Alternative Investment Management Association

The Forum for Hedge Funds, Managed Futures and Managed Currencies

A constant volatility framework for managing tail risk

Alexandre Hocquard, Sunny Ng and Nicolas Papageorgiou[1]

Brockhouse Cooper and HEC Montreal

September 2010

[1]Alexandre Hocquard is Portfolio Manager, Brockhouse Cooper. Sunny Ng is Head of Alternatives Research, Brockhouse Coopers. Nicolas Papageorgiou is Associate Professor of Finance, HEC Montreal and Director of Quantitative Research, Brockhouse Cooper.

The recent financial crisis serves as a timely reminder of the substantial risk of investing in financial markets. It also highlights the limitations of conventional, asset allocation-based risk management strategies. The swift and relentless correction in equity, commodity and real estate markets was a clear example of why diversification, both geographically and across assets classes, is neither a sufficient nor reliable risk control mechanism. During crises, historical correlations between asset classes and their volatility characteristics tend to break down; asset classes which have, in normal times, been uncorrelated, suddenly become correlated and alternative investments, which have been selected based on their ability to generate alpha without beta, suddenly appear to deliver high beta with little alpha. The phase-locking behavior that occurred during the most recent crisis, coupled with the jump in the level of market volatility, resulted in dramatic drawdowns for many investors and has put the spotlight on risk management.

In this paper we present a novel, cost-effective portfolio management approach that focuses on delivering returns that have a constant volatility and that do not unduly expose the investor to the risk of fat-tails.

Traditional tail-risk management techniques: portfolio insurance

For a tail risk hedge to be effective it should possess two important characteristics: the hedge must be negatively correlated to asset returns and exhibit convex behavior to the upside during periods of market stress. Typically, implementation of tail risk hedging has involved the use of equity put options. Unfortunately, the cost is often prohibitive and as a result the drag on the performance of the portfolio is significant So the question is how can investors protect their portfolios against large drawdowns without having to give up substantial upside? The answer lies in properly understanding and monitoring market volatility.

Re-thinking volatility: Black Swans vs. White Swans

Many researchers and quantitative strategists (including Black Swan enthusiast Nassim Taleb and his dedicated followers) have long advocated the importance of paying greater

consideration to the tails of the distribution, calling attention to the fact that traditional risk management methods typically underestimate the frequency and/or severity of tail events. Although the normality assumption of asset returns certainly makes the mathematics a lot easier, it undeniably struggles to explain the empirical evidence.

Modern Portfolio Theory (MPT) has been the crux of the 60/40 strategic asset allocation paradigm, which many plan sponsors employ in one form or another. One of the key assumptions of MPT is that asset returns follow a normal distribution with constant volatility. However if we examine Figure 1, which plots the levels of the S&P 500 index (blue line) and its implied volatility (red line) over the last 20 years, it clearly shows that volatility does not remain constant but in fact changes significantly over time.

