

A Meta-Method for Portfolio Management Using Machine Learning for Adaptive Strategy Selection

Damian Kisiel

University College London
Department of Computer Science
London, United Kingdom
d.kisiel@cs.ucl.ac.uk

Denise Gorse

University College London
Department of Computer Science
London, United Kingdom
d.gorse@cs.ucl.ac.uk

ABSTRACT

This work proposes a novel portfolio management technique, the *Meta Portfolio Method* (MPM), inspired by the successes of meta approaches in the field of bioinformatics and elsewhere. The MPM uses XGBoost to learn how to switch between two risk-based portfolio allocation strategies, the Hierarchical Risk Parity (HRP) and more classical Naïve Risk Parity (NRP). It is demonstrated that the MPM is able to successfully take advantage of the best characteristics of each strategy (the NRP’s fast growth during market uptrends, and the HRP’s protection against drawdowns during market turmoil). As a result, the MPM is shown to possess an excellent out-of-sample risk-reward profile, as measured by the Sharpe ratio, and in addition offers a high degree of interpretability of its asset allocation decisions.

CCS CONCEPTS

• **Computing methodologies** → **Classification and regression.**

KEYWORDS

Meta-methods, Machine learning, Supervised learning, Portfolio management, Adaptive strategy selection

1 INTRODUCTION

Portfolio management as we know it today has been heavily influenced by the work of Harry Markowitz [15], who formalized the idea of the risk-return trade-off, where investors are rewarded for taking on more risk. These portfolios, however, suffer from significant flaws [13, 21]. Institutional investors who used the standard Markowitz allocations suffered large drawdowns during recent crises; as a result, the model was criticized [18] and risk-based portfolio allocation techniques became more popular. However, these newer methodologies, despite having rigorous underpinning mathematical structures, are not supported by any financial theory that would promote them as the optimal portfolio choice.

Therefore, rather than trying to discover a risk-based methodology that would work ideally well for all scenarios [3], the work of this paper aims to choose the best performing strategy for given market conditions. It proposes a new portfolio construction technique called the *Meta Portfolio Method* (MPM), inspired by the successes of meta approaches in the field of bioinformatics [2, 9], among other areas. The MPM combines two risk-based portfolio allocation strategies, namely the *Hierarchical Risk Parity* (HRP) and more classical *Naïve Risk Parity* (NRP), by allocating capital to the one predicted to provide better risk-adjusted performance in the next investment period.

2 BACKGROUND

2.1 Meta Methods

The idea behind meta approaches is that the performance of a task can often be greatly enhanced by combining multiple methodologies, as opposed to using any of the constituent techniques alone. Meta approaches have been successfully applied in the field of bioinformatics [2, 9], but have not yet been widely explored in portfolio management. An exception is provided by [11], in which the authors did identify key features of the variance-covariance matrix that help in studying the behavior of different risk-based allocation strategies. However, this was an ex-post analysis where the model could only make retrospective decisions as to which strategy performed better. What is lacking from the existing literature is an ex-ante analysis that would be able to recommend which of the strategies should be followed in the future, given a set of statistical characteristics of a given asset universe. The research described here aims to bridge this gap.

2.2 Risk-Based Portfolio Construction Strategies

As mentioned in the introduction, the meta-strategy of this work combines two risk-based portfolio construction strategies, Naïve Risk Parity (NRP) and Hierarchical Risk Parity (HRP). This section will give a brief introduction to each of these, followed by a comparison of their relative strengths and weaknesses.

2.2.1 Naïve Risk Parity (NRP). Risk parity portfolios aim to maximize the portfolio diversification level in the hope that a more diversified portfolio will be less vulnerable to unfavorable market conditions. The Naïve Risk Parity [18] method is considered to be the simplest implementation of a risk-based methodology, and aims to equalize the risk contributed to the overall portfolio by each asset. Its portfolio weights can be computed as follows

$$w_i = \frac{1}{\sum_{j=1}^n \frac{1}{\sigma_j^2}}, \quad (1)$$

where σ_i^2 denotes the variance of returns of asset i .

2.2.2 Hierarchical Risk Parity (HRP). Hierarchical Risk Parity [17], similar to other risk-budgeting approaches, considers the variance-covariance matrix during its construction process. However, it in addition makes use of hierarchical clustering to group together similar assets, to avoid computing correlation estimates between all pairs of assets. This process uses three steps: tree clustering, quasi-diagonalization, and recursive bisection.

