1QBit

Long-Short Minimum Risk Parity Optimization Using a Quantum or Digital Annealer

White Paper

Long-Short Minimum Risk Parity Optimization Using a Quantum or Digital Annealer

Gili Rosenberg and Maxwell Rounds

Abstract

In portfolio optimization, many weight allocation strategies result in long-only positions. We show how it is possible to formulate and solve an optimization problem that assigns a direction (long or short) to each weight allocation, such that the variance is minimized or maximized. This optimization problem can then be solved by many solvers, including D-Wave Systems' quantum annealer, the Fujitsu Digital Annealer, and others. We present backtested results for three datasets solved using a tabu solver run on a CPU. Our results suggest that by utilizing intelligent shorting, this method is able to reduce the volatility of long-only strategies, leading to shorter maximum drawdowns and higher Sharpe ratios, albeit with a higher turnover. We also discuss possible extensions of this model such that it attempts to achieve market neutrality, sector neutrality, or takes into account a shorting aversion.
Keywords: risk parity, portfolio optimization, quantum annealing, digital annealing

1 Introduction

Risk parity is an approach to portfolio optimization that allocates funds such that each asset is weighted inversely proportional to its risk. Often, this is used in long-only portfolios. If we move to a style of investing in which one can sell short financial instruments (such as stocks or futures), we can consider a weight as representing the amount of money that is set aside as collateral for a given position, rather than money that is exchanged for an asset. If we utilize a measure of risk that is not related to the direction of the position held (long/short), then risk parity does not fully determine the portfolio composition, since a short position carries the same risk as a long position. This means that for a portfolio of n assets, there are 2^n potential portfolios all of which have equal risk assigned to the constituents, thereby opening the door to a strategy that would choose a single portfolio from among those potential portfolios.

2 Formulation

If we define a long-only weight vector w, that is, $\sum_i w_i = 1$ and $w_i \geq 0 \; \forall i$, then each position in a long/short version of the associated portfolio is either long $+w_i$ or short $-w_i$. This allows us to define an optimization problem that selects the direction of the positions such that they minimize or maximize an objective function, such as the variance of the returns or value at risk (VaR). Variance makes for a tangible problem that is quadratic in nature.

To define the problem mathematically, let us define a sign vector s, such that if $s_i = +1$ then position i is long, and if $s_i = -1$ then position i short. We can then write the signed weights as $s^T w$. To define the optimization problem of minimizing (and similarly for maximizing) the variance of the portfolio, we define a diagonal matrix W which has w on the diagonal, and then solve

$$\text{argmin}_s \left\{ s^T W^T \Sigma W s \right\}, \tag{1}$$

where Σ is the covariance matrix. D-Wave Systems' quantum annealer and the Fujitsu Digital Annealer both solve an Ising problem, defined as

$$\text{argmin}_s \left\{ s^T J s + h^T s \right\}, \tag{2}$$

where J is a matrix of couplers, h is a vector of local biases, and $s \in \{-1, +1\}^n$. We can solve the above optimization problem (1) using a quantum or digital annealer by identifying J as $W^T \Sigma W$ but with the diagonal set to zero (qubits/variables are not coupled to themselves), which is equivalent to dropping the constant $\sum_i w_i^2$, which does not change the result of the optimization problem, and setting $h = 0$.

Note that this optimization problem has the symmetry $s \to -s$. This means that every solution is guaranteed to be doubly degenerate, with those two solutions differing only by a global minus sign. In this research, we chose which of the degenerate solutions to use based on which of them gives a weights vector that is closer (via an L_1 norm) to the original weights. The intuition behind doing so is that we expect it would tend to give us lower turnover while giving the same volatility, resulting in better performance. We note that problems with $s \to -s$ symmetry tend to be harder than problems without this symmetry.

To gain some intuition on this optimization problem, let us consider some simple theoretical edge cases. For the case of two assets, problem (1) reduces to

$$\text{argmin}_{s_1, s_2} \left\{ 2 w_1 w_2 \Sigma_{12} s_1 s_2, \right\}, \tag{3}$$

where we use the fact that $s_i^2 = 1$ and ignore the constant $w_1^2 + w_2^2$. Let us consider the case of two identical assets, that is, two assets that have the same variance and perfect correlation. If the assets are identical, our positions should be identical, so $w_1 = w_2 = 1/2$, and to minimize the objective function, we must set $s_1 = -s_2$, meaning that we buy one of the assets and sell the other short. The intuition behind doing so is that for assets that are highly correlated, one can minimize the volatility by "cancelling" the volatility of the long asset against the volatility of the short asset. In contrast, if the objective is to maximize the volatility, the result is $s_1 = s_2 = \pm 1$, meaning that we buy both assets or sell both assets short. For two assets that have the same volatility, but are perfectly anti-correlated, the results are reversed.