Figure 1: S&P500 index and VIX 1990:2010

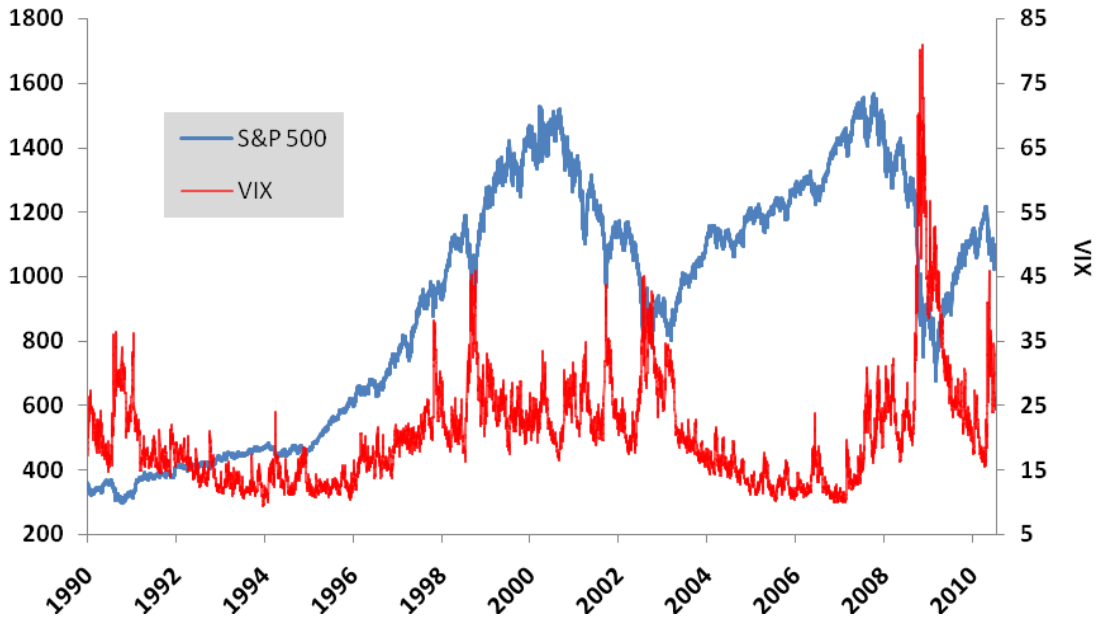

To illustrate how the time varying nature of volatility affects MPT, and in particular the associated tail risk assumptions, we can look to the most recent financial crisis. If we use the average historical volatility of the S&P 500 as our reference point, the monthly decline in U.S. equity markets during October 2008 would be considered close to a four-standard-deviation event. Under the common assumption that returns are normally distributed, a four standard

deviation has a nearly one in 10,000 chance of occurring, implying that a monthly loss of that magnitude should occur approximately once every 750 years. Such a rare event should, from a statistical point of view, be considered somewhat of a Black Swan. However, when we consider the actual returns of the S&P500 over the last 80 years, October 2008 only ranks ninth in terms of worst monthly performances, implying that such a significant drawdown is not nearly as unlikely as we would imagine. The assumption of normally distributed historical returns clearly underestimates the probability of tail events.

There are two possible avenues that can be pursued to better characterize and model the inherent risk in equity returns. The first is to use complex statistical distributions to help parameterize the true tail risk, but this represents a significant shift away from the traditional way of thinking. The second and more appealing approach is to simply re-think how we measure and interpret volatility within our traditional mean-variance framework. Rather than the historical level, we believe that the prevailing level of volatility in the market is the relevant measure. If we use the prevailing level of volatility as our reference point, the drawdown experienced in October 2008 is closer to a one-standard- deviation event. October 2008 suddenly becomes much less of a Black Swan, just an undesirable white one.

The volatility of volatility: A story of two tails

Figure 2 illustrates two Normal distributions: the blue distribution has volatility (standard deviation) of 15 per cent while the red distribution has a volatility of 30 per cent. We note that while the mean return for the two distributions is the same, the probability of a large loss (or gain) is significantly higher for the red distribution. In fact, the probability of losing 30 per cent or more is approximately eight times higher for the red distribution.

Figure 2: Comparison of Tail risk for Normal distributions with different volatilities

