GV Invest
Short Studies Series

GV INVEST | 05

Does Naive Not Mean Optimal?

The Case for the 1/N Strategy in Brazilian Equities

December, 2016

Vinicius Esposito[i]

The development of optimal approaches to portfolio construction has rendered outdated the allocation of wealth into equally weighted assets. Since the groundbreaking work of Markowitz (1952), credited with the foundations of Modern Portfolio Theory, equally weighted portfolios have acquired a distinctive alias: naive diversification. This, of course, implies the availability of more efficient, attractive ways of constructing a portfolio. This paper aims at providing a simple analysis to whether this is true for the Brazilian stock market, in the likeness of DeMiguel, Garlappi and Uppal (2009), which study the out-of-sample performance of the sample-based mean-variance model relative to the naive $1/N$ portfolio. By analyzing the performance of four portfolios, the out-of-sample performance of the mean-variance portfolio rule (as well as a few other usual portfolio construction methods) is evaluated against the performance of the equally weighted portfolio rule.

The equally weighted portfolio (also referred to as the naive or $1/N$ portfolio) is the one in which each asset receives the same share of wealth at each rebalancing date. Thus, $1/N$ of the wealth is invested in each of the N assets available to the investor at each time period. This strategy is used in the same way as DeMiguel, Garlappi and Uppal (2009): that is, as a benchmark for the assessment of some of the most widely used portfolio rules, and its pivotal role in this study is justified by its ease of implementation (no estimation of moments is required as well as no math other than a simple division) and also by the fact that many practitioners, even with today's advanced computational methods, continue to use simple rules which more than often are mere rules of thumb.

All datasets are retrieved from NEFIN, the Brazilian Center for Research in Financial Economics of the University of Sao Paulo, with the exception of the dataset of country indices, which is obtained directly from MSCI's public index information. The time series used in this study consist of $T = 189$ months and the estimation window comprises $M = 24$ months. All returns are computed in excess of the risk-free rate, which is the 30-day DI swap for Brazilian assets and the nominal US T-bill for the country indices.

Table 1: Datasets considered

Dataset	Source	# of assets	Abbreviation
Seven industry portfolios	NEFIN	7	Sectors
Five Latin American indices	MSCI	5	Latam
Five risk factor portfolios	NEFIN	5	Factors
Four size- and book-to-market portfolio and the Brazilian market portfolio	NEFIN	5	FF

Asset-Allocation Models Considered

In this short study, basic portfolio rules that are common for both the academic and the practitioner are assessed. Since the idea is to present a simple overview of the optimality of the $1/N$ portfolio, the author refrains from using more complex approaches that take into account the estimation error as the sample for Brazilian equities is not long over time.

A few definitions: the N-vector of excess returns is denoted by R_t, while μ_t is used to describe the expected returns on the risky assets (also in excess of the risk-free rate). The $N{\times}N$ variance-covariance matrix of returns is defined by Σ_t. Let M denote the length over which the moments are estimated and T be the total length of the time series. Furthermore, let x_t indicate the vector of portfolio weights invested in the N risky assets (leaving $1 - \mathbf{1}_N^\intercal \boldsymbol{x_t}$ to the risk-free asset. Then, the vector of relative weights in the risky portfolio is $w_t = x_t / |\mathbf{1}_N^\intercal \boldsymbol{x_t}|$, where $\mathbf{1}_N$ is a N-vector of ones. This makes sure that the direction of positions is preserved even in the case where the sum of weights is negative.

As mentioned previously, the benchmark model for this analysis will be the $1/N$ portfolio with monthly rebalancing. Considering a portfolio with N assets, the weights are thus defined as: $w_t = N^{-1}\mathbf{1}_N$.

