TIC

Quantitative Trading

A FRAMEWORK FOR BUILDING A ROBUST PORTFOLIO

By Afonso Barão

By Martim Castro

On the task of constructing a portfolio, one has a whole universe of possibilities. From ETFs to bonds, to individual stocks, there is plenty to choose from. Moreover, there is an uncountable infinite number of possible ways to combine different assets.

This article proposes a long-only, with no leverage, portfolio construction framework based on Hierarchical Risk Parity (HRP). It combines hierarchical clustering, intra-cluster allocation methodologies, and inter-cluster Risk Parity.

In order to validate the framework, a rolling Walk Forward validation is employed. Out-of-sample performance is then evaluated against benchmark strategies.

Hierarchical Clustering

Put simply, a cluster is a collection of groups composed of different objects, which, in this case, are assets, organized with the goal of bringing together those that share similar characteristics. To construct it, one finds the necessity of defining a distance metric, a measure that will dictate how close (or similar) two objects are from each other. In this case, the metric used is:

$$d(x_i, x_j) = \sqrt{\frac{1}{2}(1 - \rho_{ij})}$$

where ρ_{ij} is the Pearson correlation between the log-returns of the two assets.

Once the distance metric is defined, it is necessary to specify a linkage function, which determines how distances between clusters are computed during the hierarchical aggregation process. Different linkage criteria lead to different cluster structures and stability properties.

In this framework, average linkage is employed. Under this criterion, the distance between two clusters A and B is defined as the average pairwise distance between all assets in the two clusters:

$$d(A, B) = \frac{1}{|A|\,|B|} \sum_{i \in A} \sum_{j \in B} d(x_i, x_j)$$

where |A| and |B| denote the number of assets contained in clusters A and B, respectively. The normalization by |A||B| ensures that the distance between clusters is computed as the average of all pairwise distances, rather than being dominated by cluster size or extreme individual distances.

Average linkage provides a balanced compromise between single and complete linkage, reducing sensitivity to extreme pairwise distances. This leads to more compact and stable clusters, while remaining relatively robust to isolated outliers in asset correlations.

The hierarchical clustering process is represented by a dendrogram, which is constructed iteratively. At each step, the two clusters that minimize the linkage distance *d(A, B)* are merged into a single cluster. This procedure is repeated until all assets are aggregated into a single cluster.

Figure 1 illustrates an example of hierarchical clustering using the Euclidean distance. Points that are closer in the feature space are merged first in the dendrogram, while the most distant points are combined only at later stages.

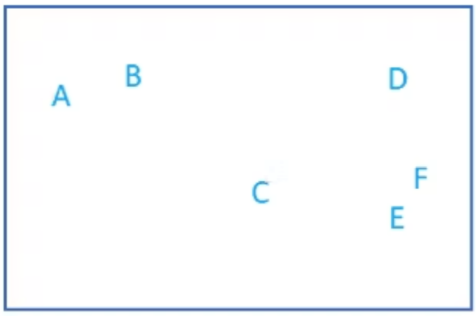

TIC Quantitative Trading

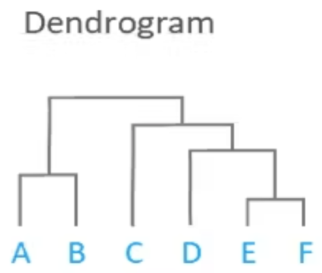

Fig. 1

Intra Clusters Portfolios

Once all assets are assigned to their respective clusters, the next step is to construct a portfolio within each cluster. This step consists of determining the allocated capital to each asset within its family of similar assets. The allocation rule introduced by López de Prado (2016) in the original HRP framework is:

$$w_i = \frac{C}{\sigma_i^2},$$

where w_i is the weight of asset i, σ_i^2 is its log-returns variance and C is a constant. Since all weights sum to 1:

$$\sum_{i=1}^{N} \frac{C}{\sigma_i^2} = C \sum_{i=1}^{N} \frac{1}{\sigma_i^2} = 1$$

with N being the number of assets in the cluster. This implies:

$$C = \left(\sum_{i=1}^{N} \frac{1}{\sigma_i^2} \right)^{-1}$$

for this specific formulation.

Some limitations of this rule include the fact that it does not explicitly minimize intra-cluster portfolio variance, as correlations between assets within the same cluster are ignored. This motivated alternative approaches such as the Minimum Variance optimization, whose downfall relates to the necessity of inverting the covariance matrix, leading to unstable weigh estimations. In order to balance stability and diversification, a Risk Parity intra-cluster allocation is adopted.

