INSIGHT.DATA.CLARITY.

ALTERNATIVE BETA MATTERS

Quarterly Report

All Data as of 31 December 2015

Introduction

Welcome to CFM's "Alternative Beta Matters" Quarterly Report for Q4-2015.

Within this report we recap major developments of the quarter within 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 two white papers and an extended academic abstract from a paper produced 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 a reference for our readers, and can stimulate conversations on these topical issues.

RISK PREMIUM INVESTING – A TALE OF TWO TAILS

This white paper discusses Risk Premium investing, and with examples, explores the idea that the premium is a compensation for downside or negative skewness risk. The existence of positively skewed, positive Sharpe non-risk premium strategies is also introduced, followed by a discussion of how this is relevant to portfolio construction.

THE MISLEADING NATURE OF CORRELATIONS

Correlation measurements are included in many investment analyses and are often blindly used to identify decorrelating and diversifying investments. This white paper takes a closer look at what is represented by correlation, and at what level a correlation should be considered statistically signficant.

DECONSTRUCTING THE LOW VOL ANOMALY

The low volatility anomaly has been observed and documented for nearly half a century; this summary of an academic paper discusses the origins and persistence of this anomaly.

Equities

Stocks had a rough year in developed countries but 2015 was particularly difficult for emerging markets. The MSCI World index, made up of 23 developed country equity indices, returned -2.7% on the year, while the MSCI EM index, with 23 emerging countries, fared much worse, returning -16.95%. The Fed's intentions to raise rates weighed heavily on countries that previously provided a hedge against the post-2008 crisis zero dollar interest rate environment. In 2015, investors looking to flee such risky destinations sought to repatriate back to less risky, newly interest paying markets on US shores.

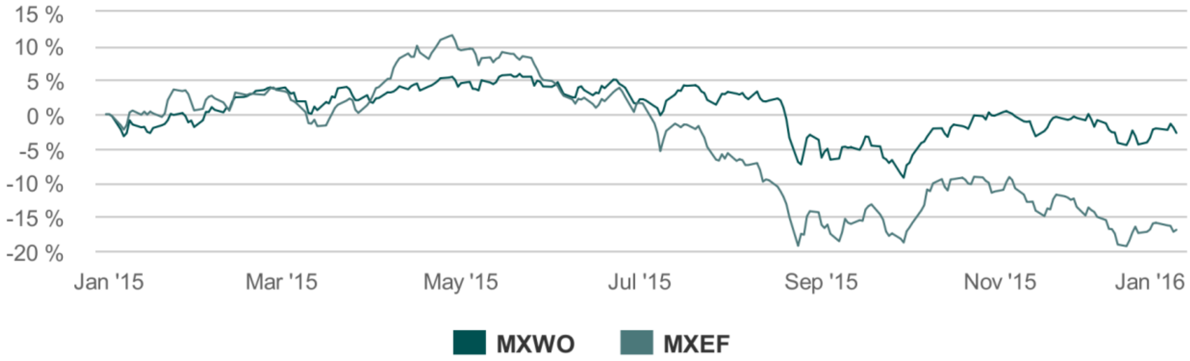

The return of the MSCI World and MSCI Emerging Market indices for the past year

Emerging and developed market equities have become more correlated over the past quarter, suggesting that the global economy has recoupled a little; the US will have a higher chance of performing if emerging markets perform well. Less developed markets also tend to be based on less sophisticated economies, often reliant on commodity exports. The selling pressure on many commodities in 2015, arising from the continuing OPEC oil glut and from an underperforming Chinese economy, also contributed to the difficult environment for emerging markets. Tumbling oil prices not only hit emerging economies, but developed market indices' energy firms also felt a negative drag. Crude and US equities have increased in correlation over the past quarter, confirming this concern of the effects of low commodity pricing on equity performance.

The past quarter has seen developed market indices mostly ranging following an initial rally in October with the MSCI World returning 5.1% in Q4. Only the FTSE100 showed a slight tendency towards the downside in the latter half of the quarter. This slight downwards trend provided a flat performance for the simple trend follower[1] with all other developed market equity indices being poor performers for the trend.

The FTSE100 achieved the lowest equity RSI[2] score in Q4 with its lowest point in the middle of December. Implied volatilities have settled back to more normal levels following the Chinese sell-off in August and the VIX finished the quarter at the very reasonable level of 18.2 points. The Chinese market rebounded in the final quarter of the year, returning 16.5%, and remained positive on the year despite the spectacular June to September drop, with an overall 2015 return of 5.6%.

[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

Equity index liquidity has dropped off to a level consistent with Q2 following the spike in August, with end of year liquidity characteristically dropping off through the festive holiday period. Individual equities followed indices with a similar pattern in liquidity that dropped post-August. In terms of industrial sectors, energy stocks continued the trend of the past couple of years in maintaining the highest level of volatility implied by the market. This is a regime that began, unsurprisingly, with OPEC's decision to maintain supply. Liquidity across industrial sectors was, however, more homogeneous with no notable differences from one sector to another.

Fixed Income and Credit

The year has been one of diverging monetary policy and also therefore of diverging fixed income pricing. The Fed has been on its way to lift-off, while the ECB and the BoJ have been desperately attempting to head in the opposite direction in an effort to stave off deflation and kick start their lethargic economies.

December was a case in point. The ECB surprised the market on December 3, by having overpromised and, subsequently, underdelivering on further easing. The Fed, on the other hand, was widely anticipated to increase rates and followed its course to lift-off, its first rate increase in almost a decade. The past quarter has seen Eurodollars moving in the direction of tightening and Euribor carrying on down to zero rates and beyond. The BoJ's incessant bond buying seems to be having the desired effect in pushing Japanese long rates lower, while in Canada the central bank has been fighting against Fed pressure in trying to keep rates low.

Interestingly, the lowest RSI score this quarter was the Australian Bank Bill future, which reached 30 on December 11 and was the best performer for our trender (Euribor was second with a rates trend on the downside), while Eurodollars were the worst trenders, reversing course in the direction of lift-off mid-quarter. The trend has generally not worked as well for long bonds, with only the JGB and Canadian 10-Year bonds finishing the quarter positively. The Commitments of Traders data showed non-commercials reversing Eurodollar positioning from long to short through Q4, pre-empting the Fed's rate rise.

Liquidity at the short end of the curve has been good, especially around the various central bank announcement dates. Liquidity changes have been less pronounced at the long end of the curve, with a more stable evolution of trading volumes. Overall, with diverging rates across regions, developed market bonds were flat through Q4 with the Barclays Global Aggregate Total Return Hedged index returning 0.1%.

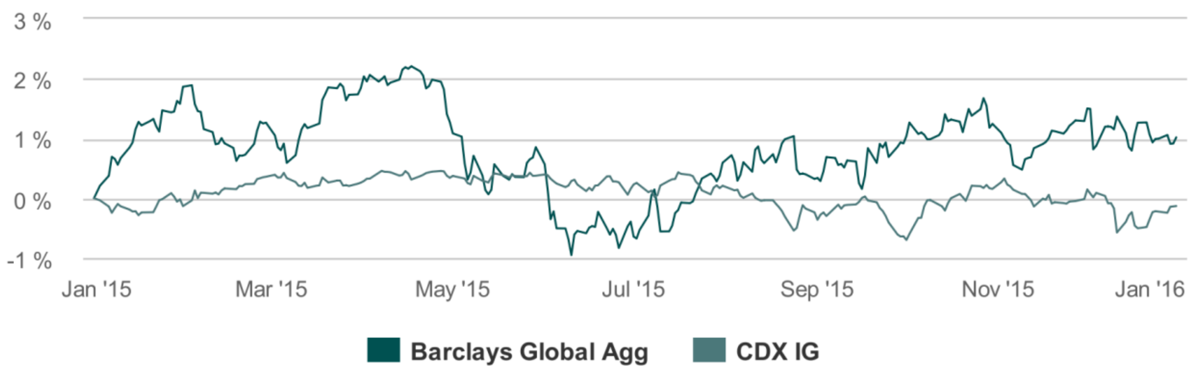

The 2015 total return of the Barclays Hedged Global Aggregate Bond and the CDX Investment Grade indices

Corporate credit conditions oscillated in the final quarter of the year as the junk bond market started to look increasingly fragile ahead of the Fed's interest rate increase. All four of the major corporate CDS indices fell sharply in December, along with corporate bonds and, in particular, high yielders on news of the biggest US mutual fund liquidation since 2008. High Yield Bond ETFs slumped and CDS spreads rose as a prominent Wall Street firm shut down its $800m credit fund after a wave of losses and redemptions. Credit conditions improved heading into the year-end festive period, but corporate leverage remains high and fund flow data points only to outflows.

