CFM
INSIGHT.DATA.CLARITY.

3rd May 2017

ALTERNATIVE BETA MATTERS
Quarterly newsletter - Q2 2017

Introduction

Welcome to CFM's "Alternative Beta Matters" Quarterly Newsletter.

Within this report we recap major developments of the quarter for Equities, Fixed Income / Credit, FX and Commodities, as well as Alternatives. All discussion is agnostic to particular approaches or techniques, and where alternative benchmark strategy results are presented, the exact methodology used is given.

We have also included one white paper and an extended academic abstract from a paper published during the quarter. Our hope is that these publications, which convey our views on topics related to Alternative Beta that have arisen in our many discussions with clients, can be used as a reference for our readers, and can stimulate conversations on these topical issues.

Contact details

Call us +33 1 49 49 59 49
Email us cfm@cfm.fr

Quarterly review

Quantitative overview of key developments in Q1 2017

Equity indices

The beginning of 2017 brought a robust quarter for equities and a reversal of fortune for emerging markets, compared to the end of 2016, as they spectacularly outperformed their developed market counterparts. Donald Trump took office in January on promises of infrastructure spending, deregulation and protectionism. The lack of follow-through on these protectionist trade policies, however, supported risk appetite and gave a lift to emerging market equities. The difference in performance between developed and emerging equity markets was clearly seen in the MSCI World/MSCI EM index spread - the MSCI World equity index, made up of 23 developed country equity indices, returning 5.9% through Q1, while the MSCI EM index, made up of 23 emerging countries, returned 11.1%. The correlation between the two indices continued to drift south, a post Trump election pattern that has emerged, slipping below 80%, a level still below the last peak of 90% seen in 2012. Equity returns in the US continued to be positive in Q1 while in Europe and Japan performance was somewhat flatter. European markets were jittery, responding to the start of political cycles in many countries. Elections in Holland saw the far-right candidate, Geert Wilders, lose out to the centre-right and the build-up to the presidential election in France brought its own economic and political uncertainty. German elections are also on their way in 2017.

The generic trender[1] applied to equity index markets was the only positive performer with the FX, commodity and interest rate sectors each delivering negatively. The best performing contract was the mini S&P traded on the CME as US equity markets continued on their upwards trend from 2016. The RSI[2] applied to this pool of equity indices reached a maximum of 68 at the beginning of March for the mini S&P and a minimum of 49 at the beginning of January for the Chinese A50 future traded on the SGX in Singapore.

Despite the uptick in political uncertainty, equity index implied volatility remained low in Q1. Various policy uncertainty indices spiked in anticipation of future volatility but the VIX, V2X and VNKY all remained low through the quarter. Liquidity conditions remained stable, falling post Q4, following the increase in activity through the US presidential election result before picking up in March. The French presidential election in April, with a second round two weeks later, may yet see a demand for insurance and a pick-up in implied volatility.

The returns of the MSCI World and the MSCI Emerging Markets indices for the past year

Stocks and factors

2016 was not a good year for equity factors with momentum, in particular, having a tough time navigating the choppiness in the markets. This was reflected in our reproduction of the Fama-French-Carhardt factors[3] with good performance for value and poor performance for momentum, in our case the HML and UMD factors respectively. The beginning of 2017 saw a good start for UMD with January being positive in Japan and Europe, a pattern seen more generally across the equity market neutral industry as a whole. This was not a sign of things to come, however, as the factor gave everything back by the close of the quarter. The HML factor delivered negatively in Q1 and value investing, looking more broadly across publicly available offerings, has delivered negatively in 2017 so far. Low volatility factors have tended to fare better through Q1 having been a drag on performance generally in 2016 for equity market neutral players.

