Technical Note

INSIGHT.DATA.CLARITY.

May 2017

MAKING OPTIMISATION TECHNIQUES ROBUST WITH AGNOSTIC RISK PARITY

Introduction

The alternative investment industry is becoming ever more accessible to those wishing to diversify away from traditional portfolios. With interest rates at record lows, bonds and equity indices look less enticing investment opportunities than they once did and alternatives in friendly formats are looking like potentially meaningful diversifiers. Alternative benchmark strategies are often simple to understand and in some cases to paper trade but in practice require the necessary implementation skill set to fulfil their potential.

In this short note we discuss portfolio construction, an area of research that has been a point of focus for many years in trading equities at CFM and which has more recently become an active part of the research program in directional strategies. The fruits of this research program are presented in this note in a procedure based on the techniques of Markowitz's Mean Variance Optimisation (MVO) and extended to the idea of Agnostic Risk Parity[1] (ARP). We use a Trend Following strategy to illustrate the effectiveness of the approach in building a robust portfolio of correlated instruments.

Contact details

Call us +33 1 49 49 59 49

Email us cfm@cfm.fr

The importance of diversification

It was Harry Markowitz who once stated that diversification is the only free lunch in finance. Our intuition, that we should not put all our eggs in the same basket, is indeed backed up by the mathematics of how risk and returns add in combining decorrelated investments.

Let's consider two arbitrary strategies and look at how the returns combine compared to the volatilities. An investment in an instrument that returns say $100 and another that returns say $150 will result in a net profit of $250, meaning that returns simply add, no matter if the investments are correlated or not. Volatilities on the other hand add quadratically in the following way:

$$\sigma_{total}^2 = \sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2$$

where σ_1 and σ_2 are the volatilities of the two investments while the last term accounts for the potential correlation between the investments, ρ. Let's say for example that σ_1=$100 and σ_2=$150 such that the Sharpe ratios of each strategy are equal to 1. If the two strategies are decorrelated with ρ=0, adding them together gives us a Sharpe ratio of $250/\sqrt{32500}=1.4$, higher than either of the initial Sharpe ratios. This is the key to diversification! If strategies are decorrelated then the return of the combination increases more quickly than the risk and we end up with better risk adjusted returns.

Diversification in a pool of (often highly correlated) instruments

Let's now begin with the simplest portfolio possible, two instruments, say the S&P 500 and the Nasdaq 100 future contracts, chosen due to the fact that they are approximately 90% correlated. With such a high level of correlation the diversification benefits of trading both are limited. There is, however, a way to decorrelate from a position which is long both contracts and that is to trade a long and short pair. If we assume that the volatility of the two contracts is the same and the contracts have the same notional size then a position of +1 lot of the S&P 500 and -1 lot of the Nasdaq 100 gives a strategy that is decorrelated from a position of +1 lot of the S&P 500 and +1 lot of the Nasdaq 100. This decorrelation property of the pair exists whether the two contracts are correlated or not (unless, of course, the correlation is 100%!). In Figure 1 we show the Profit & Loss (P&L) curves for these long/long and long/short configurations along with a correlation between the two estimated using moving averages.

