INSIGHT.DATA.CLARITY.

November 2019

OF PRESIDENTS AND HEART ATTACKS

Risk control as diversification through time

Executive summary

In this short note we briefly introduce "risk" (for those readers unfamiliar with the concept and how it is defined); show conceptually how we are able to forecast and control it (for those concerned about risk in their portfolio and how it may be mitigated); and highlight, through a set of historic events, how even a risk-controlled portfolio remains at the mercy of idiosyncratic events.

Contact details

Call us +33 1 49 49 59 49

Email us cfm@cfm.fr

Introduction

In common parlance, risk is simply the possibility of loss. Losses may stem, predominantly, from two sources: uncertainty about the *1) direction* of expected returns, and the *2) magnitude* of returns. While the former (alpha) is difficult to harvest and maintain, the latter (volatility) displays certain features that facilitate the forecasting, and thus controlling of the magnitude of investment moves.[1]

The question of risk and volatility has occupied some of the most celebrated minds in finance and economics for the better part of 70 years. Most of the seminal work still cited today reads like a who's who of Nobel Prize winning economists: Markowitz, Sharpe, and Engle to name but a few, all toiled to understand the nature of risk, and sought appropriate models to measure and forecast it.

The use of volatility as a short-hand for risk, notwithstanding its now ubiquitous acceptance, provokes much criticism.[2] Still, as volatility is easier to quantify and exhibits features we can leverage to adjust exposure to further hikes in market stress, it can readily be used to design a risk-controlled portfolio.

Therefore choosing and assigning a quantifiable proxy for the size of market moves is important to control volatility, and deliver a more predictable return profile. This note will focus on volatility and why the adoption of a systematic, volatility-controlled protocol is desirable, and even necessary.

What is risk? And how is it commonly measured?

Risk has become synonymous with volatility, and is most commonly measured as the standard deviation,[3] i.e. the degree of deviation from the mean of a price return series over a given period, ordinarily annualised and printed in percentage.[4]

Equating risk with volatility is not a new idea. Already in his 1952 seminal paper, Harry Markowitz, the father of Modern Portfolio Theory, associated risk with *variance*[5] in the value of a portfolio, stating that one should consider "variance of return an *undesirable* thing". Markowitz also discussed risk in its relation to the correlation of assets, namely that combining two correlated assets intuitively increases risk relative to the combination of two anti-correlated assets. Sitting in between these two extremes, the combination of uncorrelated assets reduces risk more than returns and thus provides an improvement in risk adjusted returns. This observation is at the heart of what Harry Markowitz described as "the only free lunch in finance".

Some features and characteristics of volatility

In finance jargon, it is well understood that volatility is said to exhibit '*autocorrelation*',[6] or '*clustering*' – that is to say that high (low) volatility in the past, is likely to be followed by high (low) volatility in the future.[7] It is thus is a measurement of the relationship between a time series, and a lagged version of itself.

Measured between -1 (perfectly negatively correlated) and +1 (perfectly positively correlated), it measures by how much volatility levels persist.[8] To illustrate this persistent feature of volatility, in Figure 1 we plot the monthly realised volatility of the Dow Jones Industrial Average, and show how a month of high (low) volatility is followed by a month with high (low) volatility.

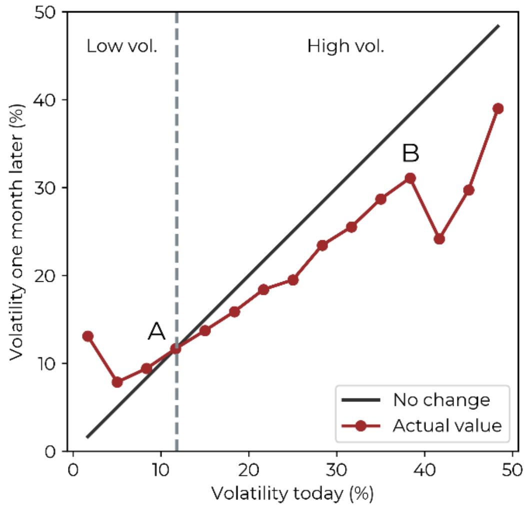

Figure 1: We plot the monthly volatility calculated as the annualised daily realised volatility using a 22-day (1-month)

[1] There are many other kinds of 'risk' that can manifest over the holding period of any investment: one need simply flip through any 'Key Investor Information Document' (KIID) or 'Product Disclosure Statement' (PDS) to take note of the myriad of various risks to which a portfolio is exposed: drawdown, currency, tax, liquidity, counterparty, etc. to name but a few.