Commodities

Q4 saw a sell-off in most major commodity markets, with the GSCI falling 13.3% through the quarter, despite beginning October with a brief rally. The glut of Crude oil has continued to pressure energy prices, in particular at the end of the year. Natural gas was the exception, having tumbled along with Crude, until the second half of December, when prices rose with the return of cold weather in North America. The natural gas market, despite the December rebound, provided the best environment for commodity trending. Crude implied volatility and liquidity peaked in December with the OPEC meeting creating uncertainty in the continuation of the excess supply policy. Uncertainty on the demand side has also provided a drag on commodity prices.

2015 return of the S&P GSCI

Chinese economic underperformance has created waves along with emerging markets struggling to keep up with the appeal of the US dollar. Industrial metals have historically benefitted from China's success and therefore have also suffered due to its woes over the past quarter. Copper, a commodity that China has historically consumed in bulk, was also one of the leaders for the trend downwards and had the lowest RSI score, hitting 36 on 24 November.

Falling prices in Q4 have made commodities the best performer for our trender, with performance on all other asset classes being flat. The Commitments of Traders data showed positioning for non-commercial Crude players as staying long but reducing positions through the quarter. Positioning was net long for the whole of 2015. This contrasts with Natural Gas positioning data, which showed non-commercial players being net short through the year. Overall commodity liquidity has fallen, in particular at year end, but conditions have reverted back to liquidity levels seen in the middle of the year.

FX

FX markets, like fixed income, were affected by the underlying themes of the quarter: diverging central bank monetary policy, the commodity sell-off and Chinese economic underperformance. The Euro sold off through Q4 until the ECB underwhelmed with future easing. The 3 December announcement resulted in the Euro's biggest single-day gain in six years, as investors covered their shorts.

It is interesting to note how the correlation of the Euro/USD rate with the S&P 500 has evolved over the past few years. When the European crisis was in its infancy, the correlation was positive. It has since steadily slipped past zero into negative correlation territory, as the ECB tries to stimulate and the Fed reverses course.

Those currencies sensitive to the price of Crude have been adversely affected; notably the Canadian Dollar, the Mexican Peso and the Russian Rouble. Repatriation of funds back to the greenback has provided downward pressure on everything quoted against it with few currencies finishing out the quarter ahead, only the Australian and New Zealand dollars have finished the quarter slightly in the black. This was reflected in the Commitments of Traders data that showed non-commercial positioning to overall favor being long the dollar. The New Zealand dollar, in particular, experienced a rise in long positioning through Q4.

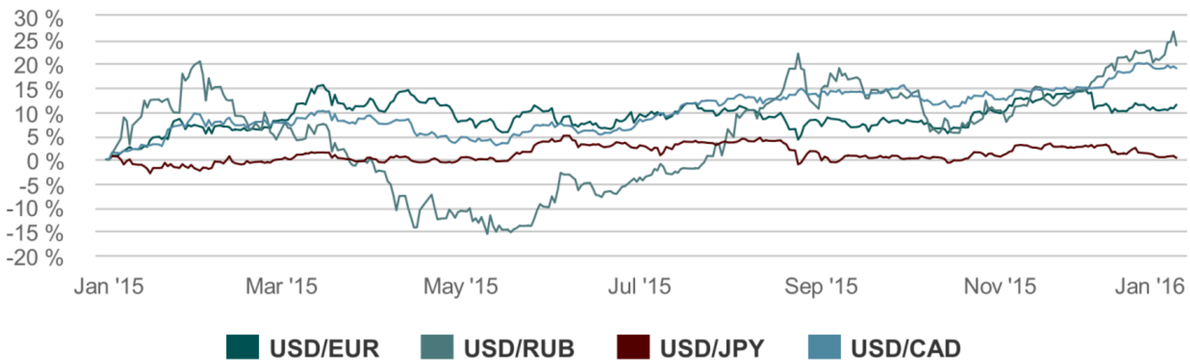

The return of one US dollar measured in euros, roubles, Japanese yen and the Canadian dollar for the past year

Implied volatilities were high coming into year-end with the high density of central bank meetings and expected policy announcements. FX liquidity has, however, stayed stable through the year. The trender has been poor in FX globally, working well for those pairs that sold off, such as the Canadian Dollar and the Russian Rouble, but poor for the ranging currencies such as the Japanese Yen and the Euro. The trender also performed well on the South African Rand, which tumbled in value towards record lows at the end of the year. Mid-December saw South African equities and the Rand fall as Finance Minister, Nhlanhla Nene, was removed from his post by the South African President, Jacob Zuma, to be replaced with a "relatively unknown member of parliament".

Alternative Industry Performance

2015 was underwhelming for the biggest CTAs as the Societe Generale CTA index finished out the year at 0.03%. The quarter fared little better, finishing at 0.06%. The damage was done earlier in the year as Q2 was a rough ride for the index. Average correlations between contracts in the CTA universe have been steadily increasing through the second half of 2015, although levels remain low compared to the peak of 5 years ago.

Equity Market Neutral funds fared better with the HFRX EMN index returning 0.87% in the quarter and finishing the year at 5.45% with the best of the performance coming from the third quarter of 2015. The best HFRX index in the quarter was for Merger Arbitrage which finished at 3.2% while the worst was the Distressed Restructuring index which saw out Q4 at -6.72%.

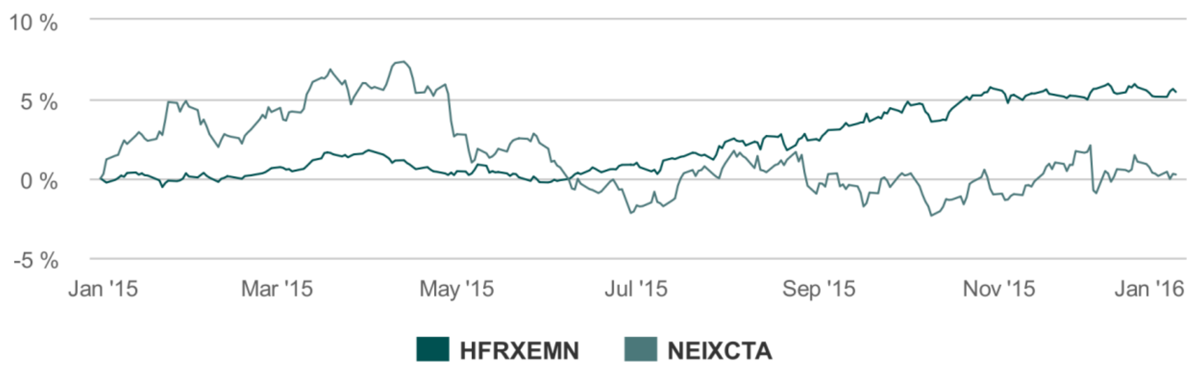

Total returns for equity market neutral (EMN) and CTA hedge fund indices over the past year. The EMN index is that calculated by HFR, while the CTA index is calculated by the Société Génerale[3]

[3] https://cib.societegenerale.com/en/sg-prime-services-indices/

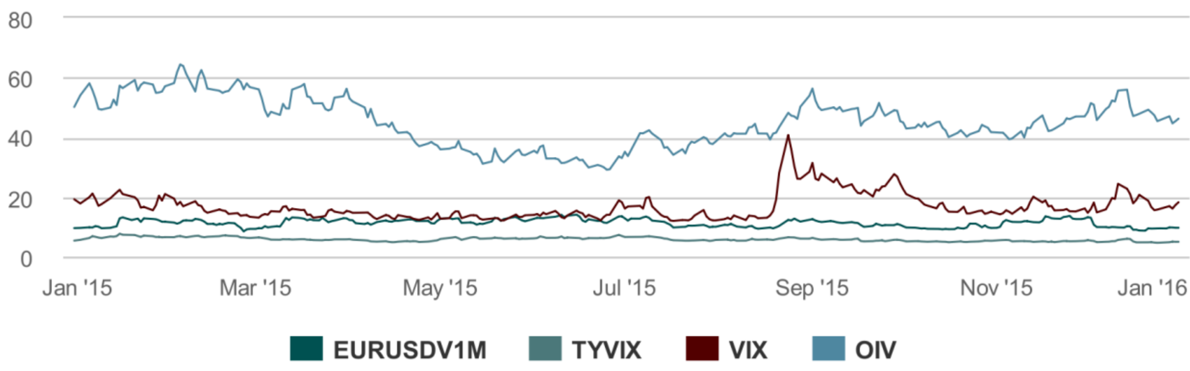

The principle implied volatility indices across the four asset classes over the past year. For the EUR/USD exchange rate we use the Bloomberg defined EURUSDV1M ticker; the TYVIX and VIX indices are both calculated and published by the CBOE while the OIV index is published by the CBOE and NYMEX exchanges.

