Drive Safely – Controlling volatility in DC default strategies

Tim Horne, DC Investment Solutions Manager**,** Schroders
John McLaughlin, Head of Portfolio Solutions, Schroders

The last decade has been a bruising experience for all those involved in financial markets, not least DC investors. After two major market crises, it is no surprise that most investors today are preoccupied with the challenge of how to protect themselves against the next market correction that awaits them, somewhere down the line.

We now realise that the accepted techniques for portfolio protection are not as robust as we once thought. What we have learned from the last financial crisis is that diversification, while very effective in normal market conditions, can let you down just when you need it most, for example when fear grips the markets and all risky asset classes fall in value at the same time.

Another important lesson learned by professional investors is that it is very difficult to protect your portfolio effectively without sacrificing too much of the upside. This is a particular challenge for DC members who need growth in excess of inflation to accumulate a reasonable savings pot but can't afford to experience severe drawdowns, particularly in the later stages of saving.

So the search is on for a better default solution, one that will provide reliable protection during those short but severe market downturns, but which will still capture growth over the medium term. One strategy that makes such a claim and has been gaining popularity recently involves protection using the so-called 'volatility cap'. What exactly is it? And does it really deliver reliable protection to DC members at an acceptable cost?

A simple braking system – the 'volatility cap'

The basic volatility cap is a remarkably simple mechanism. Starting with a portfolio of risky assets such as equities, we decide on a maximum fixed level of volatility for this portfolio e.g. 15% per annum, measure its prevailing volatility through time and whenever this prevailing volatility level exceeds the fixed 'cap' of 15% per annum, sell some risky assets for cash such that the overall portfolio volatility immediately falls back to 15% per annum again. If the volatility of the risky assets continues to rise, we continue to sell risky assets for cash to maintain the overall portfolio volatility at 15% per annum.

If subsequently the volatility of the risky assets subsides, we can gradually reinvest the cash back into risky assets, taking care not to allow the portfolio volatility to exceed the cap of 15% per annum. By following this simple, systematic rule through time, the portfolio volatility should be effectively limited to 15% per annum (See Figure 1 overleaf).

Figure 1: US equity volatility from 1928 until 2012, capped at 15% per annum

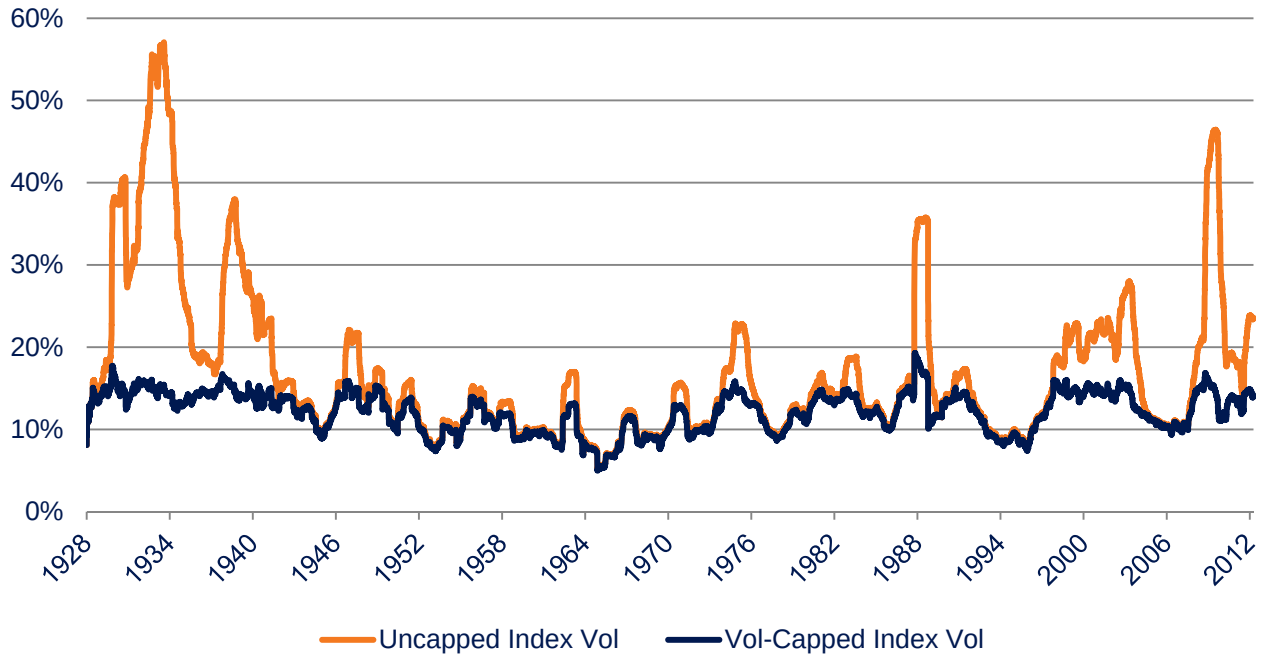

Source, Schroders, Bloomberg. Indices used are the S&P 500 [Div Adjusted] (1928-1988), S&P 500 Total Return (1988-2012). Volatility on any day is measured as the annualised standard deviation of daily returns in the prior 30 days using closing prices and de-risking, if required, is assumed to occur at the close on the same day.

