



A constant volatility framework for managing tail risk

Alexandre Hocquard, PhD, Sunny Ng, CFA and Nicolas Papageorgiou, PhD¹

Brockhouse Cooper and HEC Montreal

September 2010

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



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. Increasingly, investors are realizing the importance of mitigating tail risk in order to achieve their long-term investment objectives. Most investors can withstand an annual loss of five per cent or even 10 per cent but few are able to absorb another drawdown such as the one they suffered in 2008.

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. As an alternative to purchasing put options, investors can also resort to dynamic portfolio insurance strategies. The earliest dynamic portfolio insurance model, proposed by Brennan and Schwartz (1979) and Rubinstein and Leland (1981), consists of overlaying a synthetic put option on the existing portfolio, and delta managing the overall exposure. Other dynamic strategies include the notorious Constant Proportion Portfolio Insurance, an arguably more robust approach that was proposed by Black and Jones (1987) and Black and Perold (1992). Although all dynamic hedging strategies are exposed to some level of gap risk, there are numerous benefits to dynamic hedging over buying put options. These include no broker premiums, no up-front costs, complete flexibility to change adjust or remove the hedge and no exposure to counter party risk. The fact remains however that all portfolio insurance techniques result in a significant drag on the portfolio's performance.

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.

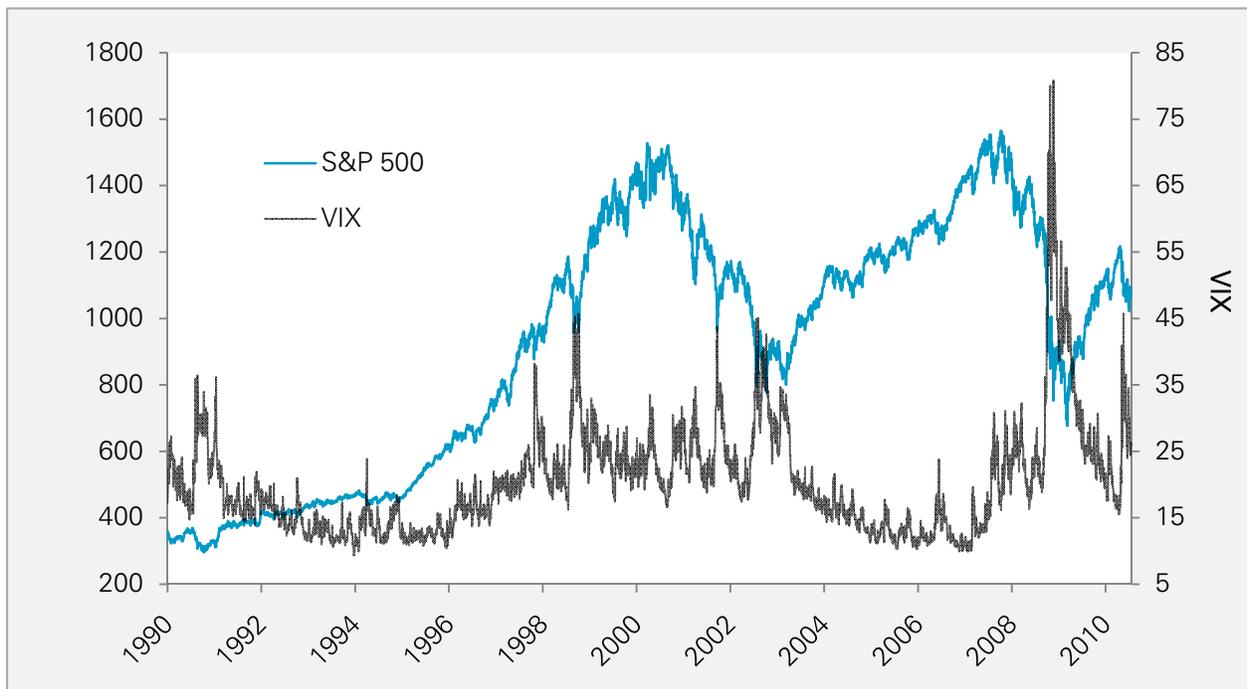


RE-THINKING VOLATILITY: BLACK SWANS VS. WHITE SWANS

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. We note that, over the measurement period, the VIX index ranged from under 10 per cent to a peak of over 77 per cent.