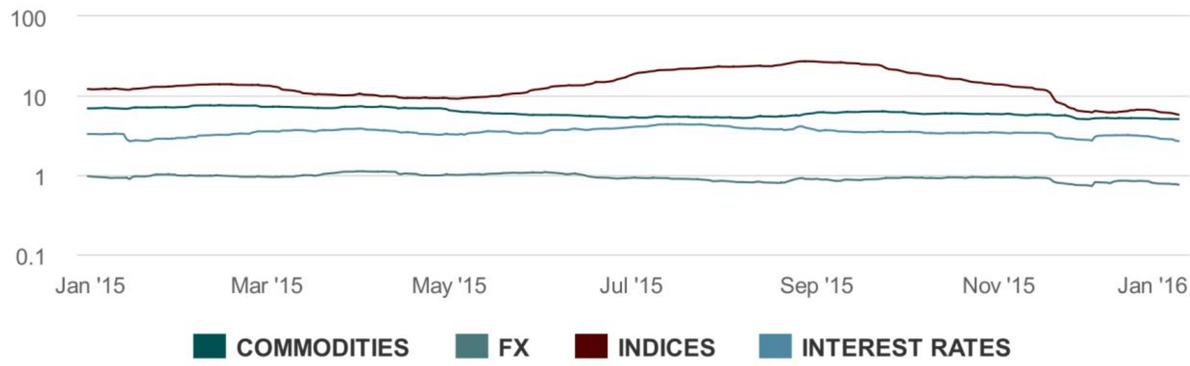

The log of the dollar risk weighted average daily volume across futures on the four asset classes over the past year. We estimate effective FX volumes to be a factor of 5-10 more than this due to the extra liquidity available through the spot markets.

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

Risk Premium Investing – A Tale of Two Tails

EXECUTIVE SUMMARY

In this note we introduce Risk Premia as generically encompassing a set of strategies where investors are compensated for assuming risk. This compensation comes in the form of a regularly received premium and results in a strategy with a positive expected return. This positive performance continues until the moment the risk, which is being assumed, is realized, resulting in a sharp negative move. The argument generally invoked is that this premium is proportional to the risk or volatility of the investment or instrument being held, a measure which uses both sides of the return distribution. In this paper we expand on the idea of a risk premium, with the help of a few example strategies, and show that the premium is in fact compensation for downside deviation or negative skewness risk, an idea which is justified with empirical evidence and also appeals to common sense.

INTRODUCTION

Presented with two investments of equal expected return, a rational investor will choose the one with the lowest risk. This statement can also be turned around as, given two investments offering differing levels of risk, the higher risk investment needs to provide a higher level of excess return in order to entice people in to invest. Within the Efficient Market Hypothesis (EMH), excess returns are *only* possible due to the existence of a risk premium. In fact EMH proponents will often explain away positive excess returns as being due to the existence of some form of hidden or latent risk. This risk, in financial parlance, is generally represented as the standard deviation (although sometimes other measures are used) of a time-series of returns. These standard measures of risk use both sides of the return distribution and do not distinguish up moves from down moves.

The existence of a risk premium as being compensation for assuming a volatile position makes sense. However, the idea does not satisfactorily account for investments exhibiting negatively skewed, asymmetric return streams. We instead, therefore, propose that a risk premium represents compensation for an investment that has a non-zero probability of selling-off sharply and we use the following example to illustrate our point. A short position in VIX futures is a profitable trade over a long back-test. At the onset of any crisis, however, experience tells us that it is a strategy that can have significant downside risk. A long position in VIX futures, though, has an equal amount of *two sided* risk or volatility to a short position, but with the opposite return profile. Are we to believe that a long position and a short position in VIX futures should have the same amount of risk premium? Clearly not and there is, therefore, a sign ambiguity which is resolved by considering risk premium as compensation for downside risk.

So, what does downside risk mean for a strategy? If we were to plot out the performance of a stream of returns with higher downside risk than upside risk we may see something like the curve plotted in **Figure 1**. We have also plotted a zero skewness (symmetric) strategy and a positively skewed (upside risk higher than downside risk) strategy for comparison. The negatively skewed investment is characterized by having frequent small positive returns and infrequent large negative returns. **Figure 1** also shows the cross sectional profile of the returns of the strategies showing the differences between the three regimes. Any investor, given the choice between these return profiles, will prefer a positive Sharpe/positive skewness strategy and it is this bias against negatively skewed strategies that creates the risk premium.

In this note we present our approach to improving the skewness profile of a risk premium offering through both diversification and the power of the combination with non-risk premium, behavioral anomaly strategies such as trend following.

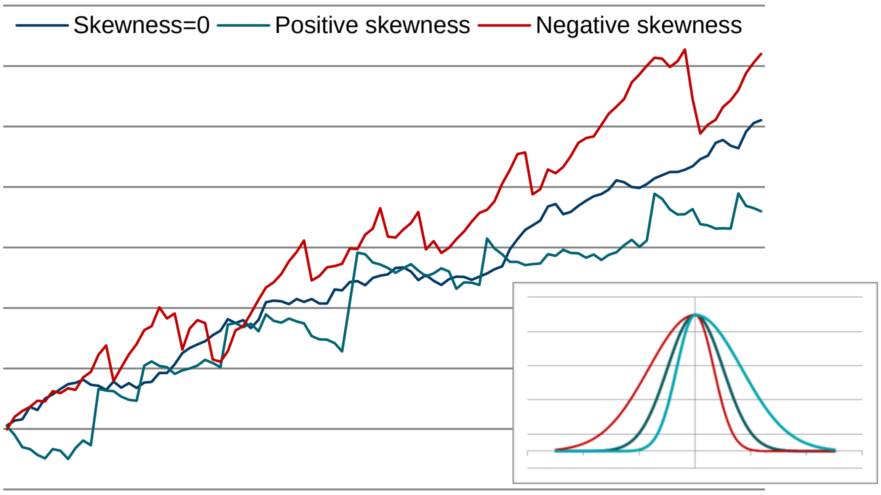

Figure 1: Three strategies with different skewness profiles (negatively skewed, zero skewed and positively skewed). The superimposed plot shows the cross sectional profile of the corresponding returns.

INSURANCE – A CLASSIC EXAMPLE OF A RISK PREMIUM STRATEGY

Insurance policies cover a multitude of outcomes from the theft of a car or a house burning down, to a professional footballer's capacity to play football. Most people are familiar with the concept of insurance, but less so with how insurance companies generate profit. Within an insurance company, actuaries attempt to evaluate the probability of an insured outcome occurring, and subsequently calculate a fair value for a given policy. A mark-up is applied to the cost of the policy representing the premium charged to the end client such that on average a profit is generated. The return profile for the insurance company is, therefore, a stream of positive payments from the premiums received and an occasional large negative move when the insured outcome actually occurs and requires a pay-out. An insurance company will not just rely on one policy to run its business, diversification is essential to the model in reducing the negative skewness profile of the returns and improving the comfort level of the company's shareholders. It is this decorrelation between the pay-outs that is crucial to reducing the downside risk of the firm, a point we will address below.

A concern of any insurance firm shareholder will be the potential for a non-diversifiable component of risk. For example, let's say that an insurance company has sold policies insuring houses across a country or continent. Let's also say, for the sake of argument, that a natural disaster occurs - a meteor destroying a large industrial area, or flooding that destroys homes and factories, or an earthquake, a tsunami etc.... It is the existence of these *correlated* events requiring an instantaneous profusion of pay-outs across many clients that represents an incompressible risk for the insurance firm. This risk has been addressed more recently with the introduction of catastrophe bonds to the market place. A hedge fund or pension fund may decide to offer the insurers insurance against these types of correlated pay-outs, but of course this insurance comes at a price![4] We discuss the idea of incompressible risk in a financial context below.

