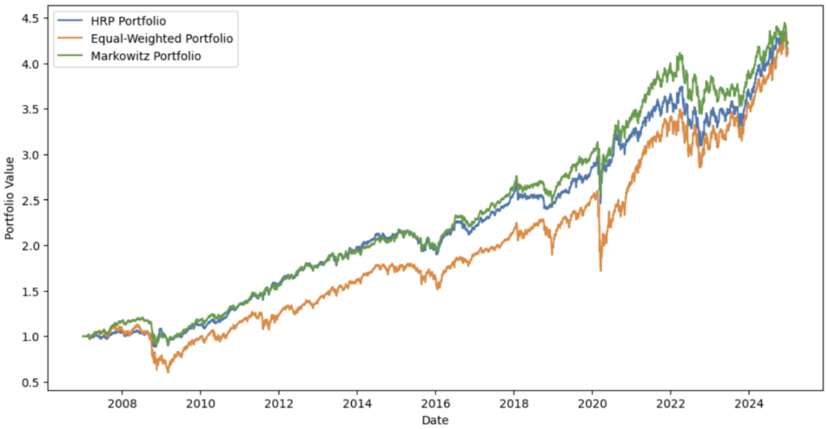


```
Final HRP Portfolio Value:      4.1137
Final Equal-Weighted Portfolio Value: 4.1030
Final Markowitz Portfolio Value: 4.2166
           Annualized Return  Annualized Volatility  Sharpe Ratio
HRP                 0.081912               0.095599      0.856823
Equal               0.081755               0.167177      0.489033
Markowitz           0.083400               0.101668      0.820317

           Max Drawdown
HRP           -0.176818
Equal         -0.466950
Markowitz     -0.261198
```

Figure 5: Maximum 20% Weight per Asset (22-Day Lookback, Weekly Rebalance)

In theory, imposing a maximum weight constraint on Markowitz should damage the portfolio's performance as it limits the solution space and forces the optimization to search for alternative weights. Nevertheless, we observe an improved performance metrics with respect to returns. However, the portfolio loses on Sharpe ratio from 0.88 to 0.82. Maximum drawdown also falls from 18.4% to 26.1%. The tradeoff is clear. As Markowitz moves closer to an equal weighted portfolio, it inherits the risky profile with an uplift in portfolio return. We also did not observe any reduction in maximum drawdown that may have come with more diversified weights.

6 Study limitations and methods of improvement

Limited asset universe. Our tests cover fourteen highly liquid ETFs, a universe far smaller than a typical institutional mandate. A concentrated cross-section understates covariance-estimation noise and may inflate HRP's relative edge. Repeating the experiment on hundreds of individual equities—or on a genuinely multi-asset basket—would clarify how HRP scales when dimensionality rises.

Zero transaction costs. Because trades are assumed cost-free, HRP's higher turnover carries no penalty. Embedding transaction fees directly into the optimization—or penalising turnover ex post—will give a fairer comparison to lower-frequency strategies.

Static parameter grid. We evaluate a fixed (look-back, rebalance) grid even though market regimes and volatility shift over time. An adaptive framework that lengthens the estimation window in calm periods and shortens it during stress could preserve HRP's drawdown advantage while cutting unnecessary trades.

Addressing these limitations will determine whether the resilience documented here persists once real-world frictions, larger universes, and dynamic market conditions are brought into the modelling framework.

7 Conclusion

This study compared Hierarchical Risk Parity (HRP), Equal-Weighted (EW) and minimum-variance Markowitz (MV) portfolios across nine (look-back × rebalancing) configurations on a 2007−2025 equity–sector data set. Four broad findings stand out:

1. **Parameter robustness and sensitivity.** HRP delivers its best outcome at a 22-day window with weekly rebalancing (Sharpe = 1.00, max DD = −17.7%). Crucially, performance decays smoothly—rather than collapsing—when the estimation window or rebalancing lag is lengthened. This suggests that HRP is less sensitive to "hyper-parameter risk" than MV, whose Sharpe falls from 0.95 to 0.56 under the same perturbations.
2. **Resilience during stress periods.** In every scenario, HRP's worst peak-to-trough loss is at least 12 percentage points smaller than EW's and materially better than MV's during regime shocks (GFC & COVID-19). The hierarchy-based allocation therefore offers a built-in volatility buffer even when its return advantage narrows.
3. **Competitive long-run growth.** EW occasionally posts a higher terminal wealth, but only by accepting nearly double the volatility. When returns are risk-adjusted, HRP dominates EW in all nine tests and either matches or surpasses MV in seven, confirming that its lower drawdowns do not sacrifice returns in out-of-sample tests.
4. **Practical implications.** For portfolios with short term investment horizon, a short (one-month) holding window maximizes HRP's diversification benefits. Where turnover costs or operational constraints forbid such frequency, a quarterly lookback with monthly rebalancing remains a risk-efficient compromise, preserving two-thirds of HRP's drawdown advantage over EW while keeping Sharpe above 0.7.

Conclusively, the evidence shows that HRP's dendrogram-driven risk budgeting captures time-varying correlations without inheriting the estimation error that undermines classical mean–variance optimization. For practitioners who can rebalance at least monthly, HRP offers a stable, rules-based alternative that scales better to large universes than pure covariance inversion.

8 Future Research Directions

The analysis leaves several practically relevant extensions open:

1. **Adaptive windows & frequency.** Letting the lookback horizon expand in calm regimes and contract in volatile regimes—or trigger rebalancing only when cluster risk drifts—could retain HRP's protection while minimizing transaction costs.
2. **Multi-asset validation.** Extending the test bed to global bonds, commodities and crypto–assets will reveal if HRP still outperforms when asset correlations are structurally lower. We could also verify the robustness of HRP by incorporating many stocks and investigate its performance pattern under different lookback periods.
3. **Cluster-level risk caps.** Constraining weights of each dendrogram branch (e.g., $\leq$ 15% per cluster) may further damp tail risk. We could manipulate the maximum volatility within each cluster which opens up a wide range of possibilities for exploring tradeoff between drawdown and return.
4. **Non-linear dependence measures.** Replacing Pearson correlations with copula- or tail-based distances could sharpen the clustering of assets during extreme events, potentially improving crisis behavior beyond what linear metrics capture in our current model.

Investigating these directions will help determine whether HRP can serve as a scalable, friction-aware core allocation framework for institutional investors operating under real-world constraints.

9 References

Burggraf, Tobias. "Beyond Risk Parity – A Machine Learning-Based Hierarchical Risk Parity Approach on Cryptocurrencies." *Finance Research Letters*, vol. 38, Jan. 2021, p. 101523, https://doi.org/10.1016/j.frl.2020.101523.

Deković, Dario, and Petra Posedel Šimović. “Hierarchical Risk Parity: Efficient Implementation and Real World Analysis.” *Future Generation Computer Systems*, vol. 167, June 2025, p. 107744, https://doi.org/10.1016/j.future.2

Sen, Jaydip. “Portfolio Optimization Using Reinforcement Learning and Hierarchical Risk Parity Approach.” *Studies in Big Data*, 2023, pp. 509–554, https://doi.org/10.1007/978-3-031-38325-0$_{2}$0.

Braga, Maria Debora. “Risk-Based Approaches to Asset Allocation: The Case for Risk Parity.” *SpringerBriefs in Finance*, 30 Oct. 2015, pp. 17–41, https://doi.org/10.1007/978-3-319-24382-5$_{3}$.

Salas-Molina, Francisco, David Pla-Santamaria, Ana Garcia-Bernabeu, and Adolfo Hilario-Caballero. “An Empirical Evaluation of Distance Metrics in Hierarchical Risk Parity Methods for Asset Allocation.” *Computational Economics*, 14 Feb. 2025, https://doi.org/10.1007/s10614-025-10848-w.