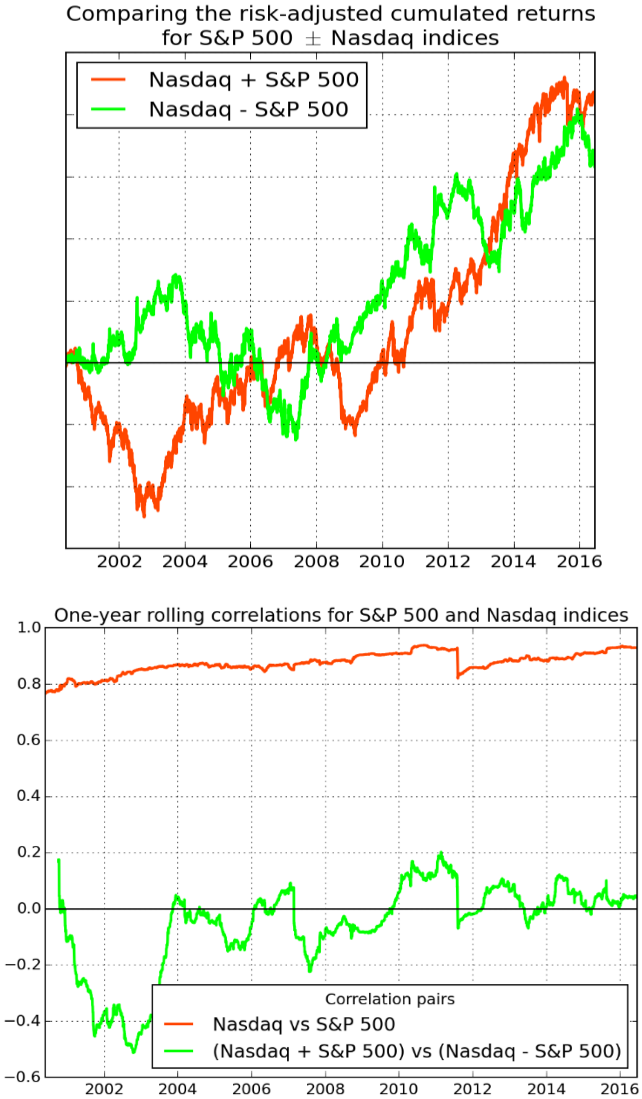

Figure 1: The upper plot shows the P&L curves of an equal long/long position in the Nasdaq 100 and S&P 500 futures and the decorrelated equal long/short positions. On the lower plot we show the correlation between these two futures and the correlation of the long/long combination with the long/short combination. This correlation is seen to consistently be zero.

A MVO will try to maximise the gain of a risk constrained portfolio, or stated more simply maximise the Sharpe ratio, by allocating optimally to these diversifying configurations of instruments. However, the caveat for a MVO to improve

[1] Please refer to the Appendix for a discussion of the ideas of Mean Variance Optimisation and Agnostic Risk Parity. This short note describes techniques and ideas detailed in our academic paper "Agnostic Risk Parity: Taming Known and Unknown Unknowns"

performance is that one needs a perfect knowledge of future performance and correlation between instruments, in which case a unique and optimal solution exists. For the given example the optimal solution allocates favourably to a portfolio configuration that is a long/short combination of the S&P 500 and Nasdaq 100 futures. Given that the correlations between these two instruments are so high, this configuration is one that heavily reduces the risk of the portfolio and ultimately the Sharpe ratio and the returns (after leveraging the risk back up again!). A MVO will always improve a non-causal simulation[2] due to its ability to allocate to portfolios that have worked in-sample, in the past, with the benefit of hindsight[3]. In this example the optimisation will allocate precisely the right amount to long/long and to the long/short positions such that the Sharpe ratio is maximised to work over the period of the back-test. Unfortunately, these returns and correlations in the past are not good indicators of the future! It is now natural to take this two instrument example and scale up to a portfolio. We can achieve diversification by allocating to combinations of instruments which are similarly decorrelated from a portfolio which is long all instruments, and in so doing will improve the in-sample Sharpe ratio for the period of the back-test having perfect knowledge of the returns and correlations between the instruments.

The key to achieving real forward looking performance improvements with MVO techniques is to get better and more robust estimates of future returns and correlations. CFM has studied this subject for many years and pioneered such research in the field of correlation matrix cleaning[4] which is an essential component to being able to allocate correctly across robustly determined Risk Factors[5].

A commonly employed heuristic for trend followers is to be equally allocated among the four sectors – equity indices, interest rates, commodities and FX. This is a diversified portfolio with the natural choice of sectors providing the decorrelation. The research we have conducted in adapting the techniques of the MVO has always used this equally weighted portfolio as the benchmark we try to beat. If we now generalise to a full portfolio of contracts from a typical CTA universe then we would like our portfolio construction algorithm to equally allocate to all sources of decorrelation in the universe, not only through exposure to decorrelated contracts but also to decorrelated portfolios or risk factors. The ideas of ARP are consistent with this, instead of allocating to contracts, we allocate risk equally (on average and over time) to the principal components of the universe, thus building a fully diversified portfolio accounting for correlations between contracts.