[4] Insurers may also invoke exceptional "Act of God" clauses in insurance contracts to avoid this risk

A FEW EXAMPLES OF RISK PREMIA (AND OTHERS THAT ARE NOT)

We begin with some familiar examples of risk premia from the world of traditional asset classes and investments. Holding equities is a well-known strategy and is often a big component of investor portfolios. Although the skewness of individual equities is quite small, when combining stocks together to form an index, the skewness becomes negative. This "leverage effect" is well known, and is due to the fact that stocks tend to correlate in their downward moves, creating the negative fat tail for the index, and thus making long-only equities a risk premium strategy. Experience of stock markets tells us that they tend to drift upwards over long periods, albeit with weak Sharpe ratios, but can sell-off sharply. One can see the measured Sharpe ratio and skewness[5] in **Error! Reference source not found.**, showing that equity indices are our first example of a risk remium strategy.

Instruments/Strategies	Skewness	Sharpe ratio
S&P 500	-0.5	0.5
DAX	-0.5	0.4
Nikkei 225	-0.1	-0.1
2 year US government bond	0.7	0.7
10 year US government bond	0.0	0.7
2 year German government bond	-0.4	0.7
10 year German government bond	-0.2	1.0
Corporate bond index (US B rated)	-1.8	1.7
Short equity index volatility	-1.6	1.3
Short 10 year bond volatility	-2.2	0.9
Short energy volatility	-2.2	0.4
Short metal volatility	-2.0	0.6
Short grain volatility	-1.4	0.5
Short FX volatility	-1.6	0.2
UMD US	-0.7	0.7
HML US	0.2	0.5
Trend following	0.4	0.8

Table 1: Skewness and Sharpe ratios for a set of instruments and strategies.

Of perhaps more interest for this note is another traditional asset class - that of fixed income. Short dated government paper, issued by the most creditworthy governments such as the US and Germany[6], exhibits positively skewed returns due to the safe haven nature of the investment. In periods of risk-off and flight-to-quality, investors seek to hold these types of investments, even more so than cash (which necessarily involves being exposed to a potentially uncreditworthy bank!). The post 2008 climate of extreme, loose monetary policy has, however, pushed rates to record lows and investors to seek other sources of return. Indeed it is precisely for this reason that central banks employ such policies! A portfolio manager may, therefore, turn to the world of corporate bonds to try to get an extra "pick-up" on his cash, with corporates needing to offer a higher level

[5] We use the "non-parametric skewness" as defined by en.wikipedia.org/wiki/Skewness. This is a more stable measure than the third moment, a point which is discussed in [1].
[6] One may argue that Germany is now less creditworthy due to the fact it no longer controls its own money supply and can therefore not just print cash to pay off its debts

of return compared to a government in order to entice investors to buy their debt. Of course, the investor is fully aware of the fact that if things don't go well for the firm then the bond principal is at risk and may not be repaid. However, intuition is strong that this risk can be diversified away (we will describe this more quantitatively later) by buying a *basket* of corporate bonds offering a wide ranging mix of corporate debt. This approach allows the investor to diversify away a lot of their credit risk, whilst benefiting from the pick-up in yield versus government bonds. The investor may be exposed to a non-diversifiable risk that defaults correlate across firms. It is for this reason that an investment in a diversified basket of corporate debt will yield more than government paper, as compensation for the existence of this correlated default risk. In **Table 1** we summarize the case of US short dated government paper exhibiting a positive skewness meaning it is not a risk premium strategy, while longer dated government bonds and, certainly, corporate bonds stand-out as clear risk premium strategies.

Insurance selling in finance comes from the options market. If a big pension fund is holding a large equity portfolio and is fearful of adverse moves due to an upcoming FOMC, the pension fund may choose to protect its portfolio by buying down strike put options. In so doing, the investor is pushing option prices up, implied volatility up relative to realized volatility, and paying a premium to the option seller. The option seller is providing protection and so requires a premium for taking on the risk that the hedger wants to offload. If nothing happens through the FOMC then the option seller has made money by receiving a high price for the insurance. If, however, the move that the investor feared actually comes to fruition then the option seller faces a potentially large pay-out. This risk can again be mitigated and the negative skewness reduced by diversifying over many different options on many different "underlyings." When applied to this example, the idea of non-diversifiable or incompressible risk that we touched upon earlier corresponds to a volatility move that correlates across underlyings. In the biggest financial crises one observes that this is indeed the case, with losses being seen across asset classes. We do, however, see an improvement in the negative skew and Sharpe ratio characteristics from a short volatility portfolio that is diversified across many underlyings, which suggests that we are more able to diversify away this risk than one might imagine.

At the beginning of the 1990s, Eugene Fama and Kenneth French identified certain market neutral portfolios in the US stock market that seemed to exhibit positive excess returns. Here we focus on two Fama and French factors – UMD[7] and HML. UMD refers to "Up Minus Down" as being a market neutral portfolio of the trend applied to individual stocks, exploiting the fact that outperforming stocks tend to continue to outperform while underperforming stocks continue to underperform. Measuring the skew of this strategy shows that the factor exhibits negative skewness. What is interesting, however, is that the large downside moves come from upside moves in the index. Despite the fact, therefore, that the strategy has negative skewness, it is a true diversifier in the sense that the negative skewness comes from risk-on environments occurring at times unrelated to the drawdowns of other risk premium strategies.

Let us now switch our attention to the HML factor. HML refers to "High Minus Low", a market neutral factor that is long stocks that have a high book value[8] to price ratio and short stocks that have a low book to price ratio. The strategy is a classic "value" strategy in that one exploits the fact that prices should move in tandem with a fundamental value metric such as book value. We find in this case that HML exhibits a positive skewness. A positive Sharpe ratio strategy with positive skewness is rare indeed. Our claim in this short note is that positively skewed, positive Sharpe ratio strategies do not come from the family of risk premia, but rather generate returns by exploiting some form of market anomaly. The case of the HML factor is curious, however, and not easily explained. The skewness of the strategy turns negative when using monthly data, adding to the confusion, and necessitating further work. **Table 1** summarizes the statistics for the two Fama French factors using stocks across a range of geographical zones.

The final example we will discuss here is another case of a strategy exhibiting positive skewness. Trend following is a special case due to its relationship with risk, as has been discussed in our paper on the subject [2]. Trend following will perform in times of market stress if the difficult period persists for a time comparable to, and ideally longer than, the timescale over which one trends. That being the case, a classic trend following system will short crashing equities and other risk sensitive instruments and will buy bonds and other flight-to-

[7] UMD was not one of the original Fama French factors but was introduced in subsequent factor literature

[8] Book value is assets minus liabilities. In the absence of earnings it is a good starting point for valuing a company

quality instruments. As long as the crisis period persists for longer than the trend timescale, there is a guaranteed payoff. This "long volatility" component of trend following is what gives us a positive skew and was the reason for its success through the 2008 global financial crisis. We would argue that the trend is in fact not a risk premium strategy, but rather a genuine market anomaly arising out of the extrapolative tendencies of investors.

Many other examples of strategies have been studied and we invite interested readers to consult [3] for a more comprehensive list and description.

IS RISK PREMIUM COMPENSATION FOR HIGH VOLATILITY OR NEGATIVE SKEWNESS?

We have presented our understanding of risk premium strategies as having positive Sharpe ratios and negative skewness, with the positive excess returns being the reward for assuming downside risk. This idea appeals to common sense, but can we show this to be the case empirically? CFM has done research on this subject [3] which culminated in the plot seen in **Figure 2**. We took as many standard, well-known, accessible strategies as we could, many of which are discussed in the previous section, and plotted the Sharpe ratio against the level of skewness. One sees a cloud of points that are dispersed around an explanatory line, suggesting that indeed, the higher the reward or Sharpe ratio, the higher the level of negative skewness. The suggestion is, therefore, that strategies that lie below the line have too much negative skewness compared to the level of excess return while, contrarily, strategies above the line have too much excess return for the level of negative skewness. This is indeed the case for the two strategies discussed above - trend following and the HML equity market neutral factor. One observes that both have positive excess returns with a positively skewed return distribution and, for this reason, our conjecture is that, rather than these strategies being Risk Premia, they instead represent genuine market anomalies.