[2] Perhaps most famously from Warren Buffet who berated the use of volatility as a risk measure in his 2014 letter to shareholders, calling it, amongst others, "*far* from synonymous with risk": https://www.berkshirehathaway.com/letters/2014ltr.pdf

[3] Standard deviation is a statistical measure of the dispersion of a set of numbers around an average. See for instance: https://en.wikipedia.org/wiki/Standard_deviation

[4] A criticism levelled at the typical calculation of volatility, is the agnostic treatment of positive and negative price changes (since both the positive and negative returns are squared). However, when one wants to forecast risk, this can be remedied by, for instance, constraining a risk estimation model by only taking negative price changes into account – often called 'downside risk measures'.

[5] Variance is the squared difference between an observation and the mean of all the observations in the same time series, the standard deviation (volatility) is the square root of the variance.

[6] Also sometimes referred to as 'serial correlation'.

[7] This feature was first documented by Benoit Mandelbrot in his 1963 paper, 'The Variation of Certain Speculative Prices'. The Journal of Business, 1963, vol. 36, 394.

[8] Please see our discussion note "Is there a 'new normal' in Volatility Markets?... Probably not!" in which we illustrate the fluctuating, yet consistent autocorrelation characteristic of both implied and realised volatility in equity markets. The paper is available on our website.

standard deviation moving average. High (low) volatility in a preceding month ('volatility today' on the x-axis), is typically followed by high (low) volatility in the succeeding month ('volatility one month later' on the y-axis). Point B for example shows that a month with volatility of ~38%, is followed by a month where the monthly volatility is ~30%. Point A is the average, annualised monthly realised volatility of the Dow since 1910: 11%.

Autocorrelation of volatility is observable, and measurable at both security, and portfolio level. However, the volatility of a single security is habitually higher than that of a portfolio with multiple (uncorrelated) securities. This is as a result of diversification (due to the way risk combines among uncorrelated instruments as mentioned above). The addition of securities in a portfolio thus reduces the overall volatility of the portfolio as illustrated in Figure 2.

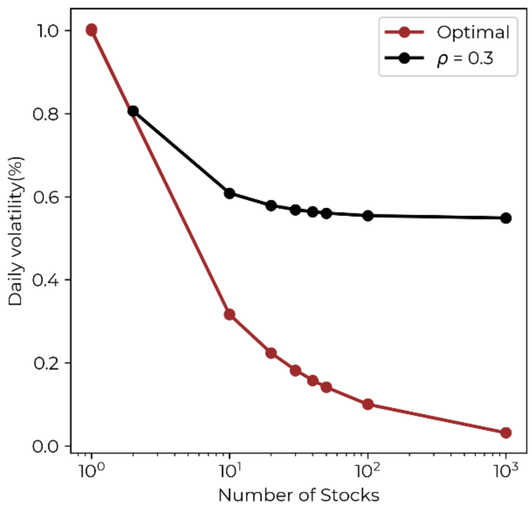

Figure 2: *A representation of how the volatility (y-axis) of a portfolio decreases with the addition of an increasing number of securities in a portfolio (x-axis in log-scale). We assume zero correlation between stocks and normally distributed returns (the red curve marked optimal). For the sake of simplicity, each stock is assumed to have a daily volatility of 1%, with the addition of each new stock also assumed to have the same 1% daily volatility. The volatility of a portfolio of N assets – given the assumption that the securities are uncorrelated – decreases at the rate of $\frac{1}{\sqrt{N}}$ If, however, as is the case in financial markets, securities exhibit varying degrees of correlation among themselves, the rate of decrease in volatility is slower. This is indicated by the black curve, marked ρ = 0.3, i.e. assuming a ~30% average correlation between securities in the market. Correlations between stocks exhibit such an effect and therefore stock market indices quickly no longer feel the benefits of diversification beyond a given threshold of N.*

Another feature of volatility is its tendency to spike unexpectedly (see for instance Figure 3). Whilst some spikes in volatility are easily attributable to an idiosyncratic or exceptional event, it is often less evident why volatility in the market increases so suddenly. Higher volatility in equities is also most commonly associated with negative return shocks, that is, a *negative* and *asymmetric* relationship exists between volatility and returns – the so called *Leverage Effect*. This ties in with the negative *skewness* observed in equity markets. Whilst interested readers can refer to a more technical explanation of skewness and different moments of the return distribution, for the purposes of this note, it suffices to say stock returns are characterised by a mix of regular, but small positive returns, dotted with less frequent, but large negative returns.