But why do this? Our objective after all is not to control the volatility of the portfolio, but to protect against sharp corrections and at the same time, maintain returns over the medium term. Controlling volatility is not the same thing as protecting against corrections; moreover basic investment theory tells us that higher portfolio returns are a consequence of a willingness to accept more risk. So if we cap the volatility of a portfolio, surely we reduce its expected return?

Remarkably, the volatility cap seems to contradict what our investment intuition tells us. A 15% per annum volatility cap applied continuously to the US stockmarket over the last 90 years (as in Figure 1), would have reduced losses in 12 of the 13 major market corrections that occurred over that period and the protection achieved would have been significant in 8 of those cases (see Figure 2 overleaf). Moreover the capped portfolio would have performed similarly to the uncapped equity market over the entire 90 year period.

Figure 2: Analysis of major corrections in the US equity market since 1928

Years with returns less than -10%	Index return	15% volatility capped return	Difference
1930	-27%	-16%	+11%
1931	-46%	-20%	+26%
1932	-13%	-2%	+11%
1937	-37%	-22%	+15%
1940	-13%	-9%	+4%
1941	-16%	-15%	+1%
1957	-12%	-11%	+1%
1966	-11%	-12%	-1%
1973	-15%	-14%	+1%
1974	-28%	-21%	+7%
2001	-12%	-11%	+1%
2002	-22%	-16%	+6%
2008	-37%	-19%	+18%

Source: Schroders, Bloomberg. Decade ending 1930 only includes data from 1928. Indices used are the S&P 500 [Div Adjusted] (1928-1988), S&P 500 Total Return (1988-2012).

Why might this be? A closer analysis of major market corrections reveals that they are quite often preceded by a jump in short term volatility. (Look again at Figure 1 and see how many of those volatility spikes you can label with a well-known market shock, for example the credit crisis, internet bubble, Long Term Capital Management collapse, Russian debt crisis, etc.). In other words, short term volatility can give early warning of an impending market correction. The volatility cap responds automatically to this warning, de-risking the portfolio just prior to the downturn so that it is better positioned to withstand the blow. Then, after the market finally finds the bottom and volatility subsides, the portfolio can gradually re-risk and participate in the recovery.

The above reasoning also makes intuitive sense. It is not hard to imagine that as the market heads for a fall, internal pressures are manifested first as uncertainty (market volatility) and afterwards explicitly (collapsing prices). Furthermore asset volatility (unlike asset prices) exhibits a behaviour known as 'mean reversion' i.e. volatility has a natural long term equilibrium level (see Figure 1). Volatility may deviate significantly from its mean, but usually not for long. So a well judged volatility cap (set just above the mean) should only demand de-risking for relatively short periods, hence having little impact on returns over long periods. Lastly, there is a school of thought that markets in general tend to rise in a slow and fairly steady fashion until they become over-valued, then fall suddenly until they become cheap again, rise steadily, only one day to fall precipitously again. Although a gross simplification, if it is even partly true that markets tend to behave in this way, then a volatility cap will assist you in the volatile bear markets but not hold you back in the low volatility recoveries.

Active management using a volatility cap

Portfolio hedging always comes down to a compromise between certainty (of downside protection) and cost (to upside returns). Systematic hedging techniques (among which the volatility cap is one example) offer the greatest certainty of not exceeding a maximum loss, but the cost of that insurance is typically prohibitive. Conventional active management on the other hand, is arguably the most cost efficient way to protect your portfolio against short term losses, because the manager is not forced to

follow a 'blind' rule to switch into cash to protect the portfolio, even when the option exists to switch instead into another lower risk asset with better return prospects. However active judgement is inherently uncertain, especially when markets become very risky and fear is driving asset class prices down.

The solution is obvious: combine judgement with system. Specifically, design an investment process such that when market prices are fluctuating within a range that can be considered normal, the active manager is free to exercise judgement without constraint, avoiding pitfalls and maximising returns. However on those rarer occasions when price fluctuations move outside of this normal range, over-ride the active process with a systematic de-risking rule. The volatility cap is perfectly suited to this purpose, as a 'last resort' mechanism to control extreme downside.