Following the usual approach in the literature, it is considered a mean-variance optimizing investor who selects a weight vector in order to maximize their utility given by:

$$\max_{x_t} \quad x_t^\intercal \mu_t - \frac{\gamma}{2} x_t^\intercal \Sigma_t x_t$$

Taking into account the solution $x_t = (1/\gamma)\,\Sigma_t^{-1}\mu$, the vector of relative portfolio weights invested in the N risky assets at time t then is:

$$w_t = \frac{\Sigma_t^{-1}\mu_t}{\mathbf{1}_N^\intercal\, \Sigma_t^{-1}\mu_t}$$

This result holds true for all procedures that involve an optimization, so that it is only necessary to plug in both the estimates of μ_t and Σ_t according to each strategy.

Apart from the equally weighted strategy, the first model to be estimated is the sample-based mean-variance portfolio, which follows the procedures first established by Markowitz (1952). The investor solves the aforementioned problem using the sample estimates, $\hat{\mu}_t$ and $\hat{\Sigma}_t$. This approach, however, has a significant caveat: it ignores estimation error, a limitation which will not be directly dealt with in this study.

The third portfolio rule estimated is the minimum-variance portfolio. This approach is similar to the previous one in the sense that it also uses sample estimates, but it completely ignores estimates of expected returns. Thus, the investor's problem may be simplified to:

$$\min_{w_t} \quad w_t^{\intercal} \Sigma_t w_t, \qquad s.t.\ \mathbf{1}_N^{\intercal} w_t = 1$$

The final strategy is a simplified risk parity rule, in which an asset's weight is inversely proportional to its volatility, thus $w_{it} \propto {}^{1}/_{\sigma_{it}}, \mathbf{1}_N^{\intercal} w_t = 1$. This is a simplified version of the risk parity strategy because it assumes that assets are uncorrelated with each other, which makes it easier to compute a portfolio where assets have equal contribution to portfolio risk.

Performance Evaluation Metrics

Following the approach of DeMiguel, Garlappi and Uppal (2009), the analysis is based on rolling-samples. In the case of this study, an estimation windows of length $M = 24$ will be used. Then, for each time period, the returns from the previous 24 months will be used in the estimation of the parameters demanded in the execution of a particular rule. This results in a series of $T - M$ out-of-sample data that can be used to gauge each strategy. It is a significantly small sample due to the relatively infant nature of the Brazilian stock market.

In order to measure the performances of each portfolio rule, the first metric to be calculated is the out-of-sample Sharpe ratio, defined as the sample mean of out-of-sample excess returns divided by their sample standard deviation. That is, for each strategy k:

$$\widehat{SR}_k = \frac{\hat{\mu}_k}{\hat{\sigma}_k}$$

In order to test if the Sharpe ratios of two rules are statistically different, it is presented the p-value of the difference between a specific strategy and the $1/N$ strategy, which is computed according to the approach of Jobson and Korkie (1981) with the correction of Memmel (2003).[1]

It is of interest to have a notion of the effect of estimation error on performance. Therefore, the in-sample Sharpe ratio is also calculated for each strategy using the whole time series of returns (i.e. with an estimation window $M = T$). For the portfolio rule k, this is defined as:

[1] This statistic is computed as follows: given two portfolios 1 and 2 with means, variances and covariance given by $\hat{\mu}_1, \hat{\mu}_2$, $\hat{\sigma}_1$, $\hat{\sigma}_2$, $\hat{\sigma}_{1,2}$ estimated over a time horizon of $T - M$, the hypothesis $H_0: \widehat{SR}_1 - \widehat{SR}_2 = 0$ is tested by:

$$\hat{z}_{JK} = \frac{\hat{\sigma}_2\hat{\mu}_1 - \hat{\sigma}_1\hat{\mu}_2}{\sqrt{\hat{\vartheta}}}, \text{ where}$$

$$\hat{\vartheta} = \frac{1}{T-M}\left(2\,\hat{\sigma}_1^2\hat{\sigma}_2^2 - 2\,\hat{\sigma}_1\,\hat{\sigma}_2\hat{\sigma}_{1,2} + \frac{1}{2}\hat{\mu}_1^2\hat{\sigma}_2^2 + \frac{1}{2}\hat{\mu}_2^2\hat{\sigma}_1^2 - \frac{\hat{\mu}_1\,\hat{\mu}_2}{\hat{\sigma}_1\,\hat{\sigma}_2}\hat{\sigma}_{1,2}^2\right)$$