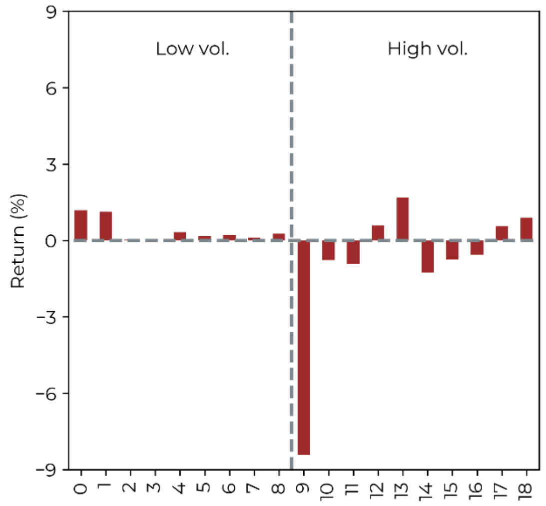

Figure 3: The daily price returns for Coca-Cola (constituent of the Dow Jones Industrial Average Index) during February 2019. Volatility was muted for the first half of the month, but a sell-off on February 14 (day 9 in the plot) following a warning issued by the company that sales growth was likely to slow in 2019, lead to a -8.4% drop in the stock price. Although this is one anecdotal example of a pronounced spike in volatility, followed by a period of higher volatility, it is a common observation in equity markets.

Why does one control for risk?

Prospect theory shows that given a set of choices, investors prefer less risk – they are for the most part said to be risk-averse.[9] Investors are as such inclined to select more certainty about the outcome of an investment, even if it entails lower returns. That is to say, they prefer returns with fewer severe negative returns (fat-left-tailed returns).

Since investors are behaviourally more risk-averse, increased volatility and drawdowns in the short-term can prove very trying for some. This often leads to irrational – or

[9] We have shown in prior work that in reality investors are rather loss-averse and actually quite like to experience large and sharp gains in a P&L, a result that explains the premium received by insurance sellers. Curious readers are directed to our discussion note "Risk Premium Investing. A tale of two tails" available on the CFM website.

even imprudent – investment decisions. Some might, in addition, require certainty of returns as and when financial risk thresholds are breached. Others might also require a limit in downside risk.

In order to address these behavioural traits it is natural to want to control the risk of an investment or a portfolio of investments. If risk can fluctuate, then inevitably there are periods of high volatility. One may be aware that these periods will exist, but living through them is a different matter! If these periods of high risk also coincide with periods of negative performance then one is in the realm of the most difficult to bear returns: those with the dreaded fat-left-tail! Risk control thus is a remedy for reducing fat tails but also, in general, an improver of risk-adjusted returns, once the biggest moves are smoothed out.

How does one control for risk?

As any Finance 101 student will instinctively tell you, diversification. It is one of the key tenets of finance and allows investors to avoid what is commonly termed 'unsystematic risk', or diversifiable risk - those risks that are not commonly shared across all industries or asset classes.

As investors spread risk exposure over various asset classes (or within a given asset class[10]), the risk of the portfolio is diluted.

Cross sectional diversification, while effective in reducing risk, is not sufficient to protect against fluctuations *across time* – see Figure 4. And while diversification across sectors is a good hedge, it moreover does not mitigate against systemic risk – those risks that affect multiple sectors. Typical examples are political instability, or geopolitical events which trigger volatility in all corners of the market.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
Dow Jones Index	**7.2**	**3.7**	**0.1**	**2.6**	**-6.7**	**7.2**	**-1.7**	**2.0**	**0.5**
Communication Services	0.3	2.9	0.6	9.3	-3.5	5.4	-0.7	1.0	-0.6
Consumer Discretionary	6.2	3.2	1.0	4.0	-4.2	7.7	1.5	4.0	3.8
Consumer Staples	4.3	-0.3	0.9	-0.7	-2.2	7.5	2.1	2.0	3.6
Energy	4.4	2.6	1.0	-2.8	-6.7	8.8	-2.1	-4.0	1.9
Financials	8.4	2.5	-0.7	6.2	-4.8	7.0	1.2	-2.1	2.0
Health Care	2.9	-0.2	0.3	-3.5	-0.2	4.3	-4.3	-2.8	-2.0
Industrials	9.6	7.1	-3.1	-1.9	-11.9	8.0	-1.5	-1.3	4.3
Information Technology	6.2	6.6	4.2	2.1	-8.4	8.6	3.5	-4.8	4.1

Figure 4: *Individual securities (or sectors at an aggregated level) exhibit different levels of return, and have historically displayed different sensitivities to, for example, economic or business cycles. Here we show a map of the average returns per sector of the Dow Jones, per month for 2019 up until the end of Q3. The squares in green show the top two performing sectors for each month, with the bottom two performing sectors for each month in red. There is clear, inconsistent performance across sectors over time – making the argument for sector diversification. However, just diversifying across the rows (sectors) is not sufficient. The resulting reduction in volatility may improve portfolio performance but will not be enough to avoid periodic exposure to heightened levels of market volatility (when correlations between securities increase, and where diversification across the columns (time) is required).*

While many investors might even be unaware, when they allocate exposure, for instance, between developed and emerging markets, or equities and bonds, they are in effect taking correlation differentials into account. As explained earlier, during periods of high market stress – especially owing to systemic risk, the correlation between financial instruments within the same asset class tends to increase, and move towards 1. The benefits of diversification subsequently get eroded, since most of the price movements of the individual securities are in the same direction.