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

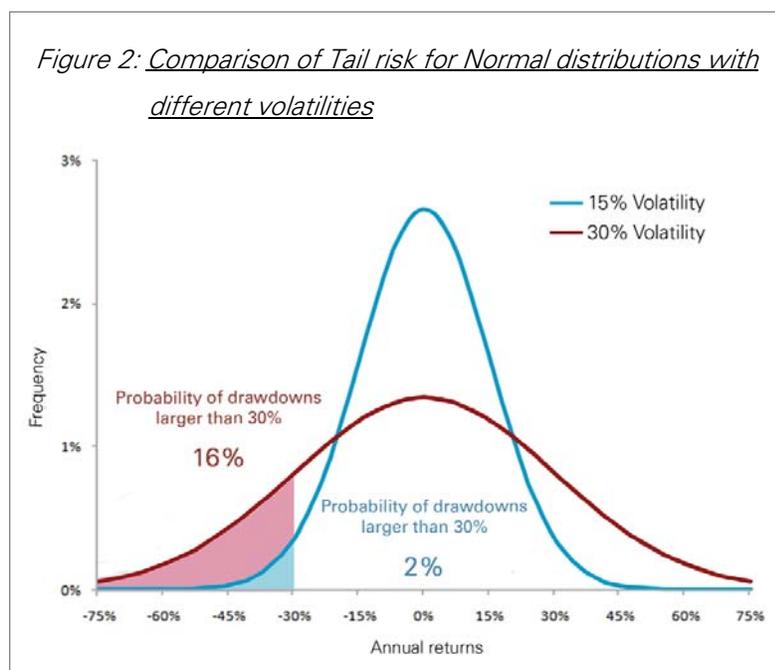


assumption that returns are normally distributed, a four standard deviation has a nearly one in 10,000 chance of occurring, implying that a monthly loss of that magnitude should occur approximately once every 750 years. Such a rare event should, from a statistical point of view, be considered somewhat 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 based, for example, on extreme value theory to help parameterize the true tail risk. This is a highly quantitative approach and 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

When we consider the historical return distribution of a given asset, the volatility provides us with a measure of the dispersion of the returns around the mean. 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.

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.

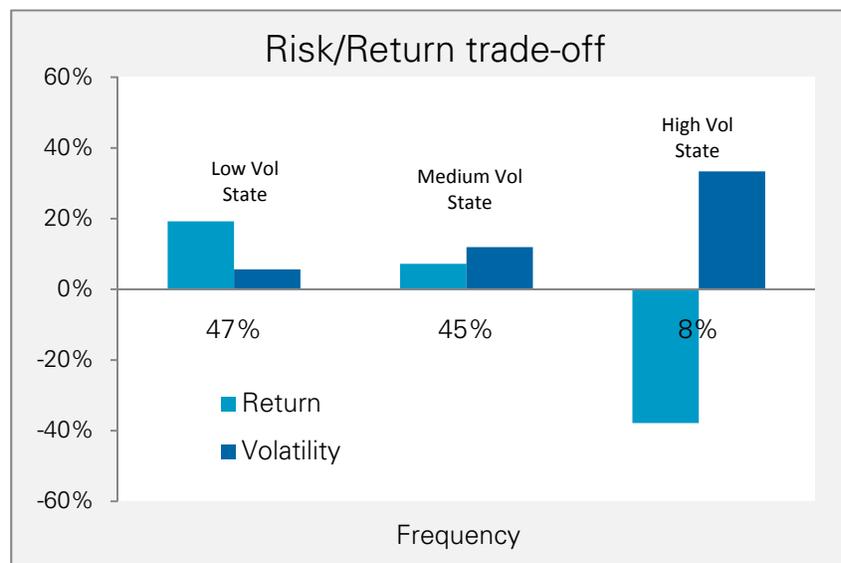
VOLATILITY REGIMES AND ASSET RETURNS

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.

There is substantial empirical evidence that market returns form volatility clusters - in other words, as noted many years ago by Mandelbrot (1963), "large changes (in returns) tend to be followed by large changes (in returns), of either sign and small changes tend to be followed by small changes."

Let us consider the last 20 years of data on the typical equity exposure of a Canadian pension fund² (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)