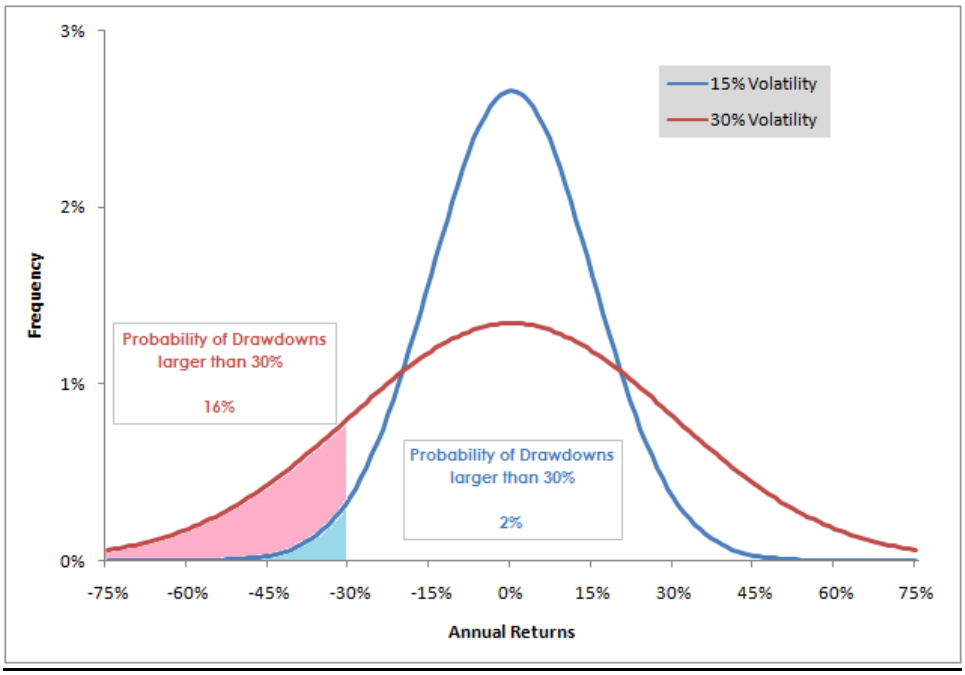

The important take away from this graph is that the two distributions do not necessarily represent two different assets; in fact, they can represent the same asset at two different points in time. Conditions in the markets change over time and, as a consequence, the risk (volatility) profile of a given asset also varies. As market volatility increases, the distribution of returns for the asset flattens and the tails appear to fatten relative to their average historical distribution. Furthermore, as volatility increases, the probability of the asset undergoing large swings becomes much greater and historical probabilities are no longer representative of actual loss potential. The temporal cumulative effect of variable volatility leads to asymmetric tails in the assets return distribution, and in particular the 'unwanted' negative fat tails. With this in mind, we note that efficient frontier analysis/strategic asset allocation based on a static measure of volatility is relatively useless as a risk management tool.

Volatility regimes and asset returns

Thus far we have shown that volatility is dynamic and therefore maintaining a static strategic asset allocation can put significant capital at risk during periods of heightened volatility. However, before we can begin thinking about how we can potentially incorporate volatility in our tail risk management programs, we must understand the relationship between market returns and volatility.

Let us consider the last 20 years of data on the typical equity exposure of a Canadian pension fund (50% TSX, 25% EAFE and 25% S&P500). A Hidden Markov Model (HMM) is used to identify the presence of three volatility regimes (high; medium; and low) and to estimate the parameters for the three regimes. Figure 3 presents the annualized average volatility for each regime and the corresponding annualized return.

Figure 3: Volatility regimes for typical Canadian pension fund equity exposure (1990:2010)