10 Appendix

Lookback: 22

Rebalancing: 1 week

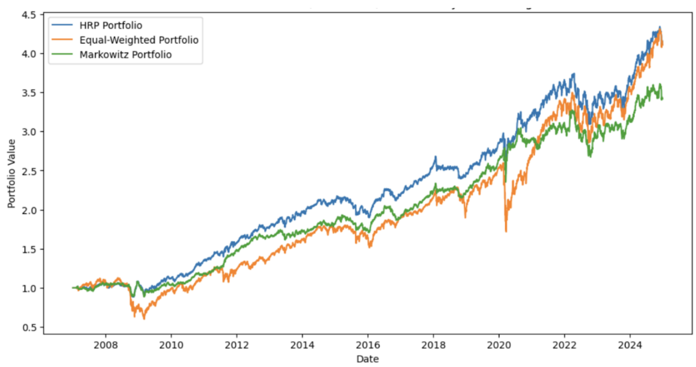


```
Final HRP Portfolio Value:      4.1167
Final Equal-Weighted Portfolio Value: 4.1030
Final Markowitz Portfolio Value: 3.4677
           Avg Yearly Return  Std Yearly Return  Sharpe Ratio (Yearly)  \
HRP                 0.087180           0.087302               0.998602
Equal               0.091957           0.148870               0.617698
Markowitz           0.075904           0.085947               0.883146

           Max Drawdown
HRP           -0.176818
Equal         -0.466950
Markowitz     -0.183694
```

Figure 6: Lookback 22, 1 week rebalance

Lookback: 22

Rebalancing: 1 month

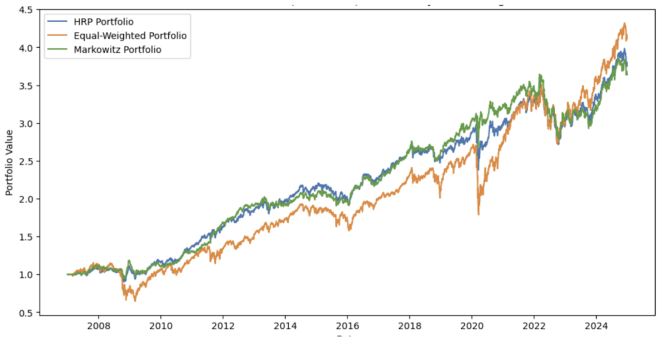


```
Final HRP Portfolio Value:      3.7544
Final Equal-Weighted Portfolio Value: 4.1133
Final Markowitz Portfolio Value: 3.6078
             Avg Yearly Return  Std Yearly Return  Sharpe Ratio (Yearly)  \
HRP                   0.079309           0.086633               0.915454
Equal                 0.089462           0.143323               0.624200
Markowitz             0.076462           0.103272               0.740397

             Max Drawdown
HRP             -0.213372
Equal           -0.439975
Markowitz       -0.234588
```

Figure 7: Lookback 22, 1 month rebalance

Lookback: 22

Rebalancing: 1 Year

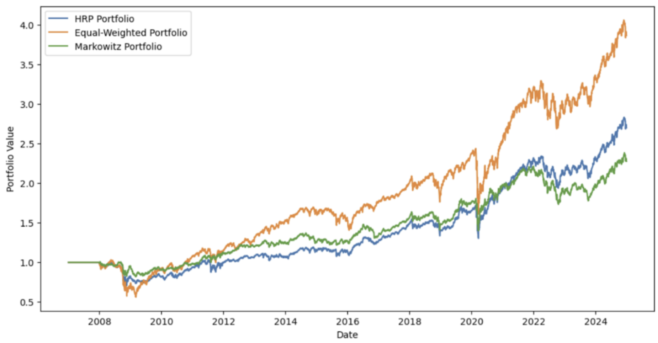


```
Final HRP Portfolio Value:      2.6993
Final Equal-Weighted Portfolio Value: 3.8645
Final Markowitz Portfolio Value: 2.2748
             Avg Yearly Return  Std Yearly Return  Sharpe Ratio (Yearly)  \
HRP                   0.065743           0.107154               0.613540
Equal                 0.093296           0.145731               0.640191
Markowitz             0.053738           0.093643               0.573858

             Max Drawdown
HRP             -0.292596
Equal           -0.454011
Markowitz       -0.227385
```