However, one may mitigate any expected market volatility by scaling down one's position to avoid periods of heightened volatility. If one can forecast future risk, one can scale positions, investing more (less) when volatility expectations are low (high). This rescaling of positions is, in practice, only possible in a world of partially financed futures. Real money positioning in any asset leads to a fixed notional sizing which is naturally capped (at the level of assets of the investment) and sometimes not easily adjusted due to illiquidity. Such effects are avoided

[10] Investors might choose to diversify their risk across various sectors in an equity portfolio for example, or, distribute their risk exposure across regions between developed and emerging markets.

through the use of futures[11] and has in more recent years given rise to the Risk Parity industry.

Can one forecast risk?

As any market participant will attest, forecasting any *direction* in financial time series is difficult. 'Forecasting' volatility, however, is easier.

This is because of the clustering feature of volatility discussed earlier: higher or lower volatility periods tend to persist. Luckily, for the purpose of risk control we are not so much interested in the *direction* of change, but rather the *magnitude* (and, being agnostic as to the sign of the change).

There are various models that can be employed to estimate risk, perhaps none more commonly used than the family of 'ARCH' models (see the appendix for a brief explainer). One can also devise an expectation of future volatility from option prices. This is commonly referred to as 'implied' volatility (implied from the prices of options).[12]

A test case – controlling the risk of the Dow Jones Industrial Average

Since we can forecast volatility (albeit imperfectly), we can distribute risk more evenly through time by exposing a portfolio less when moves are large (or expected to be large), and more when moves are small (or expected to be small).

While a myriad of techniques to control for risk exist[13], we will demonstrate the results by using a rolling standard deviation technique:

$$\sigma_t = \sqrt{std\left(\{r_{t'}^2\}_{t'=t-22D}^{t'}\right)}$$

Where σ is volatility, r corresponds to returns and t is time measured in days.

By scaling one's position based on realised volatility $(pos_t \propto \frac{1}{\sigma})$, we can improve the Sharpe ratio of being long the Dow Jones by ~25% since 1910. In Figure 5, we plot the cumulated performance of the Dow Jones with and without adjusting for fluctuating risk.

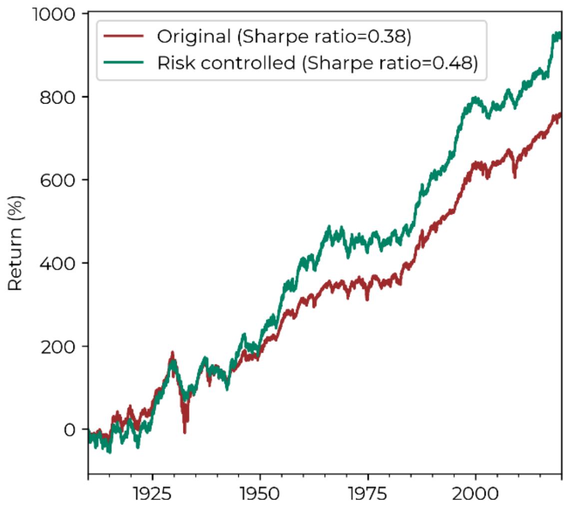

Figure 5: *The original, and risk-controlled total cumulative returns of the Dow Jones. By employing the simple rolling standard deviation to estimate, and consequently scale the exposure to achieve the same long term risk, the Sharpe ratio is increased by ~25% over the more than 100-year sample period.*

Another benefit of applying a risk control protocol is the mitigation of extreme returns in the tails of the return distribution. In Figure 6 below, the original (long only, no risk control) and risk controlled price change distribution is plotted and one observes a more normally distributed return stream after the risk control is applied.