[1] The trender used here is defined as the sign (either +1 or -1) of the difference of a 50 day exponentially weighted moving average (EWMA) and a 100 day EWMA

[2] Defined according to https://en.wikipedia.org/wiki/Relative_strength_index using 100 day exponentially weighted moving averages. The RSI varies between 0 and 100 with 70 implying an instrument is overbought and 30 implying the instrument is oversold

[3] We use a CFM version of the Fama French implementation for momentum (UMD), value (HML) and size (SMB) and have tested the convergence with the data from Kenneth French's website

Despite the strong performance of equity indices in the US there were clear winners and losers across industrial sectors. Information technology, followed by consumer discretionary led the pack while energy stocks lagged, in line with expectations following the collapse in Crude Oil late in the quarter. This pattern was also seen in Europe while in Japan cyclical sectors came out on top and Financials lagged. Implied volatility of sectors in the US exhibited unsurprising patterns with a general sell-off in volatility through the quarter and a pick-up in March for Energy and Financials.

The Fama-French factors for the last year in Europe, Japan and US

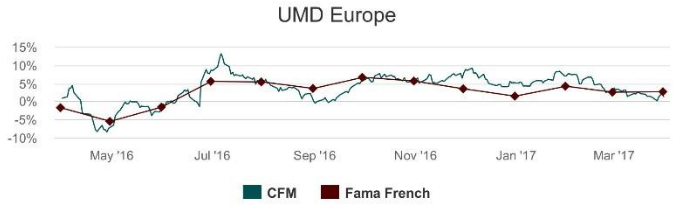

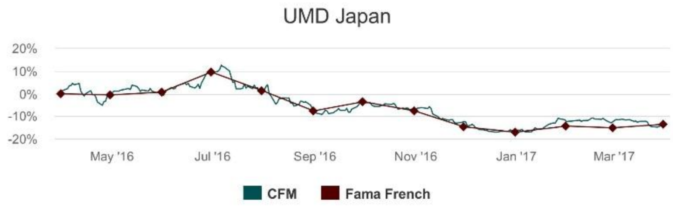

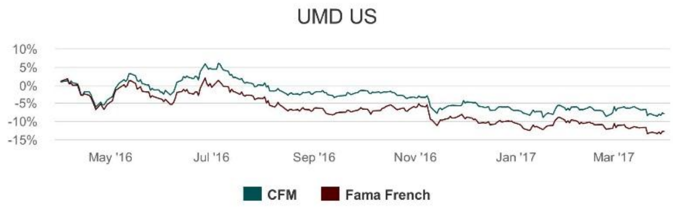

(High Minus Low (HML) corresponds to a market neutral (MN) portfolio long the high book to price stocks and short the low book to price stocks. Small Minus Big (SMB) corresponds to a MN portfolio long the small market cap stocks and short the large market cap stocks. Up Minus Down (UMD) corresponds to a MN portfolio long the historical winners and short the historical losers. In each case, the red line is downloaded from Kennenth French's website, while the blue line is the CFM reproduction of the Fama-French portfolios. The methodology can be attributed to Eugene Fama and Kenneth French and is not explicitly used in any CFM product.)

Fixed income and credit

Q1 saw diverging developed market central bank policy accompanied by an inflation up-tick against a back-drop of increasing commodity prices, at least in the first two months of the quarter. The Federal Reserve remains on an upwards path for interest rates whereas the European Central Bank and the Bank of Japan are staying the course with a continuation of their accommodating loose policies. This translated into a slightly positive quarter for bond holders with the Barclays Hedged Global Aggregate Bond Index returning a meagre 0.4% in Q1. A better growth outlook and expectations of fiscal changes in the US saw a highly anticipated 25bp rate hike from the Federal Reserve in March. Janet Yellen's post FOMC press conference convinced the market that the increase was dovish and ten year treasury prices drifted higher towards the end of the quarter. Political uncertainty in Europe contributed to diverging sovereign debt yield spreads with attention, in the second half of the quarter in particular, paid to France's cost of borrowing relative to Germany. The bond market's perception of a potentially market unfriendly French president saw French bonds fall and a flight to quality supporting the price of German debt. The BoJ meanwhile has stuck rigidly to its steady-as-she-goes message with policy settings unchanged: negative rates on excess reserves, a cap on 10-year yields at about zero per cent and purchasing 80 trillion Yen's worth annually of government bonds.