Risk Parity Portfolio

Having the intra-cluster portfolios constructed, the final step consists of building the global portfolio using a Risk Parity allocation. Risk Parity is a portfolio allocation approach in which asset weights are chosen such that each asset contributes equally to the total portfolio risk. In a traditional Risk Parity framework, each asset risk contribution is measured through volatility or marginal risk contributions derived from the covariance matrix.

Let $\mathbf{w} \in \mathbb{R}^n$ denote the portfolio weight vector and Σ the covariance matrix of asset returns. The total portfolio volatility is given by

$$\sigma_p = \sqrt{w^T \Sigma w}.$$

The marginal risk contribution of asset is defined as

$$MRC_i = \frac{\partial \sigma_p}{\partial w_i} = \frac{(\Sigma w)_i}{\sigma_p}.$$

The risk contribution of asset to total portfolio risk is then

$$RC_i = w_i \cdot MRC_i = \frac{w_i (\sum w)_i}{\sigma_p}$$

A Risk Parity portfolio is defined by the condition that all assets contribute equally to total portfolio risk:

$$RC_1 = RC_2 = \cdots = RC_N = \frac{\sigma_p}{N}$$

This condition is typically enforced numerically by solving a constrained optimization problem under full investment and long-only constraints.

In the context of Hierarchical Risk Parity (HRP), the same principle is applied at a higher level of aggregation. Rather than allocating risk directly across individual assets, each cluster is treated as a single synthetic asset, whose risk is defined by the covariance structure of the assets it contains and the previously computed intra-cluster weights. This creates a direct bridge to Traditional Risk Parity (TRP), while operating on a reduced and more structured risk space.

HRP vs Risk Parity

The used hierarchical formulation in HRP substantially improves robustness over Traditional Risk Parity. Estimating a full covariance matrix at the asset level, with a big number of assets, is notoriously unstable. By first clustering correlated assets and aggregating them into clusters, HRP reduces estimation error. As a result, the Risk Parity allocation is performed on a smaller, better conditioned covariance matrix, leading to more stable and interpretable portfolio weights.

It is also economically more intuitive to first group assets according to their correlations, thereby forming clusters of assets that share similar characteristics. For instance, when constructing an HRP portfolio using only S&P 500 equities, a natural choice is to consider eleven clusters, corresponding to the eleven S&P sectors. Since equities within the same sector tend to be highly correlated, the hierarchical clustering procedure naturally groups them together.

These clusters can then be treated as single aggregated assets in the subsequent allocation step. Continuing with the same example, the final step consists of constructing a Risk Parity portfolio across sectors. That is, the resulting portfolio is designed so that each sector contributes equally to the overall portfolio risk.

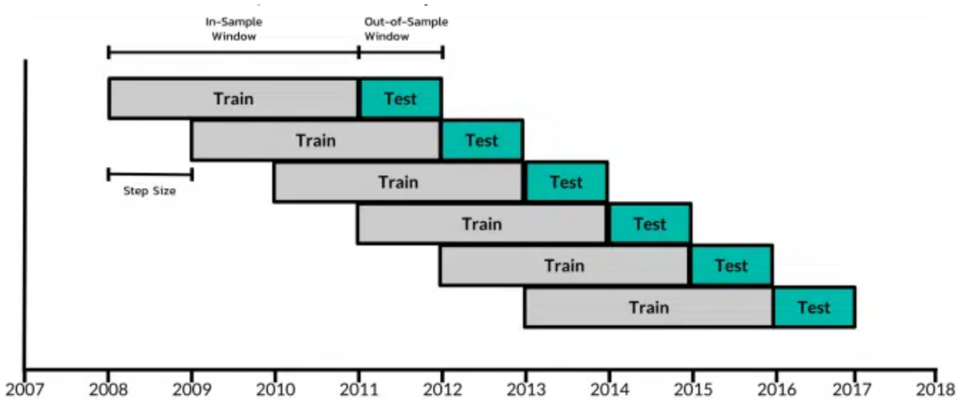

Fig. 2

Walk Forward Validation

To validate the framework and replicate real time asset weight estimation and deployment on unseen data, a rolling Walk Forward validation is employed. This is a validation procedure in which the model is trained using only past information and evaluated on future, unseen data, strictly preserving the temporal order.