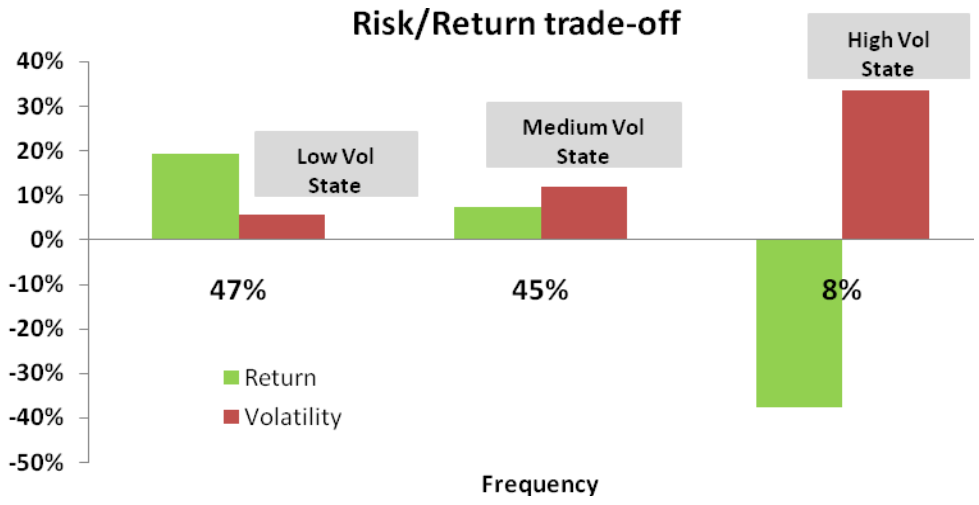

There is a clear trend across the three regimes. The high volatility regime, which occurs 8% of the time, produces an average volatility of 33% and an average annualized return of -38%. Markets find themselves in the medium volatility regime 45% of the time, when the average volatility and average annualized return are 12% and 7% respectively. The most likely regime is

the low volatility regime (47% probability) and it offers by far the best risk/reward tradeoff, with an average volatility of 6% and an average annualized return of 19%. Most of the risk premium provided by equity markets is extracted during these periods of low volatility.

We believe that there is a fundamental behavioral explanation for this persistent market phenomenon. As we can observe in Figure 4, bull markets tend to last longer and develop over time as market participants become increasingly confident in equity returns. These drawn out periods of positive returns and low volatility generate very significant capital appreciation. In contrast, most major market declines are of short duration and develop rapidly as fear takes hold of market participants. The sharp drawdowns are therefore much more dramatic and the volatility is markedly higher. Knowing the regime that we are in provides us with important information about the short term risk/return trade-off that the markets are offering.

Figure 4: Cumulative returns and volatility regimes for proxy portfolio

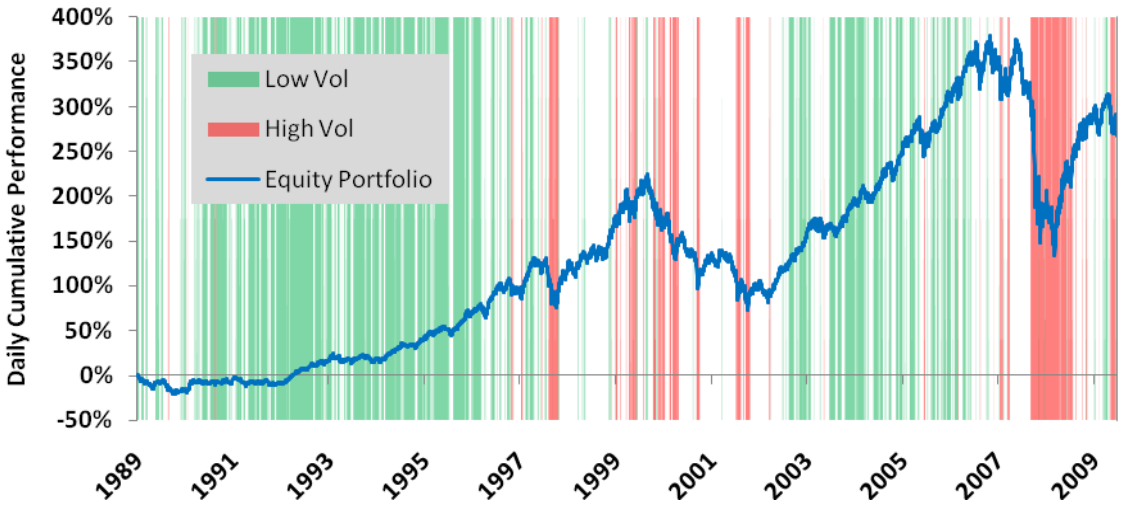

Using volatility to smooth returns and manage tail risk

The fact that equity volatility is not constant implies that the level of risk (and therefore the probability of a large drawdown) for a given portfolio is constantly changing. If, however, we can accurately measure the prevailing level of volatility and effectively hedge against changes in that volatility, we can greatly reduce tail risk and potentially improve our portfolio's risk

adjusted returns. Because most assets tend to exhibit volatility clustering, the recent (realized) volatility of an asset provides useful information about the near term risks.