The generic trender applied to interest rate markets was negative in Q1 having been caught out by a series of reversals. The Bund was the worst performer for long bonds with a series of sharp reversals on the timescale of the trender leading to a series of buy high/sell low trades. All bond markets that were considered, however, delivered negative performance. The short end of the curve was more mixed with offsetting positive and negative performance for the trend applied to a range of Short Term Interest Rate futures. The Australian Bank Bills market was the worst performer while the trender applied to Euribor was the best. The RSI indicator was characteristically close to 50 for a range trading market with a low of 40 for the Canadian 10 year at the beginning of the year and a high of 57 for UK Gilts at the beginning of March. US ten year treasury implied volatility steadily, as measured by the TYVIX[4], reduced throughout Q1 following the jump upwards around the time of the US presidential election result. Liquidity in interest rate markets fell initially, at the start of the year, before surging at the end of February, and seeing out the quarter by providing good trading conditions.

Equity market rallies in developed markets were also matched by improving conditions in credit markets with credit spreads continuing to narrow. Investment Grade credit on both sides of the Atlantic played catch up relative to High Yielders compared to Q4 2016 with High Yielders only outperforming Investment Grade counterparts by 1.4% in Q1, as measured by the Barclays Corporate indices. Energy company bonds lost ground with the collapsing price of oil late in the quarter.

The return of the Barclays Hedged Global Aggregate Bond and the CDX Investment Grade indices for the last year

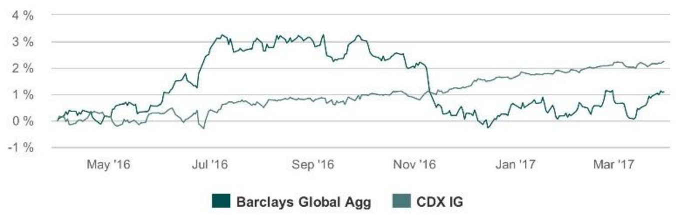

Commodities

Q1 mostly started well for commodities but saw out the quarter with a less convincing performance in March. The GSCI benchmark fell 2.5% through with the March collapse in Crude weighing heavily on the index. The oil price fall came as comments from Saudi Arabia and Russia failed to convince markets they were committed to production cuts while US inventories remained persistently high. The CFTC Commitments of Traders[5] (CoT) data for the quarter showed near record longs for non-commercials in the near WTI Crude future at the end of February and prior to the market sell-off. As we hit the end of Q1 the non-commercial positioning remains stubbornly long. Crude implied volatility as measured by the OIV[6] dipped through January and February before rising in March once more while Crude-S&P 500 correlations reduced through January and February before reversing and accelerating in March. Non- energy commodities gained slightly in Q1 albeit with a large spread between sub-sectors and products. Cocoa continued on its supply driven decline while grain prices mostly gained. Gold and Silver gained strongly in the first two months of the year on dollar weakness before selling off and rising once more in March to see out Q1 in positive territory. Gold implied volatility as measured by the GVZ[7] fell steadily through the quarter. Copper's journey was less dramatic, being generally flat following its Trumpflationary boost in Q4 2016.

The generic trender applied to commodity markets delivered negatively in the quarter with the violent March reversals for Gold and Silver being particularly difficult to navigate making Silver the worst performer. The best performance came from the NYBOT Sugar contract's decline through Q1. The RSI indicator showed no overbought or oversold contracts in Q1. The highest RSI was 61 for Copper at the beginning of January while the

[4] The TYVIX is calculated from the CBOT's options on 10 year futures, using the same methodology as the VIX, and is published by the exchange

[5] http://www.cftc.gov/MarketReports/CommitmentsofTraders/AbouttheCOTReports/index.htm

[6] The CBOE/NYMEX WTI Crude Volatility Index

[7] The CBOE Gold ETF Volatility Index

lowest value was 35 for Cocoa in the middle of February. Risk weighted liquidity continued to fall from the peak provided by the US presidential elections in November before rallying slightly at the beginning of the March.

The one year return of the S&P GSCI

FX

Q1 provided an interesting start to the year with FX markets being heavily influenced by global politics along with the more usual pressures from the most influential central banks. Europe awaited the results of the Dutch elections keenly only to see populist Geert Wilders defeated in favor of centrist politics, the first major political result for the Eurozone of 2017. Next to face the market's attention was the French presidential elections. February saw a rise in French yields as the potential election of President Marine Le Pen raised its head with markets feeling more relaxed about French elections in March.