If the asset universe consists of assets that are all positively correlated, then we can generalize the first example to understand what will happen. In the maximization version we should set all the signs of the weights to be equal, so as to maximize the quadratic product of the spins and the positive coefficients (from the positive covariance). A physicist would call such a problem a "ferromagnetic" problem, the solution to which is trivial, and yields weights that are identical to the weights of the underlying weight allocation method. In contrast, in the minimization version we must carefully decide which of the signs of the weights should be positive and which should be negative, a non-trivial problem. If, on the other hand, the asset universe consists of assets that are all negatively correlated, the situation is reversed—the minimization version gives the trivial solution, and the maximization version is a non-trivial problem.

3 Results

To compare the performance of each strategy, we performed a backtest on three sets of assets: CTA, DJIA, and SPDR. The first is a typical portfolio for a commodity trading advisor (CTA), consisting of 38 diverse futures contracts,

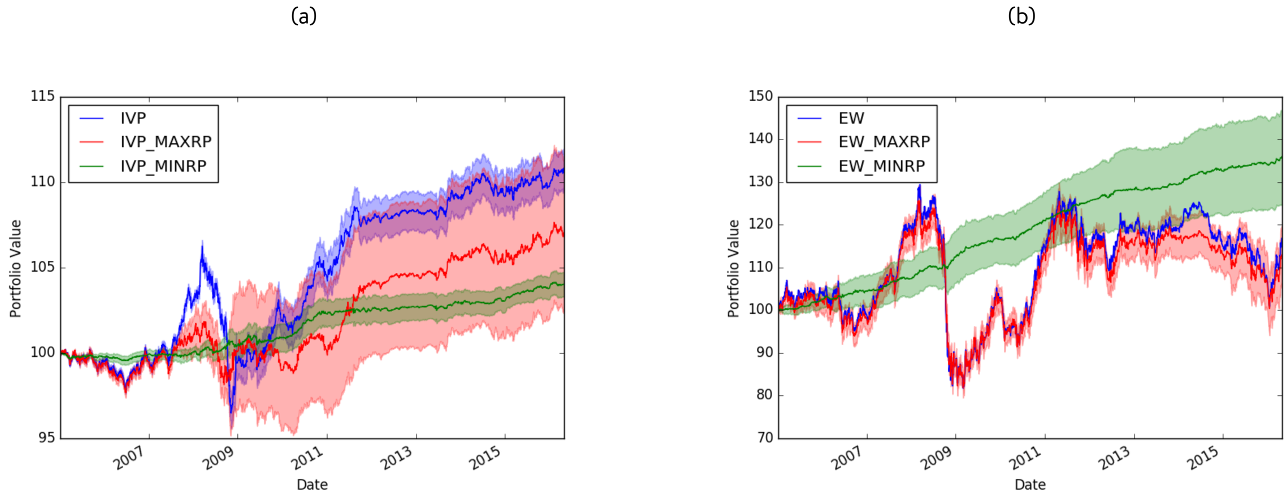

Figure 1: Results for the CTA index: portfolio values for IVP (inverse variance parity) **(a)** and EW (equal weighting) **(b)**. The minimization scheme is labelled by "MINRP" and the maximization scheme is labelled by "MAXRP".

including stocks and bonds from different countries, as well as commodities such as oil, wheat, and gold. The DJIA is the Dow Jones Industrial Average, consisting of 30 large-cap US stocks. SPDR consists of the nine S&P500 sector ETFs.

For the purposes of calculating the covariance, we used a rolling window of three months, and the portfolios were re-balanced on the first trading day of the month. In order to collect statistical data, we used bootstrapping with 25 samples, which also included the historical prices as one of the samples. To perform the bootstrapping, we numbered each date with an index, and sampled with replacement from the set of indices to obtain a sequence of indices. We then used the prices for all assets at each respective index as the price sample. The optimization problems were solved using a multi-start tabu 1-opt search with 100 starts, as a stand-in for a quantum or digital annealer.

For each weight allocation method, it is possible to apply the minimization or maximization scheme described by (1). These are labelled with the suffix "MINRP" and "MAXRP" (respectively). The results for each index for equal weighting (EW) and inverse variance parity (IVP) are presented in Figs. 1–3. In EW, a weight $1/n$ is assigned to all n assets, and in IVP, the weights are proportional to the inverse of the variance of each asset over the rolling window.

Results for all the algorithms tested are shown in Table 1, based on a backtest on the DJIA. The table includes results for the methods described above (EW and IVP) as well as: "MV"—minimum variance; "HRP"—hierarchical risk parity; and "QHRP"—quantum hierarchical risk parity. The metrics are described in the table's caption.