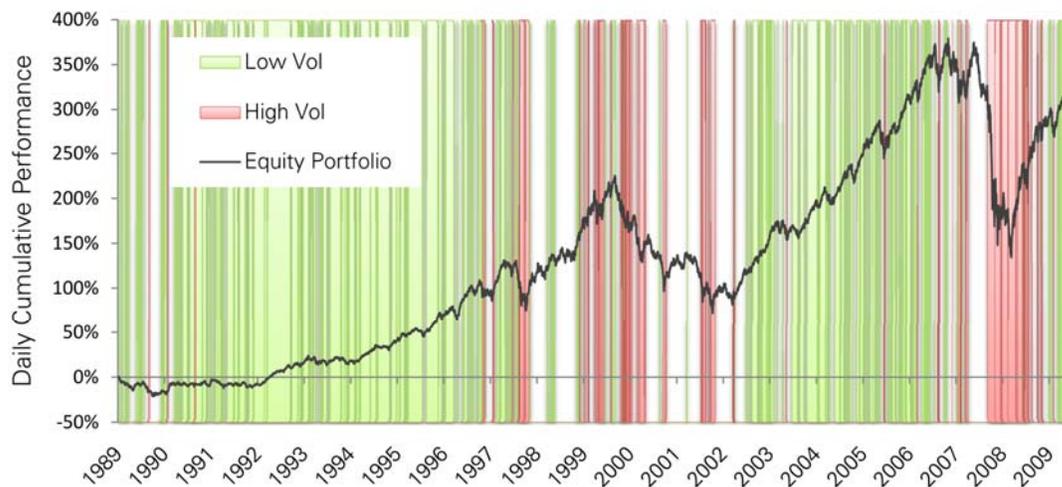
² For the sake of simplicity, we assume the portfolio is fully currency hedged



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. On this basis, 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. An abundance of literature has been published on the notable dependence and forecastability of the volatility of asset returns and its implications on asset allocation, asset pricing and risk management. Bollerslev, Chou and Kroner



(1992) provide a fine review of the academic literature on utilizing GARCH to forecast volatility, while Ghysels, Harvey and Renault (1996) surveyed the literature on stochastic volatility, Franses and van Dijk (2000) on regime-switching and volatility, and Andersen, Bollerslev and Diebold (2002) on realized volatility models.

The most common approach to using volatility for tail-risk hedging purposes involves the purchase of variance swaps. Variance swaps are over-the-counter (OTC) forward contracts on volatility in which the buyer agrees to swap a fixed variance level on a particular market index for actual realized variance from now until the maturity date. As such, variance swaps provide pure exposure to the realized volatility of an asset. They can be used to take views on future volatility, to capture the spread between realized and implied volatility and to hedge asset volatility exposure. A tail risk hedging strategy would simply involve purchasing a basket of variance swaps on the particular market which you are trying to hedge. The drawbacks of such an approach are that variance swaps are OTC products that are relatively illiquid, offer limited capacity, are priced based on the prevailing level of implied volatility and are accompanied by a generous brokerage premium. Most recently, these instruments have garnered much attention and, as a consequence, there is overwhelming demand on the long-side for variance swaps making them extremely expensive.

DYNAMIC EXPOSURE AND CONSTANT VOLATILITY

The underlying principle of constant volatility is to systematically adjust the exposure to a given asset (or portfolio of assets) conditional to its current volatility in order to maintain a pre-specified level of risk. For example if we were to target a 12 per cent level of risk for a given asset and the current volatility of the asset was 14 per cent we would lower our exposure to the asset class by a commensurate amount to yield a 12 per cent volatility and vice versa if the current volatility was lower than our target.

The rationale for maintaining a constant volatility is two-fold. Firstly, most significant market corrections have been preceded by an increase in market volatility. By conditioning one's exposure to the level of volatility in the market, the impact of the market correction will be significantly dampened. Secondly, empirical evidence shows that asset returns tend to be greater during periods of low volatility. Most bull markets have been characterized by extended periods of below-average volatility. Markets generally trend upwards in an organized and relatively smooth pattern. During these periods, investors should maximize their exposure to the asset as the risk-reward tradeoff is most favorable. As volatility increases, exposure to the asset should be reduced in order to maintain the desired level of risk.

The idea of using volatility as a trading signal is by no means a novel idea. For many years, trends in volatility have been exploited by active managers and derivatives traders. For example, option traders seek to identify volatility clusters to take advantage of the asymmetry in implied volatility



surfaces. Trend followers in the equity and currency markets often condition their strategies on volatility states, often assuming that market trends are more stable in low volatility environment. During the recent crisis, interest in volatility has reached unprecedented levels, with several volatility based products being developed for both institutional and retail clients. With the introduction of ETF's that replicate the short and mid-term Volatility Index (VIX) and volatility swaps, investors can now directly invest in volatility-based products. Increasingly, market participants are realizing the importance of incorporating volatility in the asset allocation process. 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) defines the volatility of volatility "vovo" and identifies trading strategies using leveraged ETF's to target the desired risk exposure. The author concludes that constant volatility strategies are able to profit from the upside of leverage without all of the downside. In effect, 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.