Emerging FX started the year strongly after President Trump's protectionist policies looked less likely to materialize and benefitted from the newly inaugurated President's attempts to talk down the dollar. The DXY fell 1.8% through the quarter while the Indian Rupee and the Mexican Peso, among others, rallied through to the end of Q1. The Indian Rupee gained in Q1 on political stability and healthy macro-economic factors while the Mexican central bank has repeatedly increased interest rates over the past few months and has also been selling US dollars on the market providing a cushion under the Mexican currency. The ECB has continued to stay lower for longer, pledging to keep existing stimulus in place until the end of the year while the Fed was unable to support the greenback at the beginning of the quarter against President Trump's anti-dollar rhetoric. A March rate hike from the Fed was accompanied by dovish comments keeping support fairly low on the US currency.

CoT data in March showed the largest Australian dollar long positioning for non-commercials in almost a year as the Australian currency rallied its way through Q1. The British Pound in March, on the other hand, hit record levels of non-commercial short positioning as the British currency's fate remains unknown with the triggering of Article 50 and the beginning of the Brexit negotiations now underway. Despite this uncertainty the implied volatility of the British Pound remained low in Q1 while the Euro dropped before catching up in March. The French election in the first two months of Q2 may see a demand for insurance and a rise in implied Euro volatility. Liquidity in the FX markets dropped through the quarter but remains good.

The generic trender applied to the FX markets was overall negative with the worst performance coming from the New Zealand dollar and the best performance from the Indian Rupee. The RSI indicator applied to the FX markets produced a low of 35 for the Japanese Yen at the beginning of the year following the precipitous decline of the Japanese currency in Q4 2016 and, unsurprisingly, a high for the Indian Rupee of 64 at the end of Q1.

The return of one US Dollar measured in Mexican Peso, British Pound, Euro, Indian Rupee, and the Australian Dollar for the past year

Alternative industry performance

The first quarter of 2017 produced flat performance for the industry's biggest CTAs as the Societe Generale CTA index gained 0.07% in Q1. Average absolute correlations in the CTA universe, an indicator of diversification, continued to fall, a pattern that has persisted since mid-2016, although the level remains above the previous low seen in 2015. Equity Market Neutral strategies fared better with a Q1 performance of 1.5% for the HFRX Equity Market Neutral Index. Among the other strategies within the HFRX database, the best performance in Q1 came from the Event Driven index while the worst was the Macro/CTA index.

Total returns for Equity Market Neutral (EMN) and CTA hedge fund indices over the past year[8]

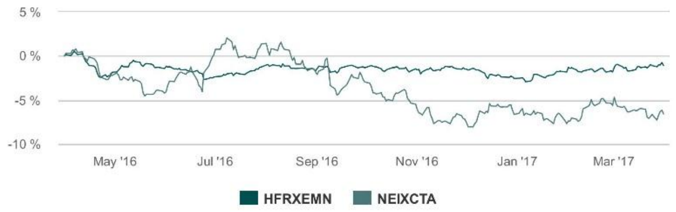

The principle implied volatility indices across four asset classes over the past year[9]

The log of the dollar risk weighted average daily volume across futures on the four asset classes over the past year[10]

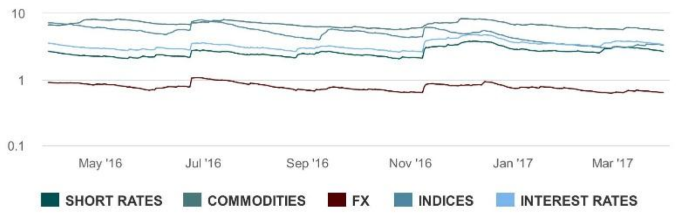

The total return of the trender defined in the text over the past year

Other news

- CFM is holding its 4th Annual Spring Seminar in New York on May 11th. Speakers include Philippe Jordan, Jean-Philippe Bouchaud and Philip Seager from CFM with guest speakers Prof. John Mulvey from Princeton University and Michael J. Mauboussin of Credit Suisse.
- Our paper "Optimal Trading with Linear and (small) Non-Linear Costs" has been published in Risk in March, and "Trading Lightly: Cross-Impact and Optimal Portfolio Execution" has been accepted in Risk.
- Guillaume Simon, a researcher in our Alpha Strategies team, has co-authored the recently published book "Engineering Investment Process: Making Value Creation Repeatable" - https://www.elsevier.com/books/engineering-investment-process/ielpo/978-1-78548-162-8
- Our research has been heavily quoted in the recently revised and extended book "Trend Following" by Michael Covel.