Applying ARP to a CTA universe of instruments

Returning to our two correlated contract portfolio we can look at the case of two trend forecasts and examine the effects of applying a standard MVO and an ARP algorithm to the portfolio. The correlation between the two contracts is measured at 90% while the volatility of each is assumed to be the same at 15% annualised. For the purposes of illustration the result of the trend forecast applied to the S&P 500 is a signal of 1 while the Nasdaq 100 has a signal of 0.5, where these forecasts are just the output of a trend algorithm applied to the price time-series of the S&P 500 and the Nasdaq 100 respectively. One can think of these forecasts as meaning, over the time horizon defined by the model, the S&P 500 will rise by twice as much as the Nasdaq 100, at least if the signal is to be trusted! If the two instruments were decorrelated then it would seem intuitively reasonable that one should allocate according to the size of the expected future return, if each has the same volatility then this is equivalent to allocating proportionally to Sharpe ratio (which is indeed the optimal solution from a MVO). In the presence of correlations, however, the MVO tells us to do something quite different! We nonetheless will retain the "allocation=forecast", equal risk weighted strategy as our benchmark.

Table 1 below illustrates the different portfolio configurations resulting from the application of an equal weighting, where the allocation is only proportional to the forecast, a standard MVO and our new ARP approach.

As one can see from the Table, the MVO procedure produces the best risk adjusted returns. However, this assumes that the correlations between the two instruments and their returns are precisely known. Unfortunately, out-of-sample, future returns and correlations are not precisely known and therein lies the problem. In fact one can see that the MVO allocates more risk to the long/short combination and is therefore less diversified than the ARP allocation, which allocates more equally across sources of risk. It is often the case that a MVO ends up with polarised or concentrated positions that end up being less diversifying rather than more.

[2] Non-causal in the sense that we use future information. We will improve any combination of strategies with a perfect knowledge of future returns and future correlations. If we build today's portfolio with returns and correlations from yesterday then we do not see any improvement in performance with these optimisation techniques

[3] See our white paper In-sample Overfitting – Avoiding the Pitfalls in Datamining

[4] See Cleaning Correlation Matrices – M Potters, J Bun, JPh Bouchaud published in Risk Magazine (April 2016) where the Rotationally Invariant Estimator (RIE) is introduced

[5] A Risk Factor (or Principal Component) generally refers to a portfolio of instruments constructed with the available universe that comes from a Principal Components Analysis – a dimensionality reduction technique that aims to explain the correlation universe as being composed of the biggest risk explaining portfolios

Measures of diversification, unfortunately, depend on the "basis"[6] in which they are defined. Consider, for example, a stock investor holding a portfolio of stocks. His most diversified portfolio will be a long only, equally risk allocated basket, meaning he defines his portfolio to be equally weighted in his coordinate system i.e. that of the stock market. But, we know that the biggest risk factor in the stock market is the market itself and this equally allocated basket is maximally correlated with the market. This highly diversified portfolio, from the point of the view of the stock holder, therefore offers very little diversifying power compared to what is possible by building portfolios exposed to other diversifying risk factors. The only coordinate system where diversification measures make sense are those in which the corresponding synthetic assets are uncorrelated such that diversification can be rationally and objectively evaluated. ARP combines this idea with the hypothesis that past correlations between predictors are fragile and should not be relied upon to hedge different bets. ARP also addresses the problem of regime shifts in the correlations. For example, although future "true" correlations are close to the best estimate of cleaned correlations, these correlations can (and do) change drastically to a new regime not observed in the past. The ARP approach serves as a layer of protection or safeguard against these statistically unexpected events.

Procedure	Positions	Gain (return)[7]	Portfolio Risk(%)[8]
Equal risk weighting	$\pi_{S\&P500}$=1 $\pi_{Nasdaq100}$=0.5	1.25	22
Mean Variance Optimisation	$\pi_{S\&P500}$=1.96 $\pi_{Nasdaq100}$=-1.42	1.25	14
Agnostic Risk Parity	$\pi_{S\&P500}$=1.38 $\pi_{Nasdaq100}$=-0.25	1.25	17

Table 1: The weights, gain (or return) and risk for the three different portfolios considered. In each case the overall return of the portfolio is kept constant and the volatility calculated for each. In each portfolio we are not adding forecasting power, merely constructing the portfolio differently such that the risk is reduced. The equally weighted portfolio is allocated according to the forecast and forms the benchmark. The MVO allocates a heavy negative weight to the Nasdaq 100 in order to reduce the overall volatility of the portfolio and increase the Sharpe ratio while still maintaining the same level of positive return. This is optimal if we assume the forecast and correlations are perfect estimates of the future, which in reality is optimistic! Furthermore, the positions taken by the MVO show that it is more preferentially allocated to the long/short risk factor. The ARP portfolio, on the other hand, is better diversified in that one allocates equally across the long/long and long/short portfolios. It is therefore less extreme in its allocation to the lowest risk long/short configuration and is more fully diversified across all sources of diversification.