4 Discussion

From the portfolio value results in Figs. 1–3, it is clear that MINRP has a remarkable ability to reduce the volatility of IVP and EW (at least for these methods and datasets), albeit at the cost of lower returns in some cases. In particular, the returns for IVP_MINRP are unimpressive for the CTA and SPDR datasets. The impressive performance of EW_MINRP on the CTA and DJIA datasets is surprising; additional investigation would be required to understand why it works so well, and why it does not work as well for SPDR.

As noted above, when all correlations are positive, MAXRP produces long-only results that are equal to the underlying method. Indeed, one can see this for the DJIA and SPDR. In contrast, CTA includes some anti-correlated assets, and there MAXRP indeed diverges from the results of the underlying methods. In the case of IVP_MAXRP, it unfortunately results in a very large standard deviation of the portfolio values. Table 1 confirms that for a wider range of methods, MINRP is able to reduce the volatility and maximum drawdown substantially, and achieve the highest Sharpe ratio seen in this experiment. We remark that although the turnover when using MINRP tends to be large, it is comparable to the turnover of MV when used alone.

(a) (b)

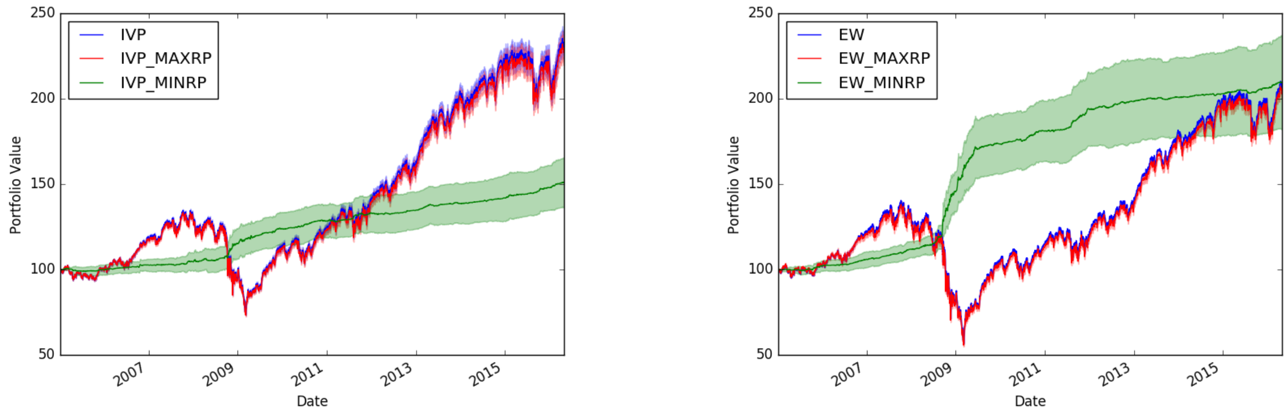

Figure 2: Results for the DJIA index: portfolio values for IVP **(a)** and EW **(b)**. Labels are as in Fig. 1.

(a) (b)

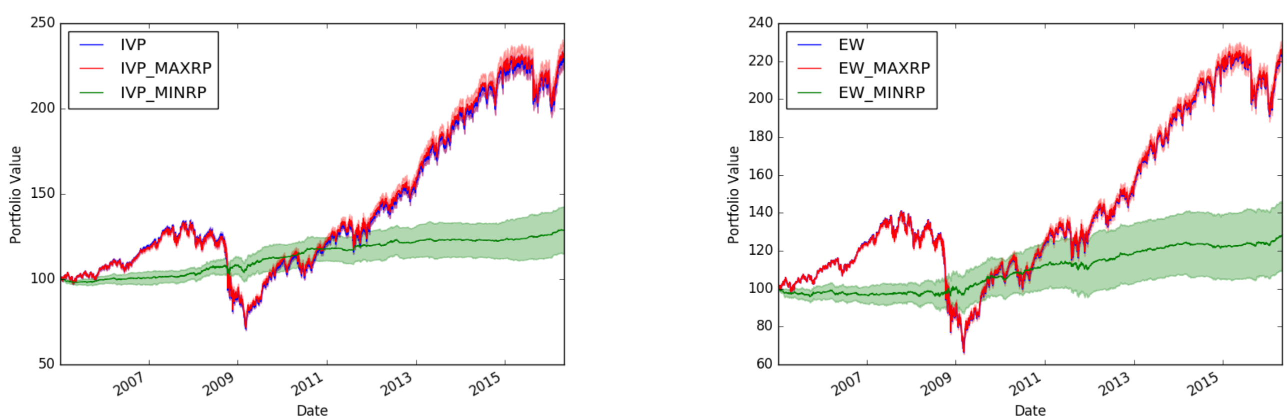

Figure 3: Results for the SPDR index: portfolio values for IVP **(a)** and EW **(b)**. Labels are as in Fig. 1.