[8] The EMN index is that calculated by HFR, while the CTA index is calculated by the Société Génerale

[9] For the EUR/USD exchange rate we use the Bloomberg defined EURUSDV1M ticker. The VIX index is calculated and published by the CBOE

[10] We estimate effective FX volumes to be a factor of 5-10 more than this due to the extra liquidity available through the spot markets

Extended abstract

Optimal trading with linear and (small) non-linear costs

Paper by A. Rej, R. Benichou, J. de Lataillade, G. Zerah, and J.-Ph. Bouchaud

Determining the optimal trading strategy in the presence of a predictive signal and transaction costs is of utmost importance for quantitative asset managers, since too much trading (both in volume and frequency) can quickly deteriorate the performance of a strategy, or even make the strategy a money-losing machine. The detailed structure of these costs is actually quite complex. Some costs are called “linear” because they simply grow linearly with the traded volume Q. These are due to various fees (market fees, brokerage fees, etc.) or the bid-ask spread and usually represent a small fraction of the amount traded. Much more subtle are impact-induced costs, which come from the fact that a large order must be split into a sequence of small trades that are executed gradually. This cost clearly increases faster than Q, since the price impact itself increases with the size of the trade, as $Q^{3/2}$.

As a simplifying assumption, one often replaces the empirical $Q^{3/2}$ behaviour by a “quadratic cost” formula proportional to Q^2. In the absence of linear costs, the optimal strategy may be found as a result of a simple quadratic optimization problem, as first explored by Garleanu & Pedersen. The optimal policy is to rebalance at finite speed towards the target portfolio. This results in a position that is an exponential moving average of the trading signal.

The pure linear cost problem (i.e. without impact) was considered independently by Martin & Shoneborn and our group. Optimal trading requires instantaneous rebalancing towards a finite “band” around the ideal position desired by the predictors, and no action inside the band, also called the no-trade region.

The case where both linear and quadratic costs are present is of course highly interesting but no exact solution is known at this stage. The aim of our work is to show that one can in fact expand around the general solution for linear cost, when the impact cost parameter is small. Using matched asymptotic expansion techniques, we find that the trading speed vanishes inside a band that is narrower in the presence of market impact by an amount that scales as a cube root of the market impact parameter. Outside the band we find three regimes: a small boundary layer where the velocity vanishes linearly with the distance to the band, an intermediate region where the velocity behaves as a square-root and an asymptotic region where it becomes linear again.

All our results above are compatible with numerical results and analytical results obtained by the group of Muhle-Karbe. Our solution enables one to estimate the impact of cost parameters on the final costs. Our asymptotic expansion method can be undertaken for arbitrary non-linear costs, for example the empirical $Q^{3/2}$ result.

Whitepaper

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[11] (ARP). We use a Trend Following strategy to illustrate the effectiveness of the approach in building a robust portfolio of correlated instruments.

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/√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.

[11] 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"