In practice, a training window is used to estimate the model parameters, in this case the portfolio weights. These parameters are then applied to the immediately following out-of-sample period. The window is rolled forward in time and the process is repeated until the end of the sample.

As seen in Figure 2, one can split the data into blocks of one year. Then, use the previous three years to obtain portfolio weights and the next year for out-of-sample backtesting. This approach simulates real time deployment, prevents look ahead bias, and allows the assessment of temporal stability.

Framework Implementation

In order to implement the proposed framework, a portfolio was constructed under the following specifications:

- The asset universe is constituted by a diversified set of instruments, including commodities (metals, energy, and agricultural products), S&P 500 equities, and the eight cryptocurrencies with the biggest market capitalization. Equity returns are adjusted for dividends.
- The choice of this asset universe has the purpose of maintaining primary exposure to the U.S. equity market while improving diversification through the inclusion of asset classes that historically exhibit low or negative correlations with equities. Commodities serve as a hedge against stock market risks, while the inclusion of cryptocurrencies provides limited exposure to an emerging asset class.
- In the Walk Forward procedure, the historical data is split in blocks of quarters, corresponding to a quarterly rebalancing of portfolio weights. The first out-of-sample evaluation period begins in the first quarter of 2017.
- To prevent data leakage and ensure a realistic validation, the equity universe is redefined at each backtesting period based on the S&P 500 constituents at that time.
- Out-of-sample results from successive windows are then concatenated to simulate a continuous backtesting from that date up to the present.
- For the computation of Sharpe ratios, annualized returns and volatilities were used, with a return's offset of 3%.

Results

Backtesting results are evaluated against two baselines and two benchmarks:

- Risk Parity portfolio (RP)
- Equal weighted portfolio (EWP)
- S&P 500 with dividend adjusted returns
- Gold spot price

The results indicate that HRP clearly outperforms the benchmarks. It achieves the highest annualized return (22.05%), with a substantial margin relative to all other baselines, while simultaneously maintaining the second lowest volatility. This combination translates into the highest risk adjusted performance, reflected by Sharpe and Sortino ratios.

From a downside risk perspective, HRP also exhibits good behavior. Although Gold displays the lowest maximum drawdown, HRP experiences smaller drawdowns than the other portfolios. Particularly in COVID-19's market shock, HRP experiences a considerably smaller drawdown (28.33%) than the equity market (33.72%).

TIC Quantitative Trading

Portfolio	Returns % (annualized)	Volatilities % (annualized)	Sharpe	Sortino	Max Drawdown %
HRP	22.05	14.90	1.278	1.776	28.33
RP	12.65	15.89	0.607	0.814	36.83
EWP	13.97	19.22	0.571	0.785	39.81
S&P 500	15.04	18.45	0.653	0.898	33.72
Gold	17.47	14.75	0.981	1.413	22.00

Additionally, these results reinforce the advantage of Hierarchical Risk Parity over traditional Risk Parity. The inclusion of the clustering process proves to provide better out-of-sample results.

Transcation costs

It's important to note that transaction costs were not included in the backtesting. When selecting a rebalancement frequency, one has to face a tradeoff between transaction costs and deviations of asset weights from their target values.

On one hand, as assets accumulate returns between rebalancing dates, portfolio weights drift away from the optimal allocation. Assets with higher cumulative returns gain weight, while those with the lowest returns lose weight. The longer the rebalancing interval, the larger these deviations become.

On the other hand, choosing very short rebalancing intervals increases turnover, causing portfolio performance to deteriorate due to transaction costs. A lower bound for the portfolio value after accounting for transaction costs is given by:

$$P_{com} \geq P\ (1-t)^{2N},$$

where $\mathrm{P}_{\mathrm{com}}$ is the portfolio value accounting transaction costs, P is the portfolio without transaction costs, t is the transaction cost rate, and N is the total number of rebalancing events.

Considering quarterly rebalancing over nine years and a commission rate of 0.05%:

$$P_{com} \geq P(1-t)^{2N}$$
$$= P(1-0.0005)^{2\cdot(4\cdot 9)}$$
$$\approx 0.965P.$$

This implies that, even in a worst-case scenario, the portfolio value after accounting for transaction costs would be approximately 96.5% of the value obtained when transaction costs are ignored.

Article written in February and updated on March.