The idea of using volatility as a risk conditioning variable to optimize the portfolio has been proven to be extremely efficient. Fleming, Kirby and Ostdiek (2001, 2002) studied the economic value of volatility timing and found that volatility timing strategies outperform a static portfolio in a mean-variance optimization framework. More recently, Cooper (2010) finds that risk smoothing can generate alpha due to the predictability of volatility. Although these results clearly support the predictability and use of volatility as a conditioning variable, we are still confronted with the issue of how to translate the prevailing level of volatility to a level of portfolio exposure. To address this issue, we propose an innovative approach based on the Payoff Distribution Model (PDM) to target a constant level of portfolio volatility and control the risks related to higher moments of the distribution.

Implementing a constant volatility strategy

The constant volatility framework can be implemented at the fund level or on the top of an existing portfolio of assets. In the latter case, tradable benchmarks representing the assets in the portfolio must be defined and an overlay of long and short futures contracts is used to adjust the portfolio's market exposures to target a pre-specified distribution and level of volatility. Importantly, the overlay does not, in any form, impact the strategic asset or manager allocation decisions or affect the alpha component of the portfolio. It simply aims to smooth exposure to market (beta) risk using exchange traded futures contracts.

Results

The main analysis in this paper is performed on the Canadian pension fund proxy portfolio (50% TSX, 25% S&P 500 and 25% MSCI EAFE) and with a 12% target level of volatility, which was selected because it is close to the median level of volatility over the sample period. All the indices are assumed fully currency hedged and all profits/losses are reinvested on a monthly basis in the underlying asset. Trading is done on a daily basis according to the forecasted daily

GARCH volatility and the cumulative month-to-date performance of the underlying asset. Performance is presented net of all costs.

Figure 6 illustrates the benefits of methodology. The blue line shows the cumulative return of the benchmark equity portfolio, while the red line represents the cumulative return of a 12% constant volatility portfolio.

Figure 6: Cumulative performance of base equity portfolio and constant volatility strategy