There is an analogy here with driving a car. Judgmental management is represented by the accelerator pedal, which is also of course a decelerator. Systematic management is represented by the brake. In most driving conditions, obstacles can be avoided and progress maintained through forward thinking and the skilful use of the accelerator pedal alone. Now and again however a sudden danger arises which could not have been anticipated and the brake has to be applied decisively. A good driver rarely uses his brake but when he does, he brakes decisively.

For a systematic strategy to be relied upon in the context of active management, it should be operated by a team that acts independently of the judgmental manager. Ideally the hedging mechanism should be implemented as a synthetic overlay[1], so that the underlying physical portfolio of the judgemental manager is rarely tampered with. A well positioned volatility cap should result in the systematic 'brake' being applied only in extreme market conditions; at all other times active judgement should predominate.

Measuring volatility

As with any systematic rule, the outcome is only as good as the inputs. For a volatility cap to be effective therefore, we need to measure the prevailing volatility of the underlying portfolio accurately. This is not as easy as it seems. If we use recent historical volatility as a measure of 'instantaneous volatility', there is the immediate problem of how far back to collect the price data. If we use a very short period then the volatility cap will have no 'memory' of distant shocks but if we go back too far, it won't be responsive enough to near term market turbulence. In either case we can also have the problem that when a past market shock drops out of the historical data set, the volatility measurement will suddenly drop, although there is no change to current market conditions.

Some practitioners prefer to use 'implied volatility', which in theory is the market's future expectation of the volatility of an asset, implied from the current trading prices of option contracts based on that asset. Implied volatility has its own problems. By definition it measures uncertainty in the future, when what we are after is a measure of uncertainty right now, before the correction occurs. Implied volatility is directly traded (e.g. using VIX futures on the S&P 500) to take positions in market 'sentiment' and is therefore often a better indicator of general market fear or greed than a measurement of the current volatility of the actual asset you are holding. Lastly, implied volatility data exists only for major traded assets (such as the S&P 500), which is therefore only really useful if your portfolio is comprised of US equities.

For this last reason, at Schroders we prefer to measure volatility using historical data, but in such a way as to iron out the imperfections described above. In particular our method recognises the 'autocorrelated' nature of volatility i.e. the fact that large changes in stock prices tend to be followed by large changes (of either sign), and small changes tend to be followed by small changes. In addition, the way we measure volatility allows us to cap complex portfolios comprising multiple asset classes and highly active asset allocation.

[1] For most conventional portfolios, volatility caps can be implemented using a basket of exchange traded equity and bond index futures. When the cap is exceeded, the systematic manager sells these futures in the correct proportion to reduce the net exposure to risk assets and guide the overall volatility back within its cap limit.

Alas, the volatility cap is not perfect

At this point you would be forgiven for thinking that you have discovered the Holy Grail of investing. Sadly that is not true and the time has come to confess to the weaknesses of the basic volatility cap.

There is still no free lunch. In periods when the pattern of the market is typified by steady recoveries punctuated by violent corrections, a capped portfolio will outperform an uncapped one, as illustrated in Figure 3.