Figure 8: Lookback 22, 1 year rebalance

Lookback: 63 days

Rebalancing: 1 week

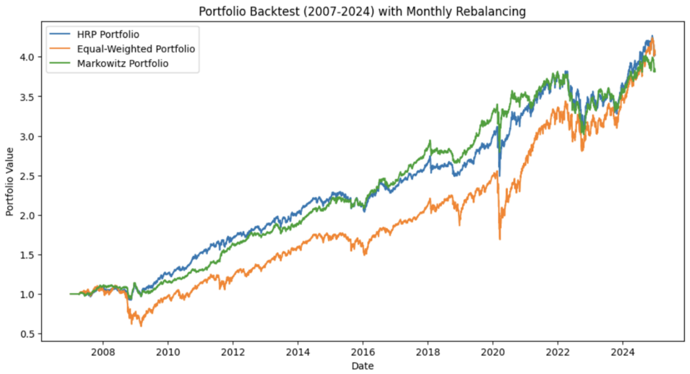


```
Final HRP Portfolio Value:      4.0217
Final Equal-Weighted Portfolio Value: 4.0358
Final Markowitz Portfolio Value: 3.8106
           Annualized Return  Annualized Volatility  Sharpe Ratio  \
HRP                 0.080551               0.095634      0.842280
Equal               0.080761               0.166768      0.484274
Markowitz           0.077312               0.081657      0.946793

           Max Drawdown
HRP           -0.201797
Equal         -0.466950
Markowitz     -0.208818
```

Figure 9: Lookback 63, 1 week rebalance

Lookback: 63 days

Rebalancing: 1 months

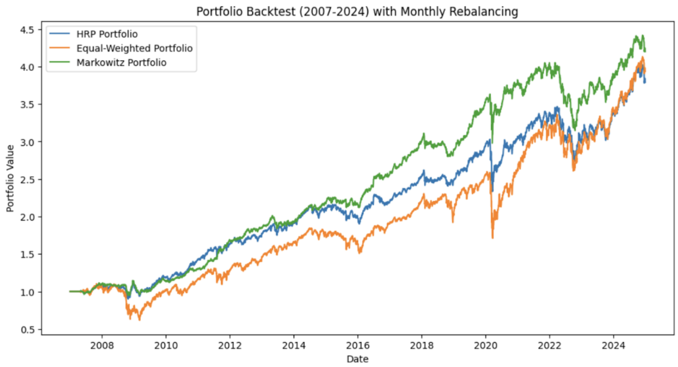


```
Final HRP Portfolio Value:      3.7911
Final Equal-Weighted Portfolio Value: 3.9345
Final Markowitz Portfolio Value: 4.2007
           Annualized Return  Annualized Volatility  Sharpe Ratio  \
HRP                 0.077005               0.104300      0.738300
Equal               0.079233               0.172386      0.459627
Markowitz           0.083173               0.085826      0.969085

           Max Drawdown
HRP           -0.229136
Equal         -0.439975
Markowitz     -0.223328
```