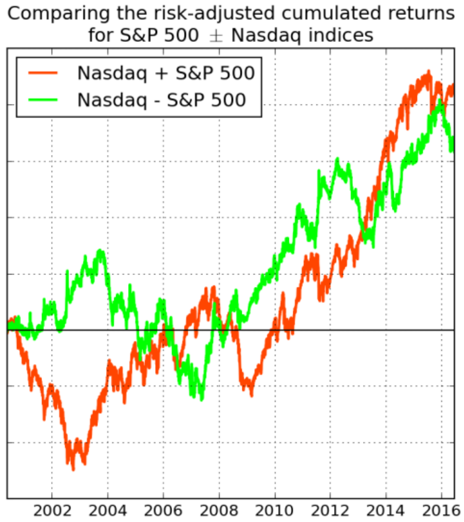

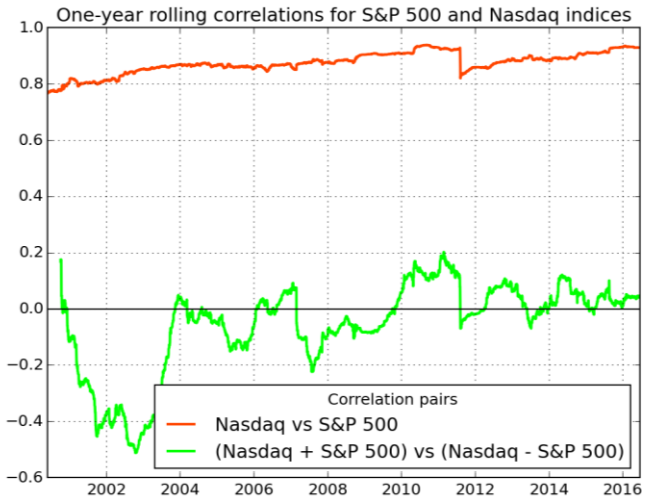

Figure 1 The upper plot shows the P&L curves of an equal long/long position in the Nasdaq 100 and S&P 500 futures and the decorrelated equal long/short positions. On the lower plot we show the correlation between these two futures and the correlation of the long/long combination with the long/short combination. 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 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[12] due to its ability to allocate to portfolios that have worked in-sample, in the past, with the benefit of hindsight[13]. 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[14] which is an essential component to being able to allocate correctly across robustly determined Risk Factors[15].

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.

[12] 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 optimization techniques

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

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

[15] 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

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 polarized or concentrated positions that end up being less diversifying rather than more. Measures of diversification, unfortunately, depend on the "basis"[16] 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)[17]	Portfolio Risk(%)[18]
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

[16] 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

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

[18] 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

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

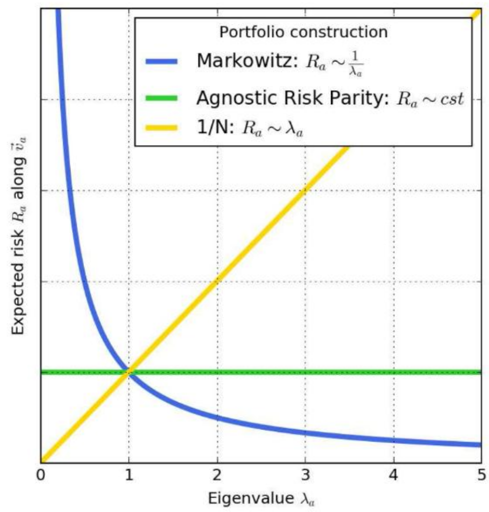

Figure 2 – The allocation of risk across the principal components of the universe for the three portfolio construction techniques considered. The x-axis corresponds to the "eigenvalue" of the principal component, which is mathematically equivalent to the variance, or the volatility squared, of each. 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.

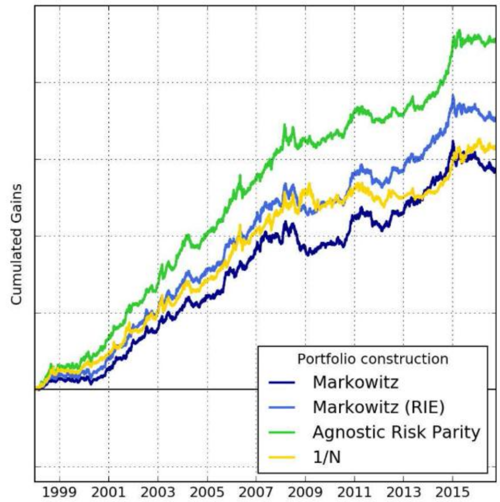

Figure 3 - The P&L performance curves of each of the allocations considered using a long term trend following signal across a universe of 110 typical CTA instruments made up of equity indices, bonds, FX and commodities. Each curve is normalised to have the same risk, showing that the ARP approach gives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

[19] 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

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 maximized 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