In the first step the algorithm performs tree clustering by first transforming the correlation matrix into a distance matrix D , where each element is given by

$$d_{i,j} = \sqrt{\frac{1}{2}(1 - \rho_{i,j})}. \quad (2)$$

The newly created matrix D is then transformed by computing the Euclidean distance between all column pairs as follows

$$\tilde{d}_{i,j} = \sqrt{\sum_{n=1}^N (d_{n,i} - d_{n,j})^2}, \quad (3)$$

creating a new matrix representation which gives the correlation of each asset with the rest of the portfolio. This allows the formation of a tree structure according to the $\tilde{d}_{i,j}$ ‘scores’. In the second step, quasi-diagonalization, the tree structure created in the previous step is used to reorganize the rows and columns of the correlation matrix in such a way as to make similar assets sit next to each other and investments considered to be different be placed further apart. The last step performs recursive bisection wherein the standard inverse-variance methodology is applied to the clusters defined in the previous steps. Because asset weights are assigned in a top-down fashion, according to each sub-cluster’s variance, only assets within the same cluster compete for allocation, achieving diversification benefits across all levels of the tree structure.

2.2.3 Comparison of the Component Strategies. Each component strategy displays a different realized volatility profile, as seen in Figure 1. HRP gives a portfolio with a much lower variance. It offers a return stream with heavily damped tails which makes it a safer option. Some negative tails can be spotted but these are much less common and less pronounced than those in NRP’s profile, which displays heavy tails in its return distributions. However, since fat tails also mean a higher probability of a large positive return, NRP has the potential of being considerably more profitable than HRP during certain market conditions.

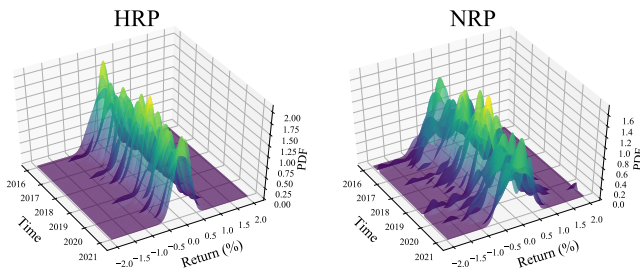


Figure 1: 3D plots showing changes in the probability density function over time.

3 METHODOLOGY

3.1 Data Selection & Preparation

Results presented in this work are based on ten different investment universes, each containing five distinct asset classes. A set of 18 exchange traded funds (ETFs), as shown in Table 1, is used as a basket from which assets can be drawn with replacement to construct a

Table 1: Basket of ETFs.

Ticker	Name	Ticker	Name
AGG	Aggregate Bonds	VNQ	Real Estate
EEM	Emerging Markets	XLV	Health Care Sector
EFA	Developed Markets	IWD	Value Equities
EWJ	Japanese Equities	IWM	Small-Cap Equities
GLD	Gold	LQD	Corporate Bonds
QQQ	Nasdaq-100 Index	TLT	20+ Year T-Bonds
SHY	1-3 Year T-Bonds	VO	Mid-Cap Equities
SPY	S&P 500 Index	VBR	Small-Cap Value
TIP	I-L Bonds	XLY	Consumer Services

specific asset universe. The selection routine is constrained so as to allow each universe to share at most two ETFs with any other asset universe. The reason for selecting market indices, rather than individual assets, is that they are generally uncorrelated and hence offer higher diversification benefits. The list of ETFs shown in Table 1 has been chosen in such a way as to represent as many distinct asset classes as possible. Data collected on each ETF includes daily returns that span the period from November 2004 until May 2021. All assets are denominated in US Dollars, the domestic currency of the investors considered in this study.

The original data contains daily price information; since we are interested in the growth of each universe component, the simple (arithmetic) return from time $t - 1$ to time t is computed as follows

$$r_{n,t} = \frac{P_{n,t}}{P_{n,t-1}} - 1, \quad (4)$$

where $P_{n,t}$ is the price of asset n at time (day) t . If we have N assets in a portfolio and their weights at time $t - 1$ are given by $w_{1,t-1}, w_{2,t-1}, \dots, w_{N,t-1}$, the portfolio returns are computed using

$$R_t = \sum_{n=1}^N w_{n,t-1} r_{n,t}. \quad (5)$$

In order to simplify calculations, it is assumed the risk-free return is zero throughout the investment period and hence the return above is our variable of interest.