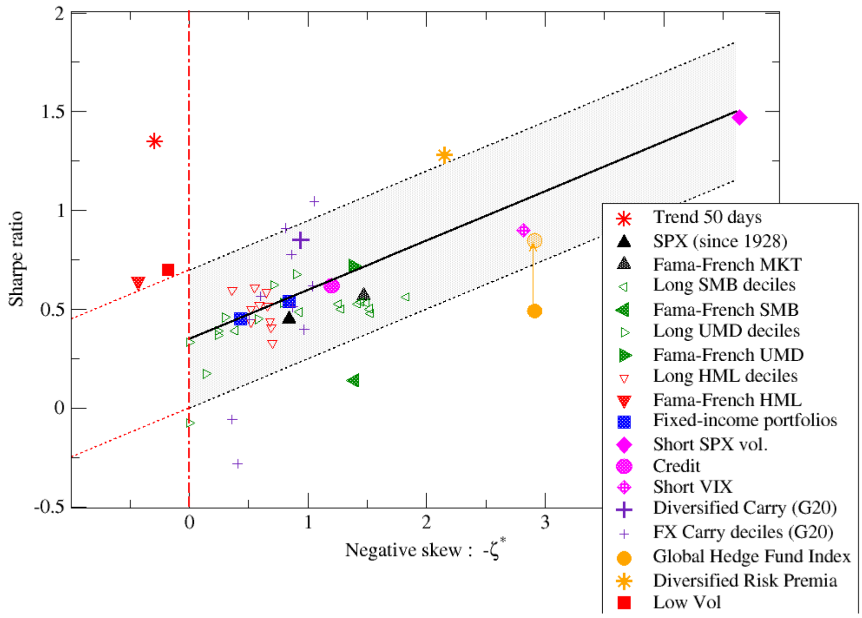

Figure 2: The "skewness rewarding line" showing a quasi-linear relationship between the Sharpe ratio (the reward) and the negative Skewness (the risk), for a set of strategies and instruments.

An interesting question is now whether volatility (the standard deviation of returns) has the same explanatory power. What we see in [3] is that volatility surprisingly tends to do the opposite of skewness in that the low volatility strategies systematically outperform those with high volatility, a result which is at odds with an explanation of a risk premium being due to volatility. This fact also explains the "low vol" effect [4], a recent "hot topic" in the academic community. Interpreting risk premium, therefore, as compensation for downside risk rather than symmetric risk seems to be consistent with the data.

CLASSIFYING AND DIVERSIFYING NEGATIVE SKEWNESS STRATEGIES

As previously stated, we are putting forward the idea that a premium exists for holding a negatively skewed investment. The origin of this negative skewness can, however, be different across a range of strategies. Identifying the nature of the risk premium is important when trying to combine strategies together to produce a higher Sharpe ratio and overall better skewness profile (ideally making it positive).

Combining negatively skewed strategies together should reduce the overall negative skewness of the combination. In order to illustrate our point, we have taken a number of negatively skewed strategies which can be seen in **Figure 3**. Before we combine them together, it is crucial to note that the skewness is not correlated across the strategies, meaning that the down moves do not typically arrive together, simultaneously, across the strategies. We have superposed the combination of these *de-coskewed* strategies on the same plot, and show that both the Sharpe ratio is improved and the level of negative skewness reduced. If we continue adding strategies then the Sharpe ratio will continue to rise and the skewness will fall to zero, meaning we end up with a symmetric return distribution. As discussed above, this is in the absence of non-diversifiable risk which may materialize itself as a correlation of downward moves across all strategies and an increase in negative skewness, something we are trying to avoid.

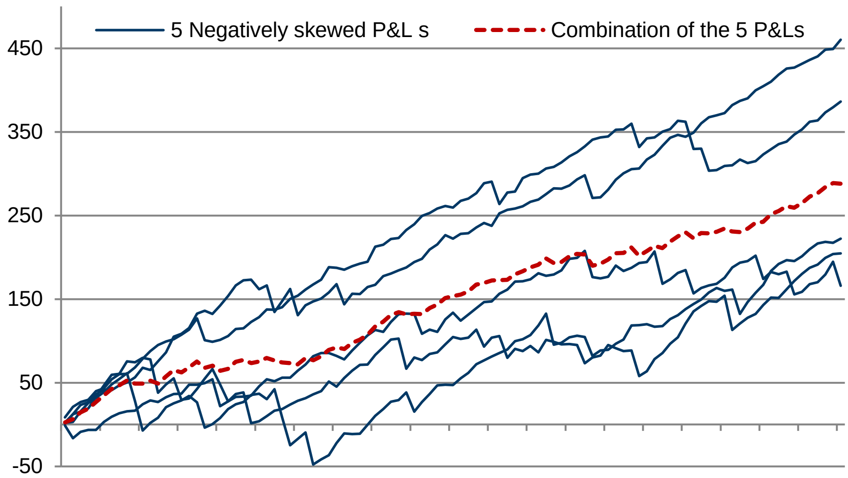

Figure 3: An illustration of the effect of diversification on skewness. The combination of a set of de-coskewed strategies with negative skews results in a combined strategy with better Sharpe ratio and close to zero skewness. The figure is based on simulated data and is for illustration purposes.

We listed several classic, well known strategies in the previous section. The question is when, on average at least, is a strategy more likely to sell-off? If the strategy expresses its risk with a sharp sell-off in a rising market, such as the UMD Fama-French factor, then we consider the strategy to be “risk-on premium” whereas if the strategy sells off during a down market, such as a short option strategy, then we are receiving a “risk-off premium”. Catastrophe bond buyers, on the other hand, take on “catastrophe risk premium” which would incur losses following a natural disaster, which may or may not be correlated with a risk-off move (depending how big the catastrophe is!). Liquidity providers, on the other hand, may be seen as assuming a “liquidity risk premium”, the risk being that they are executed and lose money on large price jumps, in either direction, in the instrument they are providing liquidity for. Identifying many different types of risk premia and allocating to them should, we believe, help to reduce the forward looking negative skewness of a risk premium portfolio. The special case of trend following, which we are claiming to be a behavioral anomaly, is also useful in combination with risk premium strategies. We deal with this case below.

It is interesting to note that the HFRX hedge fund index sits close to the negative skew/Sharpe ratio reward line, as seen in **Figure 2**, albeit sitting slightly below it. We also show in the plot, the effect of correcting hedge fund returns for fees, showing that the point then moves up to sit on the line. This is interesting and may suggest that hedge funds are exploiting risk premia in their strategies. Anecdotally, this seems consistent with a poor overall performance through the global financial crisis of 2008 when hedge funds lost a significant proportion of their assets. Although we see value in a well implemented risk premium portfolio, the advent of Alternative Beta offerings may drive fees down and the HFRX point up closer to the line.

We would also like to bring the reader's attention to the “Diversified Risk Premia” point in **Figure 2**, where we combine a portfolio of equity indices, a diversified portfolio of short options, an FX carry strategy and long the four CDS indices in the US and Europe. This point sits satisfactorily above the regression line and represents a diversified risk premium portfolio. It seems, therefore, that one is able to improve the Sharpe ratio/negative skewness characteristics of risk premium investing through diversification. The idea of non-diversifiable or incompressible risk, that we alluded to earlier in this paper, does not preclude us from improving the characteristics of the investment and, therefore, we favor this approach in building our portfolio. We have studied the possible existence of this incompressible, or “Black Swan” risk, by attempting to identify the *factors* of the combination of all risk premium strategies that went into **Figure 2**. A dominant “market” factor does not seem to exist, however, unlike the case of the market mode for a portfolio of stocks. Instead, two poorly separated and unstable factors emerge, suggesting the absence of a clear benchmark for risk premium investing. The idea of incompressible risk, nevertheless, seems plausible and we favor an approach of attempting to mitigate this risk, as previously stated, by combining and diversifying different sources of risk premia.

A TAIL HEDGE WITH THE TREND

Trend following, both empirically and intuitively, does not exhibit a negative skewness and, as such, we consider it a market anomaly rather than a risk premium strategy. The positive skewness properties of trend following make it a good strategy in combination with risk premia, in particular with strategies which perform poorly in times of market stress. We take, as an example of the latter, a diversified short volatility strategy that is short options on interest rates, equity indices, commodities and FX and is plotted in **Figure 4**. In **Table 2** we list the worst performing months of this strategy, along with the corresponding performance of a trend following system. What is striking is that for 6 out of 8 periods of the worst performance for our short volatility premium strategy, trend following delivers positive performance. This anti-correlation in the tails of these two strategies provides a form of statistical downside protection to the risk premium approach. The two strategies are combined together in **Figure 4**, showing the benefits of a skewness diversified approach.

Start date	End date	Diversified Short Volatility (%)	Trend Following (%)
Aug-08	Nov-08	-31	17.4
Oct-87	Nov-87	-13.1	2.6
Jul-90	Sep-90	-12.1	5.7
Jul-14	Jan-15	-9.6	17.9
Oct-97	Nov-97	-9.6	-0.7
Mar-13	Jun-13	-7.2	1.3
Apr-10	Jun-10	-5.9	3
May-07	Sep-07	-4.9	-0.2

Table 2: The worst performing periods for a diversified short volatility strategy and the corresponding performance of a 50-day trend following strategy in the same period.