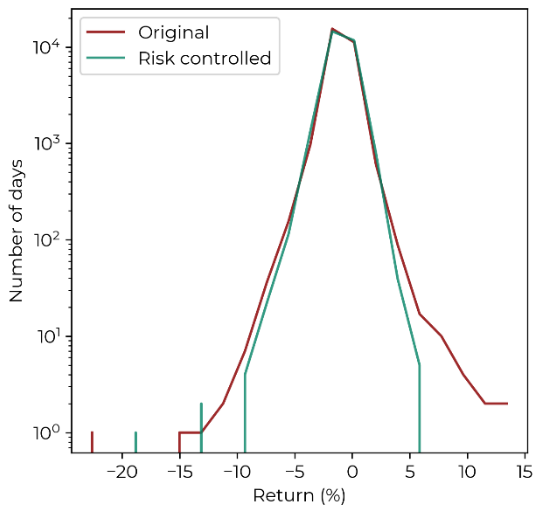

Figure 6: *The return distribution of both the original, as well as the risk-controlled index. The distribution of the risk-controlled returns shows a more normal and less fat-tailed return distribution. The frequency (Number of Days) is in log scale.*

Finally, as is shown in Figures 7.a and 7.b, the average daily price returns as well as the daily PnL volatility respectively

[11] See Appendix for a discussion of this point.

[12] Most readers will be familiar with the VIX Index, which implies the market view of market volatility of the S&P 500 over the next 30 days from S&P 500 options.

[13] Common examples include Moving Average, Exponentially Weighted Moving Average, Historical Mean, ARCH, etc. – all with their own benefits and drawbacks.

of the risk controlled investment is more stable in time than that of the original, uncontrolled investment.

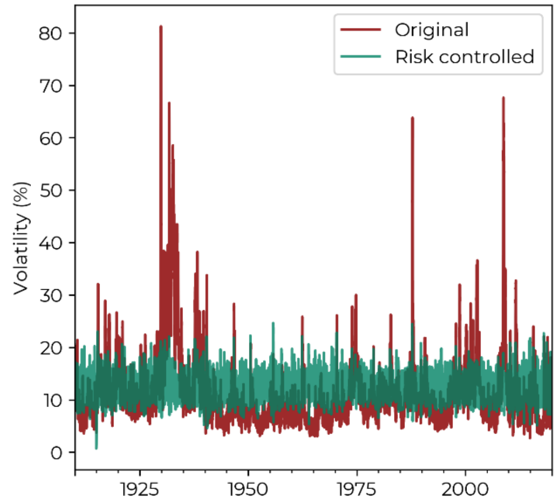

Figure 7.a: *A comparative 1-month moving average of daily volatility of the PnL (converted to annualised units) of both the original (in red) and the risk controlled (in green) investment. The PnL of the risk controlled investment is much more monotonic, while the original, uncontrolled investment features many more volatility spikes.*

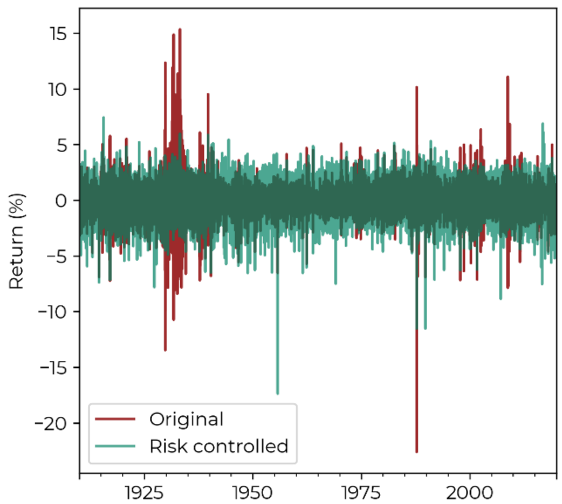

Figure 7.b: The daily price returns in percentage of both the index, and the risk-controlled version. Similar to the above, the daily price changes of the risk controlled time series features much less outsized returns – both positive and negative.

Can one perfectly control for risk?

While every effort can be made to reduce a fat-tailed return distribution, and preserve capital with disciplined systematic risk management, markets are inherently fickle.

Anticipating the idiosyncratic behaviour of markets is nigh on impossible, with the seemingly random nature of market returns well documented, if not always well understood. Proponents of the Efficient Market Hypothesis (EMH) hold that market prices reflect, at any given point in time, all the information available and relevant to any given security. If this were to be true, it should follow that only 'new' information should trigger a 'revaluation' of any given security. However, markets show too much volatility – especially in the absence of any discernible news – to be justified. Seminal work by Robert Shiller – another Nobel Prize winning economist – has shown that the volatility of stock prices is far greater than is plausibly explained by any rational view of future earnings and dividends.[14]

Another feature of financial market returns is the 'fat-tailed' distribution of returns. Fat tails simply refer to those returns that are far beyond those plausibly expected for a normal distribution. Most readers will be familiar with the 'Black Swan' analogy – the rare, *mostly* unpredictable events that trigger a sell-off and cause bursts of volatility. Accordingly, most risk control models are incapable of predicting, and as such, controlling for these events. To illustrate this phenomenon, see the below sets of tables showing the highest positive and negative returns for a risk control strategy (and the corresponding market price change) and the causes thereof.