3.1.1 Computation of Return Covariance Matrix. Both of the component strategies (HRP and NRP) require an estimate of the variance-covariance matrix between each pair of asset return series. This matrix will be used both for the computation of the features to be input to the machine learning model (next section) that selects the strategy to be used in the next time period, and for the internal construction of portfolio weights. The multivariate GARCH Dynamic Conditional Correlation (DCC) model [7] is employed here to compute the variance-covariance matrix. The DCC model is well-regarded because it allows for time-varying correlations between assets.

3.2 Computation of Features

The set of predictive features includes measures that aim to describe each constituent strategy’s most recent performance, such as average return, realized volatility [10], maximum drawdown [4] and

downside deviation [20], as well as a given investment universe’s mean return and standard deviation.

In addition, more elaborate transformations of the variance-covariance matrix are carried out. For example, *meanCORR* and *stdCORR* compute the average and the standard deviation, respectively, of all lower triangular elements of the correlation matrix and indicate overall level and heterogeneity of correlations at a given point in time. Other features include the *non-parametric k-nearest neighbor entropy estimator* [12], the *quality ratio* [1], that aims to capture how much diversification benefit each investment universe is able to offer, and the *standardized generalized variance* [16], which measures the overall volatility of a specific asset universe.

As described in section 2.2, the two component strategies (HRP and NRP) differ in the way they are constructed. Since the HRP considers an additional clustering step, its performance depends on the level of the hierarchical structure present in a given asset universe, measured by the *cophenetic correlation coefficient* [19], with the quality of clustering computed using the intra-cluster variance [8]. Finally, a set of features characterizing some properties of the correlation matrix are calculated, including its determinant, condition number [6] and the fraction of variance that can be explained by the eigenvalues that lie outside the Marchenko-Pastur distribution [14].

3.3 Outline of the Meta Portfolio Method (MPM)

The sample period denoted by T_r in Figure 2 is used as input to the multivariate GARCH DCC model that computes the variance-covariance matrix, which is then used in the construction of the two component strategies and to compute a set of model features. Observations collected during the brown sample period (denoted by T_m) in Figure 2 are used to calculate features that describe each component strategy’s most recent performance. Once all model features have been computed, the strategy selection process can take place.




Figure 2: Meta Portfolio Method (MPM) timeline.

At the beginning of the green band (denoted by T_h) in Figure 2 the machine learning model of the next section decides which of the two component strategies (HRP or NRP) should be followed in the next investment period. All available capital is then invested in the selected strategy and held during the holding period of length T_h and its returns recorded. In the next investment period (Period $k + 1$) the rolling windows are advanced by the time interval T_a and the two component strategies are separately rebalanced. Then, the process is repeated. All returns of the Meta Portfolio Method are

computed in a walk-forward manner so that the strategy selection algorithm has access only to the historical data, and transaction costs are taken into account when evaluating performance. The whole procedure is summarized in Algorithm 1, below.

Algorithm 1: Meta Portfolio Method (MPM)

```

1 import daily asset returns
2 initialise  $T_m, T_r, T_h$  and  $T_a$ 
3 while  $t < \text{last observation}$  do
4   compute var-cov matrix on  $T_r$  data using GARCH DCC
5   construct the two component strategies
6   compute XGBoost features
7   for  $i \leftarrow 1$  to  $T_m$  do
8     compute performance-related features
9   end
10  apply Bayesian hyperparameter optimization
11  predict the Sharpe ratio spread using optimized model
12  if  $\text{spread} > 0$  then
13    select HRP strategy;
14  else
15    select NRP strategy;
16  end
17  for  $i \leftarrow 1$  to  $T_h$  do
18    compute returns generated by the MPM
19  end
20  advance the sliding window by  $T_a$ 
21 end

```

3.4 Training Methodology

3.4.1 *Target Variable.* Since we are interested in predicting which of the two strategies is going to perform better (on a risk-adjusted basis) during the next investment period, the target variable is the *Sharpe ratio spread*, the difference between Sharpe ratios of the HRP and NRP strategies, given by

$$y_{\text{target}} = SR_{\text{HRP}} - SR_{\text{NRP}}. \quad (6)$$

The process of switching between the two component strategies can be formulated as a binary classification problem, since at any point in time we only have two choices at our disposal. However, we notice from equation 6 that the target variable is a real-valued number, which allows us to work in a regression setting. The advantage of this is that the magnitude of the Sharpe ratio spread carries additional information that a machine learning model can leverage during the training process. All real-valued predictions can then be converted to a binary allocation decision using the HRP strategy if $y_{\text{predicted}} \geq 0$, and the NRP strategy otherwise.