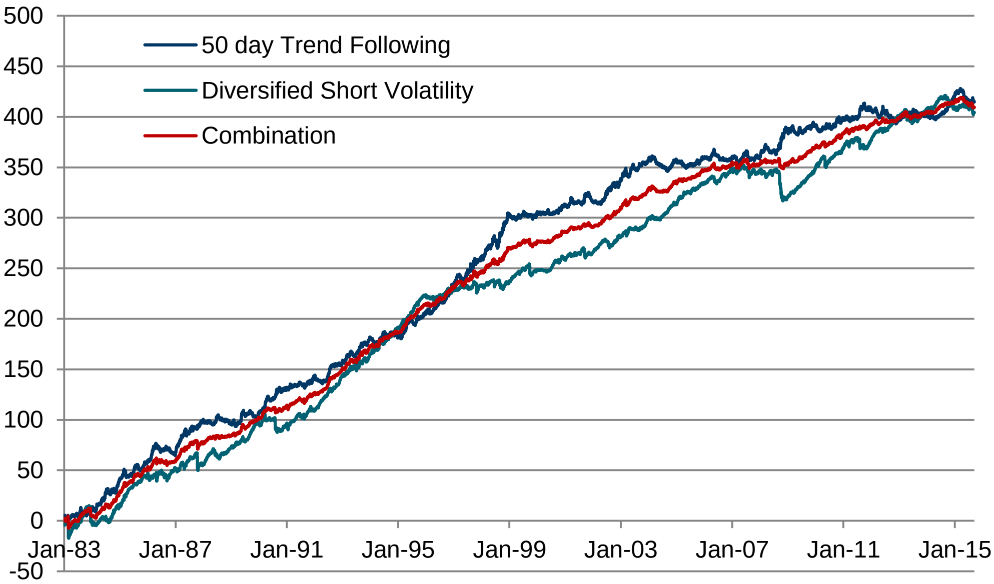

Figure 4: The performance of a diversified 50 day trend following strategy, a diversified short volatility strategy and an equally weighted combination of both. The improvement in Sharpe and in drawdown depth is clear in the combination.

CONCLUSIONS

Risk premium is compensation for assuming downside risk, and negative skewness seems to be a better explainer of the premium than volatility. If anything, there seems to be a premium for holding low volatility instruments and this is the basis of the much discussed "low volatility" effect. Negative skewness can be reduced in a portfolio by combining many de-coskewed strategies, and the skewness profile of risk premium investing improved further in combination with a trend following strategy. We believe that this approach is what all good "Alternative Beta" investment managers should strive to achieve.

REFERENCES

[1] *Errors on Statistical Measurements in Finance*. P. Seager, forthcoming.

[2] *Two Centuries of Trend Following*. Y. Lempérière, C. Deremble, P. Seager, M. Potters, J.P. Bouchaud. 2014, Journal of Investment Strategies

[3] *Tail Risk Premiums versus Pure Alpha*. Y. Lempérière, C. Deremble, T. T. Nguyen, P. Seager, M. Potters, J.P. Bouchaud. 2015, Risk Magazine.

[4] *Deconstructing the Low-Vol Anomaly*. S. Ciliberti, Y. Lempérière, A. Beveratos, G. Simon, L. Laloux, M. Potters, J. P. Bouchaud. 2015.

The Misleading Nature of Correlations

In this note we explain certain subtle features of calculating correlations between time-series. Correlation is a measure of linear co-movement, to be contrasted with the quadratic nature of risk. This can lead to misleading impressions arising from correlating two time-series. We show that the correlation of a manager with a benchmark leads to an estimate of the square root of how much exposure the manager has to the benchmark. We also show that an estimate of correlation with monthly data over 5 years has an associated error of 0.13, and therefore only a correlation of greater than 0.26 should be considered significantly greater than zero.

INTRODUCTION

When comparing two return streams, investors generally calculate correlation coefficients to identify decorrelating and diversifying investments. Correlation calculations are ubiquitous enough to be included in any reasonable time-series analysis software package and are therefore often used blindly.

In discussing correlations we will also introduce the notion of exposure. For example if we combine two independent strategies, x and y, to give a combination $sum = x + y$, then the proportion of risk taken by x is represented by its β with the total[9] and that of y is similarly represented by its β with the total. The details of why exposure is defined in this way are described in the appendix.

In this note we will model real world return streams through the use of simple "random walks" to illustrate a few counterintuitive results. A further appendix with a comprehensive derivation of the results is available upon request for the more mathematically inclined reader.

NUMERICAL SIMULATIONS - A PRAGMATIC APPROACH

In this section we will illustrate the power of using numerical methods to answer questions concerning the correlations between time-series. The following may be considered technical by some readers; it may be safely skipped in order to get to the key results. We first begin by introducing the basic tool of these simulations - the random walk. In order to keep things as simple as possible we will only study time-series with constant levels of risk and Sharpe ratio.

With this in mind, the simplest random walk for a price p can be written as follows:

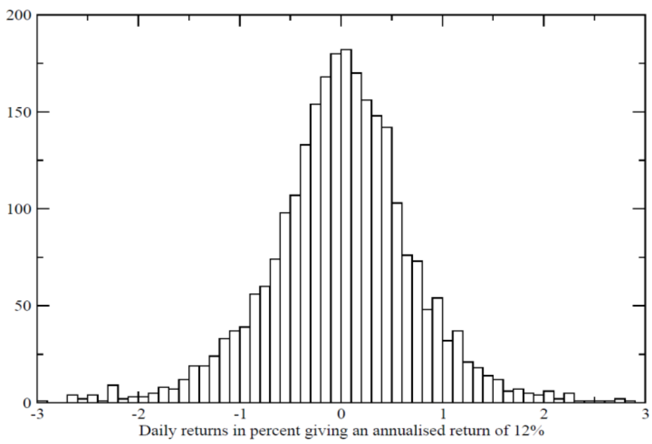

Figure 5: *A histogram illustrating the bell shaped distribution of the random numbers used in the random walks. The random numbers are centered on zero and have tails that fit financial time-series well.*

[9] The β of variables a with respect to b is defined as Covar$(a,b)/$Var(b) where Covar is the covariance between two variables Covar$(a,b) = \frac{1}{N}\sum_{n=0}^{N} a_n b_n$ while Var is simply the variance of a variable, more commonly known as the square of the standard deviation.

$$p_n = \sum_{n=0}^{N} (d + \eta_n)$$

where n is the counter, say the days for a daily return and N is the total number of days in the time-series of returns. The η term is simply a zero mean noise term or random number generator with a bell shaped distribution that best models the returns of the investment strategy. A histogram of these random numbers can be seen in **Figure 5** showing a distribution centred on zero with tails representative of financial returns[10]. The d term is a constant added to the unpredictable "noise" η_n at every time step to generate a random walk with a "drift," or positive return. **Figure 6** shows the results of generating random walks with Sharpe ratios of 0, 0.5 and 1 by varying the drift term to achieve the Sharpe ratio we require. Obviously, a Sharpe ratio of zero is generated by applying no drift term at all $i.e.$ setting d to zero and allowing the zero mean of the η_n random numbers to generate a flat (on average) random walk with a Sharpe ratio of zero.

We now have a framework within which to simulate many random walks with any particularly desired Sharpe ratio, each realisation being different due to the existence of the η_n term. The time-series in **Figure 6** shows how these random walks resemble different return streams, such as investment indices or individual funds.

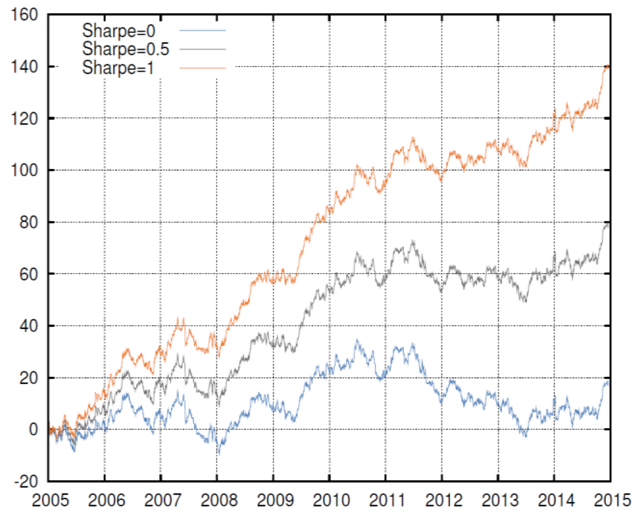

Figure 6: *Random walks generated with three Sharpe ratios, illustrating how varying the d parameter allows us to easily change the drift and hence the Sharpe ratio.*

CORRELATING TWO UNCORRELATED RANDOM WALKS WITH THEIR SUM

Let us imagine we have two time-series which are zero correlated, representing two different funds. These two time-series are shown in **Figure 7**. We have added a drift to get Sharpe ratios of 1 for each, and can now sum the two together. There is perhaps no surprise that the Sharpe ratio increases, showing the benefit of diversification, but let's now try to calculate the correlation of one of the strategies with the total. Intuitively one might expect that the correlation would be 50% due to the fact that we have 50% of each strategy in the time-series. In fact, the correlation turns out to be 71%! Correlating the sum of two time-series with either of the two strategies used in the sum gives us a higher correlation than the weight of the strategy within the mix. This could be considered a counterintuitive result. We will now show that correlation is always higher than exposure.