Highest negative *risk-controlled* returns				
Date	**Price Return**	**Sigma**	**Risk Controlled Return**	**Volatility**
26/09/1955	**-6.5%**	**-18.1**	**-17.4%**	**5.7%**
13/10/1989	-6.9%	-12.0	-11.5%	9.1%
19/10/1987	-22.6%	-12.0	-11.5%	30.0%
27/02/2007	-3.3%	-9.2	-8.9%	5.7%
25/04/1927	-2.2%	-8.1	-7.8%	4.2%
Highest negative *absolute* returns				
19/10/1987	**-22.6%**	**-12.0**	**-11.5%**	**30.0%**
28/10/1929	-13.5%	-4.7	-4.5%	45.5%
05/10/1931	-10.7%	-3.1	-2.9%	55.8%
12/08/1932	-8.4%	-2.5	-2.4%	53.0%
15/10/2008	-7.9%	-1.8	-1.7%	70.7%

Table 1. We rank the highest negative risk controlled returns, and show the corresponding returns of the index (which are the

[14] Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? The American Economic Review, vol. 71, No. 3, June 1981, pp. 421-436. Our own research has yielded similar results. Interested readers can see our paper, entitled 'Stock price jumps: news and volume play a minor role' available on our website.

returns if an investor was simply holding the index). We furthermore isolate and rank only the single highest negative returns for each year, the second highest negative return for each year is ignored. The highest daily negative risk-controlled return (top table) was on 26 September 1955, where the negative risk-controlled protocol returns would have exceeded those of simply holding the index, on account of using leverage. Using the rolling standard deviation as risk estimate would have dictated higher exposure given the relatively low volatility at the time. Volatility is calculated as the 22-day rolling standard deviation.

Key events:

26 September 1955: The NYSE lost 6.6% after President Eisenhower suffered a heart attack – see Figure 8 below.

13 October 1989: A failed leverage buyout attempt for the parent company of United Airlines, owing to an inability to obtain credit, prompted traders to dump 'take-over' related and blue-chip stocks. A fear that the junk-bond market party was coming to an end (up till then fuelling much of the take-over activity) sent investors chasing for the exits. At the time, many publications hinted (and feared) similarities to the event that caused the highest negative absolute return in our example, asking, as the New York Times did: "Is It 1987 Again?"

19 October 1987: Black Monday. Probably one of the most infamous market melt-downs in post WWII history. The Dow Jones crashed 22.6% - the single largest daily percentage loss in the history of the index. While the exact catalyst is still debated, and the extent to which nascent computer-generated algorithmic trading contributed disputed, it is commonly accepted that a variety of issues contributed and escalated the sell-off. For comparison, the perhaps equally – if not more – infamous black Tuesday of 1929 (October '29) 'only' had a 11.7% drop. This sell-off was however preceded by an even bigger 12.8% decline the day before: the less famous Monday (October '28).

Highest positive *risk-controlled* returns				
Date	Price Return	Sigma	Risk Controlled Return	Volatility
16/08/1915	**3.7%**	**7.7**	**7.4%**	**7.6%**
07/11/2016	2.1%	7.1	6.9%	4.6%
15/03/1933	15.3%	6.2	6.0%	39.3%
05/09/1939	9.5%	6.1	5.9%	24.8%
12/06/1944	1.4%	5.6	5.4%	4.0%
Highest positive *absolute* returns				
15/03/1933	**15.3%**	**6.2**	**6.0%**	**39.3%**
06/10/1931	14.9%	3.6	3.5%	65.4%
30/10/1929	12.3%	2.7	2.6%	72.4%
21/09/1932	11.4%	3.7	3.5%	49.2%
13/10/2008	11.1%	3.3	3.2%	53.5%

Table 2. Contrary to Table 1 above, here we rank the highest positive risk controlled and absolute returns – again ranking only the single highest positive returns for each year.

Key events:

16 August 1915: Global markets tanked after the onset of World War I in 1914. A later boon in market prices is commonly ascribed to the opportunities that American firms were presented by selling industrial and related goods and materials to Europe, with the Dow Jones staging a strong rally right throughout 1915. Whilst there is no readily available trigger that explains the return on this day, amidst a low volatility period, with the use of leverage, the returns of the risk controlled protocol would have exceeded those of the market.[15]

15 March 1933: On March 4, 1933, freshly inaugurated President Roosevelt declared a nationwide bank holiday so that lawmakers could address the worryingly increase in bank runs. A few days later, on March 9, the US Congress passed the 'Emergency Banking Relief Act' allowing each of the twelve regional Reserve Banks in the Federal Reserve System to issue additional currency on good assets. Effectively creating a deposit insurance scheme, the law had the desired effect – on 13 March, the first day of business after the bank holiday, Americans were seen standing in line to make deposits, and markets responded favourably. After being shuttered for nearly two week, the Dow jumped 15.3% - remaining, at the time of writing, the single largest daily percentage change in the history of the index.