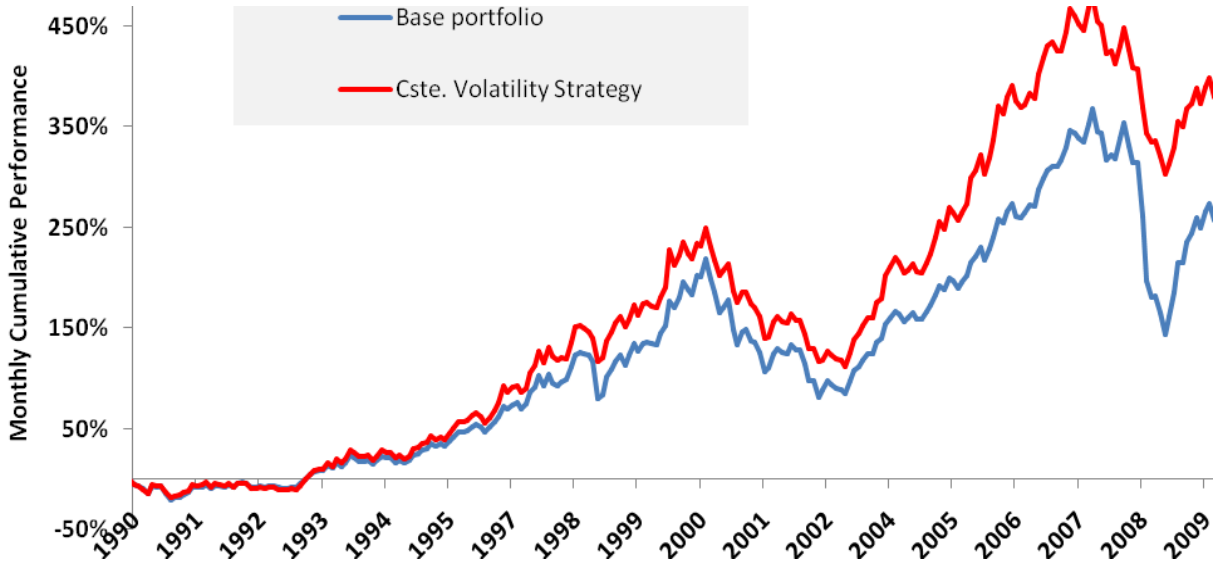

Table 1 summarizes the performance of the two strategies. The numbers confirm the superior risk-adjusted return for the constant volatility fund. The magnitude of the worst drawdowns are also greatly reduced, both on a monthly and annual basis. Finally it is interesting to note that the constant volatility overlay essentially eliminates the higher moments of the return distribution, essentially rendering the distribution Normal.

Table 1: Descriptive statistics for equity portfolio and constant volatility strategy

	Base Equity Portfolio	12% Const. Vol. Strategy
Ann. Return	7.40%	8.58%
Ann. Volatility	14.50%	13.08%
Sharpe Ratio	0.23	0.34
Omega Ratio	2.28	2.70
Skewness	-0.88	0.03
Excess Kurtosis	2.08	-0.30
Correlation	--	93.42%
Worst Month	-17.57%	-9.70%
Best Month	10.08%	12.28%
Worst Year	-36.43%	-20.94%
Best Year	32.46%	35.73%
95% VaR 1-Month	-7.51%	-5.46%
Jarque Bera p-value	0.1%	50%
Liliefors p-value	1.2%	18%

Table 2 highlights the performance of the strategy during the two largest market drawdowns during the sample period, specifically the collapse of the tech bubble and the recent financial crisis.

Table 2: Performance during drawdowns

	Tech Bubble (Aug 2000 to March 2003)	
	Maximum Drawdown	**Months to Recovery**
Base Equity Portfolio	43.32%	35
Cons. Vol. Strategy	38.35%	21

	Financial Crisis (Oct 2007 to March 2009)	
	Maximum Drawdown	**Months to Recovery**
Base Equity Portfolio	48.10%	--
Const. Vol. Strategy	30.97%	--

There are several interesting observations to take away from this analysis:

- During the recent credit crisis, the constant volatility fund greatly reduced the size of the drawdown. As volatility rose in 2008, the strategy progressively decreased market exposure in order to maintain the volatility at 12 per cent, thereby protecting the portfolio when markets subsequently plunged;
- During the bull markets in the late 1990s and from 2002-2007, the strategy actually over performed the base portfolio. This is simply because the level of realized volatility during these up-trending markets was below the 12 per cent target and therefore leverage was added to bring the risk exposure back to the desired level;
- During the 2000-2003 recession and market correction, the strategy only provided marginal downside protection. This is not surprising as markets drifted downwards over an extended period of time with no sustained increase in volatility.

Conclusion

Since the collapse of Lehman Brothers in 2008, tail-risk hedging has become an increasingly important concern for investors. Traditional approaches such as purchasing options or variance swaps as insurance are often expensive, illiquid and result in a substantial drag on performance. Furthermore, due to the time-varying nature of volatility, asset returns have been shown to behave in a non-Normal fashion which increases the likelihood of negative tail events for portfolios which maintain static asset allocation. A more cost-effective and prudent approach to managing risk is to actively manage the exposure of a portfolio, based on the prevailing level of volatility within the underlying assets, in order to maintain a constant risk exposure. This approach of portfolio and risk management can help investors obtain the desired risk exposure over the short-term and long-term, reduce exposure to tail-risk and, in general, increase the risk–adjusted performance of the portfolio.

References

Cooper, C. (2010), Alpha Generation and Risk Smoothing using Volatility of Volatility*, Working paper,* Double Digit Numerics.

Fleming, Kirby and Ostdiek (2001), The Economic Value of Volatility Timing, *Journal of Finance*, 56(1): 329–352.

Fleming, J., Kirby, C. and Ostdiek, B. (2003), 'The economic value of volatility timing using "realized" volatility', *Journal of Financial Economics* 67, 474–509.