[10] The choice of the distribution of returns can change the results of the study. Here we use a Student's distribution with 4 degrees of freedom, a distribution which is naturally "fat tailed" and fits financial time-series well. For the purpose of this short note, however, we will neglect the effects of these fat tails on the calculation of correlations. One could use the commonly known Gaussian distribution to achieve very similar results.

CORRELATION TO EVALUATE A MANAGER'S EXPOSURE TO A BENCHMARK

We now turn our attention to another example. This time we have a manager with a small exposure to a well-known benchmark strategy, such as trend following, equity momentum, carry, value etc., but claiming he has decorrelated strategies running in parallel that make up the bulk of the risk of his returns. In order to estimate the contribution of a manager's return arising from a standard factor, an analyst may choose to correlate the benchmark or factor with the manager's returns. We can now use the example of the previous section (correlating the sum of two random walks with one of the two components) to illustrate how this can yield misleading results. We now allocate a proportion f of the benchmark strategy to the manager and combine it with $(1 - f)$ of the uncorrelated non-benchmark strategy that the manager claims to be employing. Here we have a potential source of confusion as f does not reflect exposure, but it is instead the β of the strategy with respect to the total that is a true indicator of the risk taken by the strategy in the combination (please refer to the appendix for more detail on this point). We now have two time-series to correlate: the manager's returns $r_{man} = fr_{BM} + (1 - f)r_{NBM}$ and the benchmark strategy r_{BM}, where r_{BM}, and r_{NBM} represent the returns for the benchmark and for the manager's decorrelated non-benchmark return streams respectively.

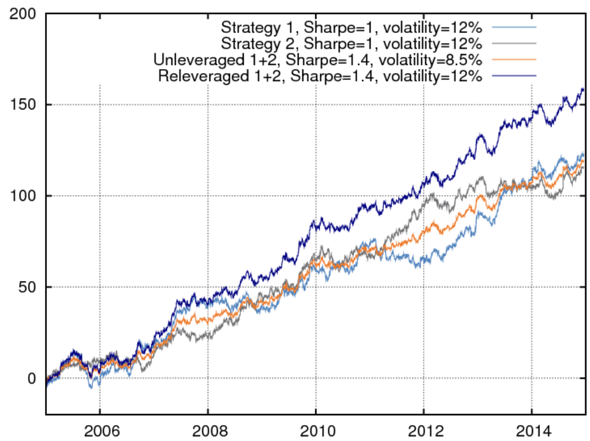

Figure 7: *Two strategies, each with a Sharpe of 1, added together to illustrate the power of diversification. We first add each strategy with a weight of one half, thus obtaining the same level of drift but a lower volatility. We then leverage the volatility to be the same level as the two inputs, thus demonstrating that we reach a higher overall gain over the period. Correlating either of the two initial strategies with the sum gives a correlation of 71% rather than 50%, as naively expected.*

Let's begin with the case of $f = 0.5$, which reproduces the result of the previous section, meaning a manager who has 50% of his risk allocated to a benchmark and 50% allocated to a non-benchmark strategy will correlate 71% with that benchmark strategy. Let's now try varying the weight f and observe how the correlation varies and, more interestingly, how the risk exposure to the benchmark varies. Because of the fact that risk sums quadratically, exposure to the benchmark strategy does not scale linearly with f (please see appendix for details). In **Figure 8** we plot the variation of the correlation *and* exposure as a function of f. One can see that the correlation does not follow the exposure, as stated, but is consistently above it. Correlation is, in fact, the square root of exposure. If we come back to the example of a 50/50 split between strategies giving a 71% correlation with the total, one can now observe that in fact the exposure of r_{man} to r_{BM} is $0.71^2 = 0.5$ which seems indeed logical. It suffices, therefore, in such situations to consider the square of the correlation as the best estimate of exposure to a particular strategy within a combination rather than just the correlation itself. We have shown this result empirically here but it can also be derived mathematically. Interested readers are invited to contact us for further details of the derivation.

THE UNCERTAINTY ON THE MEASUREMENT OF CORRELATION

Let us now turn our attention to the problem of the significance of a measurement. For correlations close to zero, the error on the measurement goes as $\sim 1/\sqrt{N}$ where N is the number of points used in the estimate[11]. If we assume that we are correlating managers with benchmarks using ~5 years of monthly data, then the error on the estimate is accordingly $\sim 1/\sqrt{12 \times 5} = 1/\sqrt{60} \sim 0.13$. Using daily data gives a far more significant result due to the fact that ~ 20 times more data is used in the estimate (as is the case in the analysis above). One needs to be careful in estimating correlations with monthly data where for a sample size of ~5 years, a correlation of 0.26 cannot (and should not) be considered positive (or negative!) with an acceptable level of significance.

CONCLUSIONS

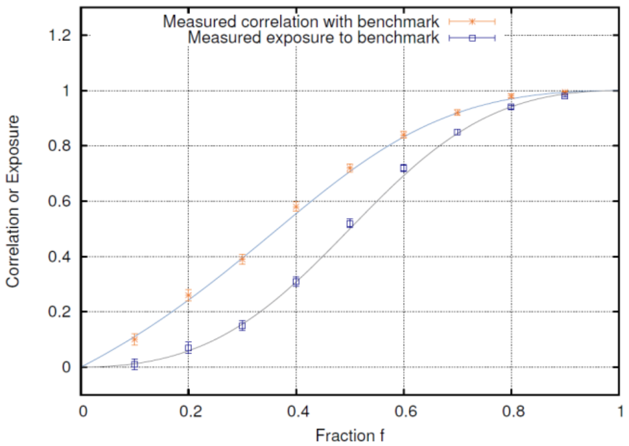

***Figure 8**: The plot shows the effect of varying the weight of the benchmark strategy that the manager is running (x-axis) against the corresponding correlation that the combination has with the benchmark and the exposure the combination has to the benchmark (y-axis). The parameter f is simply the weight allocated to the benchmark, not the proportion of risk in the combination. This "exposure" is being encapsulated in the β (see text and appendix). Correlation is not the same as exposure, the two being related such that exposure is equal to the square of the correlation. The lines through the points are the result of an analytical solution to the problem, the details of which are available upon request.*

When comparing a manager with a benchmark, correlation is not a good direct indicator of the exposure that the manager has to the benchmark. The square of the correlation is actually an estimate of the exposure the manager has to the benchmark, which can be very different to the correlation itself. One should also be aware of the fact that any correlation, especially using monthly data needs to be considered along with its statistical error. Using 5 years of monthly data means that one needs correlations of greater than 0.26 to be considered statistically significantly different to zero.

[11] The error is actually $\frac{1-\rho^2}{\sqrt{N}}$ for non-zero values of ρ

Appendix

The correlation ρ generally takes the form:

1. $$\rho = \frac{Covar(r_x, r_y)}{\sigma_x \, \sigma_y}$$

where x and y represent two investments, r is the return and σ is the standard deviation or volatility of the investment. The "Covar" represents covariance which is an averaged quantity over the two variables:

2. $$Covar(r_x, r_y) = 1/N \sum_{n=0}^{N} (r_x(n) - \overline{r_x})(r_y(n) - \overline{r_y})$$

where n is the counter, say the days for a daily return and N is the total number of days in the time-series of returns. $\bar{r}_x$ and $\bar{r}_y$ are the means of the returns r_x and r_y respectively.