Figure 3: Anatomy of a market crash and recovery – the global financial crisis

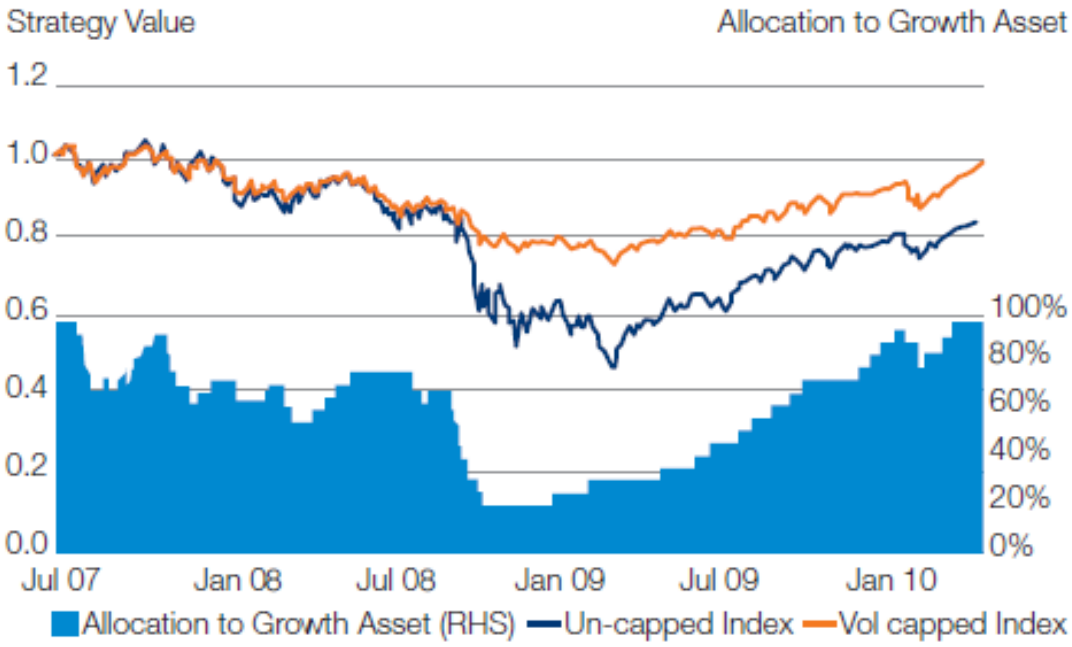

Source: Schroders, Bloomberg. Indices used are the S&P 500 (Dividend Adjusted) (1928–1988), S&P 500 Total Return (1988–2012). Left hand axis shows Index level, right hand axis the net exposure to the Index within the volatility capped strategy. Volatility cap is set at 15% per annum.

However the volatility cap is not fail-safe: heightened volatility usually precedes a market correction, but not always. On occasion, market collapses are truly unexpected, typically when the catalyst for the surprise event is completely unanticipated by the financial markets e.g. a declaration of war or natural disaster. A volatility cap offers no comfort in the 'calm before a storm' scenario. Furthermore, a market collapse and recovery never follows a neat inverted parabola, in particular the recovery is often preceded by several false starts or 'bear market rallies'. If your portfolio has been de-risked at the point when such a rally suddenly takes off, you will miss the boat.

Arguably this is not the most important boat to miss in the long run however (especially if one has survived the initial storm), provided one is able to capture the true recovery when it eventually materialises. Volatility control does not equal downside protection. Although the evidence does suggest that capping volatility does help to limit losses during most severe bear markets, it is nevertheless impossible to predict the size of that loss limitation. The basic volatility cap cannot target an explicit maximum loss amount. Often however, the end client wishes the manager to target an explicit downside target commensurate with that client's absolute risk tolerance.

We can achieve this via a simple adaptation to the basic volatility cap, so as to target the explicit downside objective while avoiding 'cash lock[2]. We call this a 'variable volatility cap'. This mechanism does not guarantee that the risk objective will be achieved, although the probability of not achieving it should be low. Compared with 'hard' protection strategies however (such as option overlays), it will be much more cost efficient to implement and much less restrictive on the underlying active manager. We will discuss the 'variable volatility cap' in more detail in a separate Thought Piece.

Conclusion

DC members require long term growth on their savings with protection from severe market corrections, at a reasonable cost. In our view the optimal approach to downside risk management is one where active judgment is applied in 'normal' markets, but with a systematic 'brake' also in place that is applied only in extreme conditions, when emotion prevails and de-risking should no longer be discretionary but rather automatic and decisive.

The volatility cap is not a panacea, however it does appear to be an effective way to mitigate portfolio losses from severe market corrections, with limited impact on medium term expected returns.

To discuss the themes in this article further, please contact Tim Horne, DC Investment Solutions Manager, on +44 (0)20 7658 4877 or email ukpensions@schroders.com.

www.schroders.com/ukpensions

Important Information

The views and opinions contained herein are those of Tim Horne, DC Investment Solutions Manager and John McLaughlin, Head of Portfolio Solutions at Schroders, and may not necessarily represent views expressed or reflected in other Schroders communications, strategies or funds.

For professional investors and advisers only. This document is not suitable for retail clients.

This document is intended to be for information purposes only and it is not intended as promotional material in any respect. The material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The material is not intended to provide, and should not be relied on for, accounting, legal or tax advice, or investment recommendations. Information herein is believed to be reliable but Schroder Investment Management Limited (Schroders) does not warrant its completeness or accuracy. No responsibility can be accepted for errors of fact or opinion. This does not exclude or restrict any duty or liability that Schroders has to its customers under the Financial Services and Markets Act 2000 (as amended from time to time) or any other regulatory system. Schroders has expressed its own views and opinions in this document and these may change. Reliance should not be placed on the views and information in the document when taking individual investment and/or strategic decisions.

Issued by Schroder Investment Management Limited, 31 Gresham Street, London EC2V 7QA. Registration No. 1893220 England. Authorised and regulated by the Financial Conduct Authority.

For your security, communications may be recorded or monitored.

2 'Cash lock' is an undesirable feature of a well known protection strategy called Continuous Proportion Portfolio Insurance. In a deep bear market, the portfolio becomes so heavily invested in low risk assets that it is unable to participate meaningfully in any subsequent recovery of the risky asset class.

Schroders