EISENHOWER IS IMPROVING; CHANCE OF FULL RECOVERY CALLED 'REASONABLY GOOD'

DOCTOR IS HOPEFUL

Says President May Resume Some Work in Two Weeks

Transcript of news conference is printed on Page 21.

By RUSSELL BAKER
Special to The New York Times.

DENVER, Tuesday, Sept. 27—Barring complications, President Eisenhower should be well enough within the next two weeks to resume some personal control of the nation's affairs, a noted heart specialist said yesterday.

Dr. Paul Dudley White of Boston, the specialist, examined the President yesterday morning. He told General Eisenhower that his chances for complete recovery were "reasonably good."

He cautioned, however, that the two-week period following an attack of coronary thrombosis, such as President Eisenhower suffered Saturday morning, was a critical one in which complicating setbacks might occur.

If no setback developed in this period, he said, the President should be able, for example, to confer briefly with John Foster Dulles, Secretary of State, and read official documents.

Progress Is Satisfactory

Dr. White found the President in satisfactory condition in the morning and said his morale was "wonderful." His optimistic analysis was echoed in the 4 P. M. (7 P. M. New York time) medical bulletin. That reported General Eisenhower had spent "a very comfortable day" and had eaten a lunch of meat, vegetables, fresh fruit cocktail and milk.

Just after 9 P. M., the attending physicians reported:

"This has been a comfortable day for the President. His satisfactory progress as previously reported has continued throughout the day."

Associated Press Wirephoto

REPORT ON PRESIDENT'S CONDITION: Dr. Paul Dudley White, left, Boston heart specialist, and Maj. Gen. Howard Snyder, President Eisenhower's personal physician, as they addressed press conference yesterday in Denver.

Stock Prices Off Sharply; Loss Is Put at $14 Billion

By BURTON CRANE

In a collapse ascribed to the illness of President Eisenhower, the stock market yesterday suffered its heaviest dollar loss in history. With most leaders losing 4 to 8 points a share, and one as much as 33¼, the value of New York Stock Exchange stocks lost approximately $14,000,000,000, half their 1955 gain. The entire loss for the month of October, 1929, was $16,000,000,000.

At the close of trading, prices were rising moderately on declin-

CABINET MEETING CALLED BY NIXON

Figure 8: By employing this generic risk control strategy, the highest negative relative return would have occurred on September 26, 1955 – the Monday after the weekend that US President Eisenhower suffered a heart-attack. Following news that the President suffered a heart attack, the New York Stock Exchange tumbled 6.6%. As can be seen in this newspaper clipping, the NYSE lost $14bn on the day – a loss that was triggered by an event that was perfectly unpredictable, and, given the low volatility at the time (leading to a levered position),

[15] In a real trading setup one might choose to limit the size of notional positions as a safeguard against unnaturally low volatility environments.

offered substantial losses. September 26 was, as was reported at the time, "the heaviest dollar loss in history" for the NYSE.

Conclusion

We have shown how employing a risk control protocol can increase the Sharpe ratio, reduce the fat-tailed return distribution, and thus improve on the skewness and kurtosis. By controlling for risk, we provide a 'smoother' PnL profile for clients, with risk control being thought of as 'diversification in time', in the same way that investors diversify across asset classes. We have also shown, that despite employing a risk control protocol, investors remain vulnerable to idiosyncratic moves in the markets.

Yet, forecasting volatility is still challenging. We have illustrated a simple risk-control model using a moving average of squared returns. There are, however, more sophisticated tools and techniques, such as machine learning, that can be deployed to forecast risk, but, a discussion of such techniques is beyond the scope of this text.

It is also important to note that risk control is generally assured over the timescale of the trading strategy. If one is employing a longer term strategy on the timescale of several months, say, then risk is generally controlled on that timescale. Controlling for risk on the timescale of days, in this instance, is costly (in execution costs) and futile given the slow nature of the trading strategy. On shorter timescales, therefore, one expects risk fluctuations that, as long as moves are not too extreme, are harmless in being able to provide a constant level of risk on the strategy timescale.