Which is asymptotically distributed as a standard normal.

$$\widehat{SR}_k^{IS} = \frac{\hat{\mu}_k^{IS} \widehat{w}_k}{\sqrt{\widehat{w}_k^{\intercal} \,\hat{\Sigma}_k^{IS}\, \widehat{w}_k}},$$

in which $\hat{\mu}_k^{IS}$ and $\hat{\Sigma}_k^{IS}$ are the in-sample estimates of the mean and variance, respectively.

Secondly, the certainty-equivalent (CEQ) return is calculated, which is the return that an agent would accept rather than undertaking an uncertain strategy such as an investment in a risky portfolio. Accordingly, for each strategy k the CEQ return of is given by:

$$\widehat{CEQ}_k = \hat{\mu}_k - \frac{\gamma}{2}\,\hat{\sigma}_k^2$$

where $\hat{\mu}_k$ and $\hat{\sigma}_k^2$ are the out-of-sample estimates of the mean and variance, respectively, and γ is the investor's risk aversion. For the sake of simplicity, only the case where $\gamma = 1$ is reported.

Finally, in order to gauge the amount of trading necessary in the execution of each portfolio rule, the portfolio turnover is calculated. Interpreted as the average share of wealth traded in each period, the turnover of strategy k is calculated as:

$$Turnover = \frac{1}{T-M} \sum_{t=1}^{T-M} \sum_{j=1}^{N} \left(\left| \widehat{w}_{k,j,t+1}^{*} - \widehat{w}_{k,j,t+1} \right| \right)$$

where $\widehat{w}_{k,j,t+1}^{*}$ is the desired portfolio weight at time $t+1$ after rebalancing, while $\widehat{w}_{k,j,t+1}$ is the weight before rebalancing.

Results

For each strategy and dataset, the Sharpe ratios, the certainty equivalent return, and the turnover are computed and presented in Table 2. While all three measures are presented, the analysis will stick mainly to the Sharpe ratio since the other metrics provide similar insights.

Table 2: Results from empirical datasets[2]

Metric	Strategy	Sectors	Latam	Factors	FF
Sharpe ratio	1/N	0,024	0,020	0,007	0,044
	mv (in sample)	0,179	0,248	0,270	0,139
	mv	-0,018	0,081	0,109	0,037
		(0,47)	(0,09)	(0,17)	(0,47)
	minvar	-0,028	0,080	0,264	0,072
		(0,29)	(0,15)	(0)	(0,29)
	risk parity	-0,096	0,024	0,058	0,048
		(0,12)	(0,27)	(0)	(0,12)
CEQ	1/N	0,000	0,012	0,000	0,000
	mv (in sample)	0,001	0,017	0,000	0,002
	mv	-0,018	-0,042	0,001	0,000
	minvar	0,000	0,013	0,000	0,000
	risk parity	0,000	0,012	0,000	0,000
Turnover	1/N	0,003	0,003	0,002	0,003
	Relative turnover of each strategy (versus 1/N)				
	mv (in sample)	-	-	-	-
	mv	27,44	5,97	3,38	38,67
	minvar	2,18	1,64	1,28	2,40
	risk parity	0,75	0,95	0,93	0,98

There are two differences of special interest in the analysis of Sharpe ratios. One of them is the difference between the in-sample mean-variance and the $1/N$ strategies, which provides a proxy for the loss from naive rather than optimal diversification, in an environment with no estimation error. This difference is substantial for the datasets

[2] In parentheses is the p-value of the difference between the Sharpe ratio of a given strategy versus that of the $1/N$ benchmark, which is computed using the Jobson and Korkie (1981) methodology previously described.

considered: all datasets have a much smaller Sharpe ratio for the $1/N$ strategy when compared to the mean-variance optimization with the whole sample.