For the purpose of the note, we would also like to introduce a couple of further ideas concerning risk. Let us take the example of two uncorrelated strategies x and y that are combined together to give a strategy $sum = x + y$. Let us now employ a weighting scheme allocating say 0.2 to x and 0.8 to y and write the risk of the strategy sum as follows:

3. $$\sigma_{sum}^2 = 0.2^2\sigma_x^2 + 0.8^2\sigma_y^2 + 2\rho 0.2\sigma_x 0.8\sigma_y$$

Note that ρ is the correlation between the strategies that we have set at zero and therefore we can neglect the last term. What one can see immediately is that with a portfolio weight of 0.2 we do not have a contribution to the total risk of 20%. In fact the contribution, because it is squared is much smaller at $0.2^2 = 0.04$ compared to the weighting of the strategy y which takes a weight of $0.8^2 = 0.64$ in the calculation of the variance[12]. If we set the volatility of each to be $\sigma_x = \sigma_y$ in order to simplify further then on a stand-alone basis with such a weighting the risk of strategy x is indeed 0.2 and that of strategy y is 0.8. However, it is clear that in the combination the contribution to the risk of strategy x is much smaller than 20%. So, the question one should ask is how much of the risk is being explained by x and y in the total?

Let's now therefore try to find a definition for exposure. Consider instead a regression analysis similar to that used in the CAPM, regressing the returns r_x of the strategy x and the returns r_y of strategy y on the returns r_{sum} of the combined strategy. We can define a weight f as the allocation weight to strategy x (this corresponds to the 0.2 in the above example) and a weight $1 - f$ to the strategy y (equally this corresponds to the 0.8 in the above example), such that the sum of these portfolio weights is equal to 1. As above however, we note that the portfolio weights are not a measure of the contribution of risk of x and y in the sum. Now, analagously to the CAPM we regress the strategy returns on $r_{sum} = fr_x + (1 - f)r_y$ to give:

4. $$fr_x = \beta_x r_{sum} + \text{"unexplained"}$$

5. $$(1 - f)r_y = \beta_y r_{sum} + \text{"unexplained"}$$

[12] variance is just volatility squared

Given that f is known and that sum is known to only be composed of r_x and r_y then the unexplained part of the regression collapses to zero and we can explain sum fully with x and y. The βs also conveniently add linearly to 1 and are proportional to the amount of risk carried by the corresponding strategy as a fraction of the total. In such a case the βs are calculated as follows:

6. $$\beta_x = \frac{Covar(f r_x, r_{sum})}{Var(r_{sum})}$$

7. $$\beta_y = \frac{Covar((1-f) r_x, r_{sum})}{Var(r_{sum})}$$

and assuming again that $\sigma_x = \sigma_y = 1$ the βs can be written as:

8. $$\beta_x = \frac{f^2}{f^2 + (1-f)^2}$$

9. $$\beta_y = \frac{[1-f]^2}{f^2 + (1-f)^2}$$

Interested readers are again asked to contact us for the derivation of the above analytical solutions which are used in figure 4 in the paper. So, "Covar" is covariance as previously defined and the unexplained noise terms in the regression are fully explained by the fit. Let us now once more draw on the analogy with the CAPM in terms of interpreting the βs, looking at equations (4) and (5) one sees clearly that on average $f r_x = \beta_x r_{sum}$ and (1-f)$r_y = \beta_y r_{sum}$ meaning that for a given return of the sum r_{sum}, the amount by which the strategies x and y move relative to sum is encapsulated by the βs. Since the two βs sum to one, they represent how much of the risk (or strictly speaking the variance) is being explained by each strategy in the total. As mentioned previously, this relates back to the CAPM. When we regress stock returns r_i over the index returns r_I we obtain the following

10. $$r_i = \beta_i r_I + \text{"unexplained"}$$

and β is interpreted as the exposure one gets to the index by holding the stock. If the β is 2 and the index goes up by 1%, we would earn 2% on average from our investment. Here we are doing much the same thing except the index is now the sum of two investments rather than an arbitrarily weighted index and the returns we are regressing are those of the strategies we used to construct the sum.

Deconstructing the Low-Vol Anomaly

Alexios Beveratos, Jean-Philippe Bouchaud, Stefano Ciliberti, Laurent Laloux, Yves Lempérière, Marc Potters, Guillaume Simon

The low-volatility anomaly was observed and documented at least as early as 1970 by Fisher Black – who failed to convince Wells Fargo to launch a levered fund that would buy low-volatility and sell high volatility stocks – and in 1972 by Robert Haugen – who equally failed to have his paper published before his results contradicting the CAPM model were expunged from the record. That low-volatility stocks should perform better than their high-volatility counterparts is indeed counter-intuitive, and in blatant contradiction with the idea, deeply rooted in economic theory, that risk should be somehow rewarded by some excess return. Still, concurrent empirical evidence has accumulated since the early seventies, and broadly confirm that this low-volatility "puzzle" is a robust, universal stylized fact of stock markets. The effect has indeed been persistent over time, and is documented on a variety of stock markets throughout the world (developed countries or emerging markets alike). Whereas all studies confirm that the low-vol anomaly is strong and pervasive in stock markets, the origin of the effect is still debated. We have conducted our own investigation of this effect and come to the following conclusions:

1. We find that a large proportion of the low-vol performance is in fact eked out from dividends. This is our central result that follows from the strong negative correlation between volatility and dividend yields which, oddly, does not seem to be clearly documented in the literature. However, the low-vol anomaly persists for ex-dividend returns. These returns are found to be roughly independent of the volatility level, i.e. risk-adjusted ex-dividend returns are higher for low-vol stocks, which is in itself an "anomaly".
2. We find that the skewness of low-vol portfolios is small but systematically positive, indicating that the low-vol excess returns cannot be identified with a hidden risk premium. This ties up with intuition: it would be hard to imagine that shorting high-vol stocks and holding safe stocks is a risky strategy.
3. A decile analysis reveals that the anomaly progressively builds up slowly from high- vol to low-vol stocks. The effect is therefore genuine and not concentrated on stocks undergoing extreme movements.

Our overall practical conclusion is that, while the low-vol effect is indeed compelling in equity markets, it is not a real diversifier in a factor driven portfolio that already has exposure to Value type strategies. In a nutshell, a dividend yield factor explains (as expected) the dividend part of the low-vol performance, while an Earning-to-Price factor explains its ex-dividend part.

Finally, the underlying reasons for the low-vol anomaly to persist in equity markets are still, by and large, obscure. Although the behavioural/institutional stories that have been put forth are persuasive and compatible with the bias observed in the holdings of mutual funds, there is no empirical smoking gun. We tend to believe in a universal "lottery ticket" or embedded option mechanism that affects institutional investors and private investors alike, leading them to over-focus on potential spectacular upsides and forget much smaller but significant regular dividends.

A full copy of the paper is available here.

Contacts

Paris
Capital Fund Management S.A.
23, rue de l'Université
75007 Paris, France
Tel: +33 1 49 49 59 49
Fax: +33 1 47 70 17 40

London
Capital Fund Management LLP
Queens House 6F
64 St James's Street
London SW1A 1NF, United Kingdom
Tel: +44 207 659 9755
Fax: +44 208 001 4684

New York
CFM International Inc.
The Chrysler Building
405 Lexington Avenue - 55th Floor
New York, NY, 10174, USA
Tel: +1 646 957 8018
Fax: +1 646 957 8019

Tokyo
CFM Asia K.K.
9F Marunouchi Building
2-4-1, Marunouchi, Chiyoda-Ku
100-6309 Tokyo, Japan
Tel: +81 3 5219 6180
Fax: +81 3 5219 6185

Disclaimer

ANY DESCRIPTION OR INFORMATION INVOLVING INVESTMENT PROCESS OR ALLOCATIONS IS PROVIDED FOR ILLUSTRATIONS PURPOSES ONLY.

ANY STATEMENTS REGARDING CORRELATIONS OR MODES OR OTHER SIMILAR STATEMENTS CONSTITUTE ONLY SUBJECTIVE VIEWS, ARE BASED UPON EXPECTATIONS OR BELIEFS, SHOULD NOT BE RELIED ON, ARE SUBJECT TO CHANGE DUE TO A VARIETY OF FACTORS, INCLUDING FLUCTUATING MARKET CONDITIONS, AND INVOLVE INHERENT RISKS AND UNCERTAINTIES, BOTH GENERAL AND SPECIFIC, MANY OF WHICH CANNOT BE PREDICTED OR QUANTIFIED AND ARE BEYOND CFM'S CONTROL. FUTURE EVIDENCE AND ACTUAL RESULTS COULD DIFFER MATERIALLY FROM THOSE SET FORTH, CONTEMPLATED BY OR UNDERLYING THESE STATEMENTS.