3.4.2 *ML Model and Hyperparameter Optimization.* XGBoost [5] is selected to solve the regression problem above. The main advantage of XGBoost, in this problem setting, comes from the fact that it is able to construct non-linear relations between the input features while, at the same time, preserving a high degree of interpretability through its built-in feature importance functions, an

Table 2: Performance comparison.

	Asset Universe	HRP	NRP	MPM
Sharpe Ratio	1	1.37	1.25	1.64
	2	1.20	1.17	1.32
	3	1.35	1.23	1.82
	4	1.07	1.03	1.22
	5	1.30	1.38	1.59
	6	1.68	1.34	1.87
	7	1.67	1.44	2.31
	8	1.12	1.07	1.41
	9	1.01	0.99	1.23
	10	0.84	0.87	1.05
Cumulative Return (%)	1	9.06	15.79	16.37
	2	60.50	74.84	82.98
	3	9.31	16.47	17.48
	4	66.23	95.52	116.60
	5	42.39	68.46	73.18
	6	11.19	18.05	19.73
	7	9.18	20.92	23.84
	8	64.89	84.11	107.14
	9	54.05	79.85	97.23
	10	22.50	51.00	60.12

important consideration when making portfolio allocation decisions. However, to fully harness the predictive power of XGBoost it is necessary to carry out careful hyperparameter optimization. Each set of model-building data (T_r sections in Figure 2) is split into a training part which represents 70% of observations and a test part which contains the remaining 30% of data. Only the training data is used for hyperparameter tuning purposes, with 5-fold cross-validation implemented to get the average regression scores across all folds. Due to the large number of hyperparameters, a classical grid search would be too time-consuming, and Bayesian hyperparameter optimization is therefore used instead to select the best set of hyperparameters for each asset universe.

4 RESULTS

4.1 Cumulative and Risk-Adjusted Returns

Table 2 gives a performance comparison between the Meta Portfolio Method (MPM) and its constituent strategies (HRP & NRP). The MPM shows the strongest performance in all investment universes, with Sharpe ratio improvements over the HRP and NRP having p-values of 6.01×10^{-9} and 2.45×10^{-7} , respectively, and cumulative return improvements having comparable p-values of 1.53×10^{-7} and 4.50×10^{-4} . The improvements provided by the MPM are therefore clearly statistically significant.

Figure 3 shows an example wealth curve for investment universe 9 (a middle-performing example), over the entire investment horizon. The MPM (green circle) is able to harness the fast growth opportunity offered by the NRP (red triangle) during market up-trends while at the same time offering a good level of protection

presented by the HRP (black square) during more turbulent periods, such as the onset of the COVID-19 pandemic in 2020. It is by such means the MPM achieves the highest cumulative returns and Sharpe ratios, as evidenced by the results in Table 2.




Figure 3: Comparison of cumulative returns for investment universe 9.

4.2 Feature Importance

One benefit of XGBoost [5], as previously mentioned, is that it is relatively easy to extract information about feature importance, which allows a better understanding of what is driving the output of the MPM. Figure 4 shows a boxplot of the importance scores for each feature across all investment universes. We notice that among the top three features is ‘hrp_down_dev’, the downside deviation of the HRP strategy. This makes intuitive sense since this feature informs the model about the historical downside risk experienced by the HRP strategy and hence the overall risk at any point in time. What is less obvious, however, is the high importance score assigned to ‘intra_cluster_var’ and ‘cophenetic_average’. These features correspond to the intra-cluster variance and the cophenetic correlation coefficient, respectively, and, as described earlier, quantify the level and the complexity of the hierarchical structure in a given investment universe. However, since this information is taken into account when constructing the HRP strategy, it is understandable these two features play a big role in predicting the relative performance of the HRP strategy.

5 CONCLUSIONS

This work has introduced the Meta Portfolio Method (MPM), which uses XGBoost to combine two different portfolio allocation paradigms, the Hierarchical Risk Parity (HRP) and the Naïve Risk Parity (NRP), both frequently used in practice, into a single portfolio allocation framework. The MPM is able to take advantage of both the fast growth opportunities offered by the NRP in favorable market conditions, and the protection against large drawdowns during market downturns offered by the HRP. As a result, the MPM gives substantially higher cumulative returns than either of the component strategies individually, while also maintaining the highest level of risk-adjusted return, as measured by the Sharpe ratio. Moreover, the use of XGBoost within the MPM results in an easily interpretable model, desirable to industry practitioners.