Figure 10: Lookback 63, 1 month rebalance

Lookback: 63 days

Rebalancing: 1 year

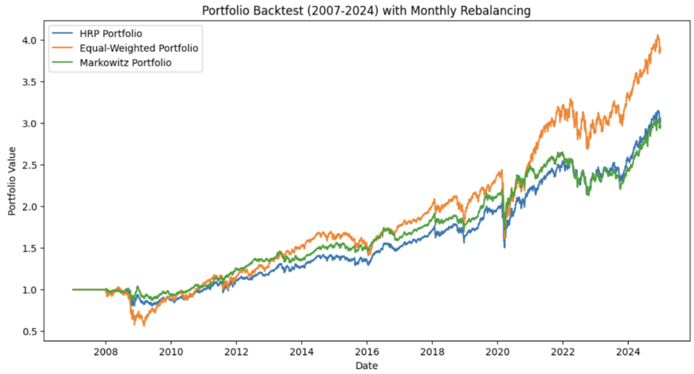


```
Final HRP Portfolio Value:      3.0162
Final Equal-Weighted Portfolio Value: 3.8645
Final Markowitz Portfolio Value: 2.9439
           Annualized Return  Annualized Volatility  Sharpe Ratio  \
HRP                 0.063384               0.113953      0.556224
Equal               0.078155               0.164857      0.474075
Markowitz           0.061947               0.097752      0.633720

           Max Drawdown
HRP           -0.264199
Equal         -0.454011
Markowitz     -0.220979
```

Figure 11: Lookback 63, 1 year rebalance

Lookback: 252

Rebalancing: 1 week

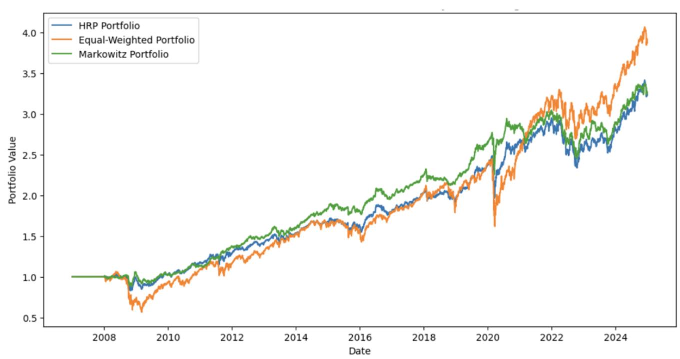


```
Final HRP Portfolio Value:      3.2342
Final Equal-Weighted Portfolio Value: 3.8695
Final Markowitz Portfolio Value: 3.2748
           Avg Yearly Return  Std Yearly Return  Sharpe Ratio (Yearly)  \
HRP                 0.074787           0.084535               0.884685
Equal               0.093066           0.145850               0.638090
Markowitz           0.075547           0.085224               0.886448

           Max Drawdown
HRP           -0.209905
Equal         -0.466950
Markowitz     -0.193779
```

Figure 12: Lookback 252, 1 week rebalance

Lookback: 252 days

Rebalancing: 1 year

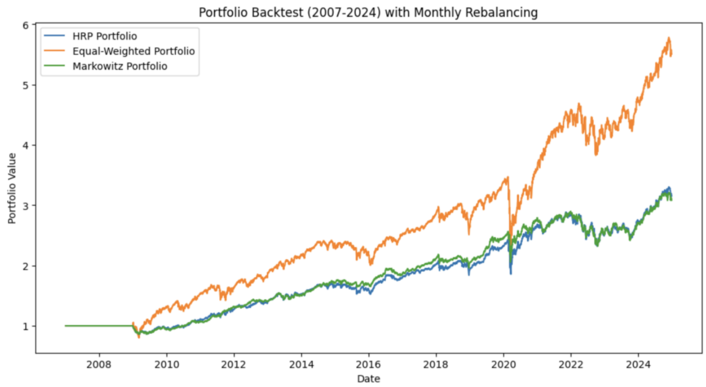


```
Final HRP Portfolio Value:      3.1534
Final Equal-Weighted Portfolio Value: 5.5067
Final Markowitz Portfolio Value: 3.0837
           Avg Yearly Return  Std Yearly Return  Sharpe Ratio (Yearly)  \
HRP                 0.073790           0.092169               0.800601
Equal               0.111070           0.112271               0.989309
Markowitz           0.072711           0.096513               0.753380

           Max Drawdown
HRP           -0.236738
Equal         -0.339487
Markowitz     -0.199739
```

Figure 13: Lookback 252, 1 year rebalance