We can now take these ideas, as applied to the simplified 2 instrument universe above, and extend to the multi-instrument universe that a typical CTA trades. The portfolio used is composed of 110 instruments made up of futures on equity indices, bonds, short term interest rates, FX and commodities. Our benchmark allocation (that we are trying to beat!) is an equal risk weighting across all assets[9] and our forecast is a 1-year average of returns, or in other words a plain vanilla long term trend forecast. This benchmark allocation is considered the standard way to allocate in the CTA industry and indeed gives us the most correlation with CTA indices made up of the biggest managers in the space. We now try comparing the benchmark portfolio to the following:

- A standard MVO approach back-tested in a causal fashion with an optimisation which is run every day, d, and applied to build the portfolio on day, d+1, in such a way as to ensure no future information is included in the simulation and robustness to in-sample biases is tested. The correlation matrix is un-cleaned and based on empirical measurement
- A standard MVO approach back-tested in a causal fashion with a cleaned correlation matrix based on the RIE technology
- An ARP approach back-tested in a causal fashion with a cleaned correlation matrix based on the RIE technology

As discussed above, one can consider each case to be different examples of allocating risk across the principal components of the correlation matrix of the universe considered. An equal weighting in risk across individual assets corresponds to a linearly increasing allocation across principal components while a MVO favours allocation to the smallest risk principal components. The ARP portfolio, on the other hand, allocates equally across sources of risk in the universe, in other words, across the principal components. This is illustrated pictorially in Figure 2.

The back-tests using the long term trend forecast are presented in Figure 3, comparing the benchmark, equally weighted portfolio to each of the different allocation styles. The performance curves are normalised to have the same risk thus demonstrating an improvement in risk adjusted returns (or Sharpe ratio) using the ARP approach. This improvement changes as a function of the window of

[6] A basis refers to a set of portfolios defined to be orthogonal (decorrelated) to each other. The "eigenbasis" is the orthogonal portfolio set that comes out of a Principal Components Analysis with each portfolio being the eigenvector with a corresponding eigenvalue that is simply a measure of the volatility carried by the eigenvector

[7] Gain is defined as the sum of the products of position and prediction $\sum_i \pi_i p_i$

[8] Portfolio risk is defined as $\sqrt{(\pi_{S\&P}\sigma_{S\&P})^2 + (\pi_{Nasdaq}\sigma_{Nasdaq})^2 + 2\rho\pi_{S\&P}\sigma_{S\&P}\pi_{Nasdaq}\sigma_{Nasdaq}}$, where ρ is the correlation between the S&P 500 and the Nasdaq 100

[9] With an approximately equal number of contracts in each sector this is equivalent to an equal weighting on each of the four sectors – equity indices, interest rates, FX and commodities

time used in the back-tests but remains stubbornly and robustly present across many periods.

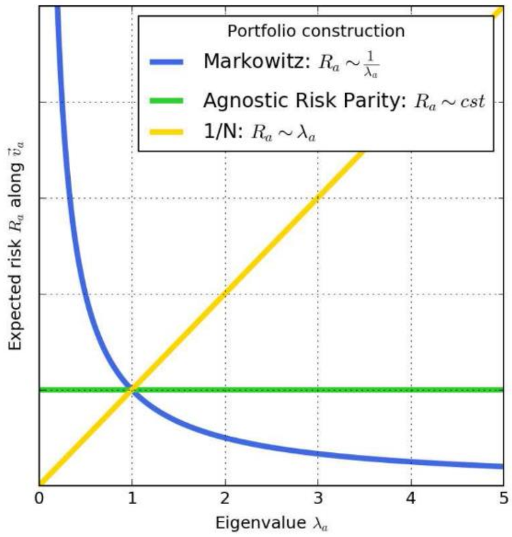

Figure 2: The allocation of risk across the principal components of the universe for the three portfolio construction techniques considered. The x-axis corresponds to the “eigenvalue” of the principal component, which is mathematically equivalent to the variance, or the volatility squared, of each. The plot demonstrates that a MVO allocates risk favourably to the lowest volatility principal components, while the equally weighted, 1/N, allocation allocates favourably to the highest volatility principal components. The ARP approach, however, allocates equally and uniformly to each of the principal components, no matter the volatility of each.