The performance of the MPM strategy has been analyzed on ten different multi-class asset universes, giving confidence the MPM is




Figure 4: Boxplot with feature importances.

not overfitted to a particular data set. It should be noted it would not be possible to construct significantly more asset universes based only on ETFs. However, while ETFs are desirable in that these assets can be chosen to be less correlated in performance than, say, individual equities, it is not an absolute necessity to use them, and future work will look at the alternative of having larger asset universes composed of less weakly correlated components.

REFERENCES

- [1] Adam Butler. 2018. *Portfolio Optimization: Simple versus Optimal Methods*. Retrieved June 18, 2021 from <https://www.advisorperspectives.com/articles/2018/10/15/portfolio-optimization-simple-versus-optimal-methods>
- [2] Phasit Charoenkwan, Chanin Nantasenamat, Md Mehedi Hasan, and Watshara Shoombuatong. 2020. Meta-iPVP: a sequence-based meta-predictor for improving the prediction of phage virion proteins using effective feature representation. *Journal of Computer-Aided Molecular Design* 34, 10 (2020), 1105–1116.
- [3] Denis Chaves, Jason Hsu, Feifei Li, and Omid Shakernia. 2011. Risk parity portfolio vs. other asset allocation heuristic portfolios. *The Journal of Investing* 20, 1 (2011), 108–118.
- [4] Alexei Chekhlov, Stanislav Uryasev, and Michael Zabarankin. 2005. Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance* 8, 1 (2005), 13–58.
- [5] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [6] Alan K. Cline, Cleve B. Moler, George W. Stewart, and James H. Wilkinson. 1979. An estimate for the condition number of a matrix. *SIAM J. Numer. Anal.* 16, 2 (1979), 368–375.
- [7] Robert Engle. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics* 20, 3 (2002), 339–350.
- [8] Maria Halkidi, Yannis Batistakis, and Michalis Vazirgiannis. 2001. On clustering validation techniques. *Journal of Intelligent Information Systems* 17, 2 (2001), 107–145.
- [9] Takashi Ishida and Kengo Kinoshita. 2008. Prediction of disordered regions in proteins based on the meta approach. *Bioinformatics* 24, 11 (2008), 1344–1348.
- [10] Marwan Izzeldin, M. Kabir Hassan, Vasileios Pappas, and Mike Tsionas. 2019. Forecasting realised volatility using ARFIMA and HAR models. *Quantitative Finance* 19, 10 (2019), 1627–1638.

- [11] Markus Jaeger, Stephan Krügel, Dimitri Marinelli, Jochen Papenbrock, and Peter Schwendner. 2021. Interpretable Machine Learning for Diversified Portfolio Construction. *The Journal of Financial Data Science* 3, 4 (June 2021).
- [12] Damiano Lombardi and Sanjay Pant. 2016. Nonparametric k-nearest-neighbor entropy estimator. *Physical Review* 93, 1 (2016), 013310.
- [13] Sébastien Maillard, Thierry Roncalli, and Jérôme Teiletche. 2010. The properties of equally weighted risk contribution portfolios. *The Journal of Portfolio Management* 36, 4 (2010), 60–70.
- [14] Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. 1967. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik* 114, 4 (1967), 507–536.
- [15] Harry Markowitz. 1952. Portfolio Selection. *The Journal of Finance* 7, 1 (1952), 77–91. <http://www.jstor.org/stable/2975974>
- [16] Dariush Najarzadeh. 2019. Testing equality of standardized generalized variances of k multivariate normal populations with arbitrary dimensions. *Statistical Methods & Applications* 28, 4 (2019), 593–623.
- [17] Marcos Lopez De Prado. 2016. Building diversified portfolios that outperform out of sample. *The Journal of Portfolio Management* 42, 4 (2016), 59–69.
- [18] Thierry Roncalli. 2013. *Introduction to risk parity and budgeting*. CRC Press.
- [19] Robert R. Sokal and F. James Rohlf. 1962. The comparison of dendrograms by objective methods. *Taxon* 11, 2 (1962), 33–40.
- [20] Frank A. Sortino and Lee N. Price. 1994. Performance measurement in a downside risk framework. *The Journal of Investing* 3, 3 (1994), 59–64.
- [21] Herold Ulf and Maurer Raimond. 2006. Portfolio choice and estimation risk. A comparison of Bayesian to heuristic approaches. *The Journal of the IAA* 36, 1 (2006), 135–160.