Risk control is an essential part of any investor's tool kit, for systematic and non-systematic managers alike. In our experience a poorly implemented risk control engine can change one's view of the utility of a given trading strategy in a portfolio. The quality of the implementation of the risk overlay is thus of utmost importance in performing research on trading strategies.

Appendix

Why the use of futures markets is crucial in delivering risk controlled returns

It seems obvious that if a stock yields say a 10% return for a 10% volatility then a $100 investment in said stock also has these same investment characteristics of 10% return and 10% volatility. So far so good! Now, what about if we instead get $100 of exposure to the stock through a futures contract? A future on the stock has a notional size, say $100 for the sake of simplicity, and earns and loses money based on the fluctuations of that notional size in the underlying stock. The investor, in theory, puts nothing down to 'buy' the future but, nonetheless, loses the $100 if the stock price collapses to zero. In reality the investor does deposit some cash in a margin account and this margin covers the potential loss for the holder of the future.

The FCM (Futures Clearing Merchant) requires this cash outlay to protect herself from the client not being able to cover potential losses. In the case of a future on a well-diversified equity index future such as the S&P 500 in the US this margin deposit would normally be of the order of 15% of the notional size of the future. In the case of our $100 future on a stock therefore we would be required to deposit 15 cents with the FCM and would be able to hold 85 cents in any investment we like – let's say we deposit it in an interest yielding bank account. We therefore have 85 cents yielding the risk free rate (the 15 cents may also be yielding the risk free rate at the FCM) plus the P&L of the stock corresponding to a $100 of investment. Because we are receiving interest on our $100 this then is removed from the P&L of the future as it is marked to market all the way to expiry (rolled futures provide excess returns rather than total returns) such that the total of $100 risk free rate plus the P&L plus the negative financing drag is equal to just holding $100 in the stock itself.

So, why go through all this detail? Because futures give us a perfect framework within which we can apply risk control to an investment. As described in the above text we would like to invest for a fluctuating notional sizing as $pos \propto \frac{1}{\sigma}$ where σ refers to our risk estimate. Because of the need to adjust our positioning we need access to an instrument that allows for a position that is sometimes higher and sometimes lower than $100 based on the most recent volatility estimate. Let's say for example that our estimate of volatility drops (increases) from 10% to 7.5% (12.5%). We then need to increase (decrease) our notional position from $100 to $125 ($75). This can be done with futures by simply adding (subtracting) notional size from the position with the addition (subtraction) of cash from the client margin account. One can safely increase notional sizing by 25% without requiring more than the $100 we actually have. This is the foundation of the Risk Parity industry that provides risk controlled exposure to bonds, equity indices, commodities etc. based on very liquid derivatives markets and provides improved levels of fat-tail control and risk adjusted returns, for the reasons discussed in this note.

A dummy's guide to GARCH

One can find many explainers of GARCH on the Internet and our intention here is not to reinvent the wheel. Instead we would like to provide some intuition behind the motivations for GARCH modelling of volatility. The classic way to write GARCH is as the following:

$$\sigma_t^2 = \gamma\sigma_0^2 + \alpha\sigma_{t-1}^2 + \beta\eta_{t-1}^2$$

For a random process where the returns, $\eta_t = \sigma_t\epsilon_t$. The ϵ here is a bell shaped random number and the index t corresponds to time, let's say days.

So, already we've probably lost you! Instead of the above equation we should consider what the GARCH model is trying to do. The main feature of volatility that we would like to model is the autocorrelation or persistence and this is what GARCH is capturing. If we set $\gamma = 0$ and $\alpha = 1 - \beta$ we obtain the following:

$$\sigma_t^2 = (1-\beta)\sigma_{t-1}^2 + \beta\eta_{t-1}^2$$

Now this may or may not look familiar but this equation is simply stating that volatility (or more precisely volatility squared or variance) is explained by an Exponentially weighted Moving Average (EMA) of squared price returns. This then produces in the volatility estimate an autocorrelation on the timescale of the EMA as approximately $\tau_{EMA} \sim 1/\beta$. So, if $\beta = 0.001$, for example, then the autocorrelation timescale would be approximately $\tau_{EMA} \sim \frac{1}{0.001} = 1000$ business days or approximately 4 years.

Indeed, this is the model we chose above to build our volatility forecast albeit with a timescale for the EMA of 22 days. Our volatility forecast model is a particular form of GARCH that captures short-term autocorrelation or persistence in a reasonably efficient way. This suffices for the purposes of this short note but clearly one can go further using more elaborate techniques to model the dynamics of volatility that are beyond the scope of this discussion.

Disclaimer

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

Capital Fund Management LLP – Sydney branch
Level 16, 333 George Street
Sydney, NSW, 2000, Australia
T +61 2 9159 3100
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