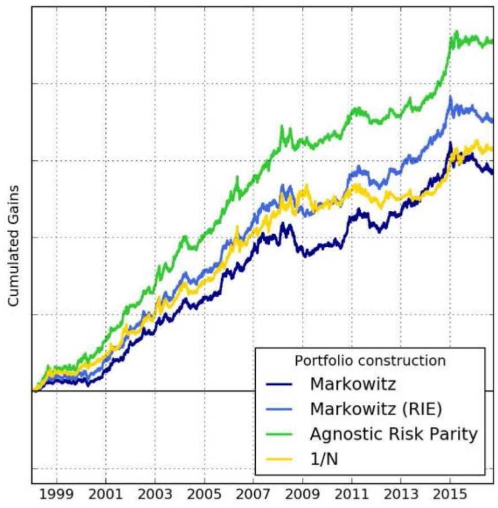

Figure 3: The P&L performance curves of each of the allocations considered using a long term trend following signal across a universe of 110 typical CTA instruments made up of equity indices, bonds, FX and commodities. Each curve is normalised to have the same risk, showing that the ARP approach gives the best risk adjusted returns.

Conclusions

The MVO techniques employed in the financial industry are known to be flawed and often prove themselves to be lacking in robustness. The optimal solution to the problem takes everything at face value and assumes that the past is a perfect indicator of the future. This unique optimal solution often builds up large positions on small bets, counterintuitively reducing the potential diversification in the portfolio. We have presented a portfolio construction framework based on an adapted MVO procedure that finds decorrelated contracts and configurations, and allocates equally across them in order to construct a more robust (to out-of-sample performance) portfolio. The approach has been tested using a long term trend following strategy applied to a universe of 110 futures contracts typically traded by the CTA industry. Our work in cleaning correlation matrices, to get the best out-of-sample estimates of correlation, is an important input to this procedure. Further research in this field will now take us in the direction of applying these techniques to other directional strategies such as Value and Carry applied to futures contracts.

Appendix

The MVO attempts to maximise expected gain with a variance penalty. For a portfolio of N tickers with positions π and returns η the problem involves maximising the following utility function:

$$\max_{\pi_i}\left(\sum_i \pi_i E(\eta_i) - \lambda \sum_{i,j} \pi_i \pi_j E(\eta_i \eta_j)\right)$$

where λ is a risk control parameter: increasing (decreasing) values of λ will reduce (increase) the optimal positions π. We assume that a predictor p_j is an estimator of the expected return $E(\eta)$ while the empirically measured covariance matrix C_{ij} is an estimate of the covariance $E(\eta_i\eta_j)$. The utility function which is actually maximised is then:

$$\max_{\pi_i}\left(\sum_i \pi_i p_i - \lambda \sum_{i,j} \pi_i \pi_j C_{ij}\right)$$

For which the maximum is classically obtained for:

$$\pi_i = \frac{1}{2\lambda}\sum_j C_{ij}^{-1} p_j$$

This solution focuses risk on smaller risk portfolio configurations in order to maximise Sharpe ratios with a precise knowledge of correlations and returns. A more robust construction involves an equal allocation of risk (on average) across the principal components or risk factors of the universe. The positions in this case are found to satisfy:

$$\pi_i = \frac{1}{2\lambda}\sum_j C_{ij}^{-1/2} p_j$$

Disclaimers

CFM has pioneered and applied an academic and scientific approach to financial markets, creating award winning strategies and a market leading investment management firm.

Contact us

Capital Fund Management S.A.
23, rue de l'Université, 75007
Paris, France
T +33 1 49 49 59 49
E cfm@cfm.fr

CFM International Inc.
The Chrysler Building, 405 Lexington Avenue - 55th Fl.,
New York, NY, 10174, U.S.A
T +1 646 957 8018
E cfm@cfm.fr

CFM Asia KK
9F Marunouchi Building, 2-4-1, Marunouchi, Chiyoda-Ku,
100-6309 Tokyo, Japan
T +81 3 5219 6180
E cfm@cfm.fr

Capital Fund Management LLP
64 St James's Street, London
SW1A 1NF, UK
T +44 207 659 9750
E cfm@cfm.fr