Method	Volatility [%] Mean	STD	Mean returns [%] Mean	STD	Sharpe ratio Mean	STD	Max DD [%] Mean	STD	Turnover Mean	STD
IVP	16.80	0.05	7.77	0.28	0.46	0.02	45.23	0.46	0.14	0.01
IVP_MAXRP	16.77	0.05	7.64	0.30	0.46	0.02	45.24	0.45	0.14	0.01
IVP_MINRP	3.18	0.10	3.79	0.93	1.19	0.29	5.22	1.21	0.54	0.01
EW	20.78	0.00	6.75	0.00	0.32	0.00	59.75	0.00	0.00	0.00
EW_MAXRP	20.72	0.03	6.60	0.13	0.32	0.01	59.70	0.16	0.01	0.00
EW_MINRP	4.42	0.18	6.78	1.19	1.53	0.27	5.11	1.34	0.48	0.01
MV	15.52	0.28	8.26	1.27	0.53	0.08	34.80	3.88	0.56	0.03
MV_MAXRP	15.03	0.31	7.47	1.80	0.50	0.12	34.58	3.75	0.59	0.04
MV_MINRP	9.19	0.49	1.99	1.82	0.22	0.20	25.26	4.54	0.74	0.03
HRP	17.04	0.27	6.84	0.87	0.40	0.06	51.09	2.01	0.36	0.01
HRP_MAXRP	16.96	0.27	6.39	0.79	0.38	0.05	50.94	2.04	0.37	0.02
HRP_MINRP	4.52	0.17	4.12	1.36	0.91	0.30	7.68	2.19	0.65	0.01
QHRP	15.24	0.28	6.56	1.44	0.43	0.10	42.76	5.42	0.59	0.03
QHRP_MAXRP	15.11	0.28	6.03	1.67	0.40	0.11	42.60	5.54	0.61	0.03
QHRP_MINRP	7.51	0.32	2.87	2.09	0.38	0.28	18.49	5.70	0.76	0.03

Table 1: Metrics from the backtesting of various portfolio allocation methods on the DJIA with a three-month window and monthly rebalancing. The means and standard deviations (STD) were taken over 25 bootstrapped samples, as explained in the main text. The metrics are: "Volatility"—the standard deviation of the returns; "Mean returns"; "Sharpe ratio" (annualized); "Max DD"—the maximum drawdown; and "Turnover"—the fraction of the portfolio that shifted at each rebalancing step. The names of the methods are defined in the text. The backtest period was chosen arbitrarily as being from January, 2005 to May, 2016.

5 Possible Extensions

There are multiple ways in which the problem in (1) could be extended. For example, consider market neutrality, in which an investor seeks to maintain a neutral exposure to a particular factor, such as maintaining a roughly equal investment in long and short positions. One could incorporate a market neutrality bias into the optimization problem, such that solutions are preferred if they are closer to market neutrality. This could be achieved by adding a penalty term $\alpha \left(\sum_i w_i s_i\right)^2$, where $\alpha > 0$ controls the relative strength of the market neutrality term versus the variance term. The new term is zero for a market-neutral allocation, and positive otherwise. One could also impose sector neutrality, which would involve adding a similar term as above for each sector, with the sum running over only the assets in that sector.

Another possibility is to take into account a shorting aversion. Having noticed that the total weight allocated to short positions can be written as $\frac{1}{2}\sum_i w_i(1-s_i)$, we can add a proportional term in order to penalize large total shorting fractions. Note that the turnover, defined as the fraction of the portfolio that is shifted, can be written as the very same term, yielding the result that imposing a bias against high turnover is the same as imposing a bias against short positions. The reason is that long positions do not incur any turnover (under the assumption that the underlying weights are all positive). Also note that this term is linear in s, such that adding it to the model would break the $s \to -s$ symmetry of the original model (and give a non-zero h).

6 Conclusions

We have presented a method for lowering the volatility of long-only portfolios by converting some positions to short sales. The results suggest that this might be an effective way of reducing the volatility and maximum drawdown for a range of portfolio allocation methods, albeit at the cost of a higher turnover. Note also that short positions are not permitted in some assets. Further study would be necessary to determine how transaction costs, slippage, and leverage affect these results. We have formulated the problem for solving on a quantum annealer, and expect to take advantage of such machines as they become larger and more powerful.

Acknowledgements

The authors would like to thank Andrew Milne and Clemens Adolphs for their useful comments, Clemens Adolphs for the use of his portfolio optimization benchmarking code, and Marko Bucyk for editing the manuscript.