Secondly, in order to compare these strategies from the perspective of a real investor, the out-of-sample estimates have to be assessed. Thus, the severity of the estimation error in such scenarios can be inferred from the difference between the mean-variance's in-sample and out-of-sample Sharpe ratios. This analysis also provides compelling results: all datasets have significantly lower Sharpe ratios for the out-of-sample mean-variance strategy versus its in-sample counterpart. It is also noteworthy that for two of the datasets analyzed, the $1/N$ performs better in terms of Sharpe ratio than the out-of-sample mean-variance optimization, which means that, in such cases, the consequences of estimation error offset the gains from optimal diversification.

With regards to the other strategies addressed in this short study, the minimum-variance and the simplified risk parity strategies, mixed results are found. In all but one dataset (the Sectors one), both the minimum-variance and the risk parity rule outperform the $1/N$ strategy. However, when compared to the out-of-sample mean-variance optimization, no conclusion can be made (mean-variance beats the others in two of the datasets, while being beaten in the remaining ones).

Finally, it is worth mentioning that the turnover for the simplified risk parity strategy is lower than all of the other strategies, even the $\frac{1}{N}$. This may be caused by the fact that the volatility of the assets in a given dataset are not very spread out and that these volatilities are computed on a rolling basis, therefore changing very little from one period to another. In the case of the equally weighted portfolio, even though $\widehat{w}_{k,j,t+1} = \widehat{w}_{k,j,t}$ for all $t \in [1, T]$, $\widehat{w}^*_{k,j,t+1}$ is different. Then, returns may be able to create significant deviations from the desired equally weighted portfolio, resulting in increased turnover.

Conclusion

The aim of this short study is to assess the performance of the $1/N$ strategy (also called naive diversification and equally weighted portfolio) versus a few widely used portfolio rules. In order to gauge the estimation error and the loss from adopting a theoretically simpler strategy, the Sharpe ratios, CEQ returns and turnovers of each strategy are compared across four different datasets related to the Brazilian equity market.

From the discussion presented, the results of the analysis of empirical data suggest that there are no clear winners. For two of the datasets considered, the $1/N$ rule beats the out-of-sample mean-variance optimization, suggesting that, for some cases, the practitioner is better off not estimating anything at all, but simply weighing each asset equally. However, the strategies which focus solely on risk, the minimum-variance portfolio and the simplified risk parity rule, seem to improve on the performance versus both the $1/N$ and the mean-variance rule. As a side note, by comparing the Sharpe ratios of the in-sample and out-of-sample mean-variance optimization, the expected estimation error proves to be very significant. In this short study, however, no strategies that directly deal with this problem are addressed, due to the brevity of the financial time series for Brazilian equities.

Therefore, while such results are a mixed bag, it seems right to be skeptical of the estimations of moments of asset returns as they are usually performed. While naive diversification may seem like an ordinary rule of thumb, the results in this study suggest it to be a powerful tool, even though it remains hard to back its case for optimality.

[i] **Vinicius Esposito; GV Invest, São Paulo School of Economics (Fundação Getulio Vargas).**

References

DeMiguel, V., Garlappi, L., Uppal, R. "Optimal Versus Naive Diversification: How Inefficient Is The 1/N Portfolio Strategy?". Review of Financial Studies. 22(5):1915-1953 (2009).

Jobson, J. D., Korkie, R. "Performance Hypothesis Testing with the Sharpe and Treynor Measures". Journal of Finance. 36:889-908 (1981).

Markowitz, H. M. "Portfolio Selection". Journal of Finance. 7:77-91 (1952).

Memmel, C. "Performance Hypothesis Testing with the Sharpe Ratio". Finance Letters. 1:21-23 (2003).