THE PAYOFF DISTRIBUTION MODEL

The PDM was introduced by Dybvig (1988) to price and evaluate the distribution of consumption for a given portfolio. The main idea was to propose a new performance measure that allowed preferences to depend on all the moments of a distribution, providing a richer framework for decision making than the traditional mean-variance approach. In this paper we extend the PDM to a more general portfolio and risk management methodology. The PDM allows us to derive and price any contingent claim on an underlying asset or pool of assets. Specifically, we use the PDM to solve for the payoff function that provides us with the target (desired) return density conditional to the distributional properties of underlying asset. In effect it provides us with the necessary distortion that must be applied to the distribution of the underlying asset in order to generate the desired distributional properties. We employ the methodology proposed in Papageorgiou et al. (2008) to replicate such distribution payoffs by delta managing the underlying asset. By construction the aggregation of the monthly payoffs will deliver the specified target density over long term.

The first step in the PDM approach is to derive the monthly payoff structure for the target distribution. In the case of a constant volatility fund, we target a Normal distribution and level of volatility. Once the monthly payoff structure is determined, we dynamically adjust the portfolio exposure in the underlying asset to achieve two key objectives: 1) a constant level of volatility,



regardless of the prevailing level of volatility in the market, and 2) a Normal distribution of monthly returns, in order to “normalize” the fat-tailed distribution of the underlying asset. In Amin and Kat (2003), the authors show that given an underlying asset S_{Under} with monthly returns R_{Under} and a target distribution F_{Target} , there exists a function $g(R_{Under})$ such that the distribution of $g(.)$ is the same as the distribution F_{Target} . This payoff's return function g is calculated using the distribution function F_{Under} of the underlying asset and the marginal distribution function of the targeted distribution F_{Target} .

The exact expression for g is given by

$$g(x) = F_{Target}^{-1}(F_{Under}(x)); \forall x \in R$$

Instead of being written on the price of the underlying like traditional call and put options, this payoff function g is written on the underlying asset's monthly returns. This implies a more adapted payoff functions that integrates the entire risk profile of the asset.

IMPLEMENTATION OF THE MODEL

In this section, we present a brief overview of the Payoff Distribution model and demonstrate how the model can be used to derive the required exposure so that the target volatility strategy can be implemented.

The steps required to generate a synthetic fund with a target Normal distribution and constant volatility are as follows:

1. Define the underlying asset or fund and if required, its tradable proxies. In our case we will restrict our study to equity and commodity indices where listed futures contracts are available.

2. Select the desired statistical properties of the target fund. We target a Normal distribution of monthly returns and a pre-specified level of volatility to illustrate the benefits of the strategy.

3. Estimate the daily process of the underlying asset returns and infer its monthly distribution.

To adapt the methodology of Papageorgiou et al. (2008) to a dynamic volatility environment we model the daily returns of the underlying assets as a simple GARCH (1,1) process. This model allows us to capture two specific features of the volatility, namely the short term serial correlation and the long-run mean reversion. GARCH family models have been widely employed in the finance industry to characterize the evolution of return variability. We could have used more adapted GARCH models, such as NGARCH or EGARCH, but we opted to keep the modeling approach relatively simple to better highlight advantages of the hedging methodology.



The GARCH (1,1) can be written under the physical measure such as:

$$\log \frac{S_{Under,t}}{S_{Under,t-1}} = R_t = \mu + \sigma_t \epsilon_t, \quad \epsilon_t \sim i.i.d. N(0,1)$$
$$\sigma_t^2 = \omega + \beta \sigma_{t-1}^2 + \alpha (R_{t-1} - \mu)^2, \quad \text{with } \alpha + \beta < 1$$

Where $S_{Under,t}$ is the level of the underlying asset at time t (in days), and R_t is the daily log-return.

The parameters will be estimated using standard maximum likelihood maximization. The estimation is performed every month using all available data.

4. Derive the monthly payoff of the targeted distribution. The payoff, or the function g , can be written in closed form such as:

$$g(x) = \mu_{Target} + \sigma_{Target} * \Phi^{-1} \left[\Phi \left(\frac{x - \mu_{Under}}{\sigma_{Under}} \right) \right], \quad \text{with } x \text{ the monthly underlying return.}$$

Where

μ_{Target} is the monthly targeted expected return. For the sake of simplicity we set $\mu_{Target} = \mu_{Under}$.

μ_{Under} is the monthly expected return of the underlying, computed as the historical expected return.

σ_{Under} is the monthly volatility of the underlying, set as the possible levels of forecasted volatility of the underlying at the end of the month.

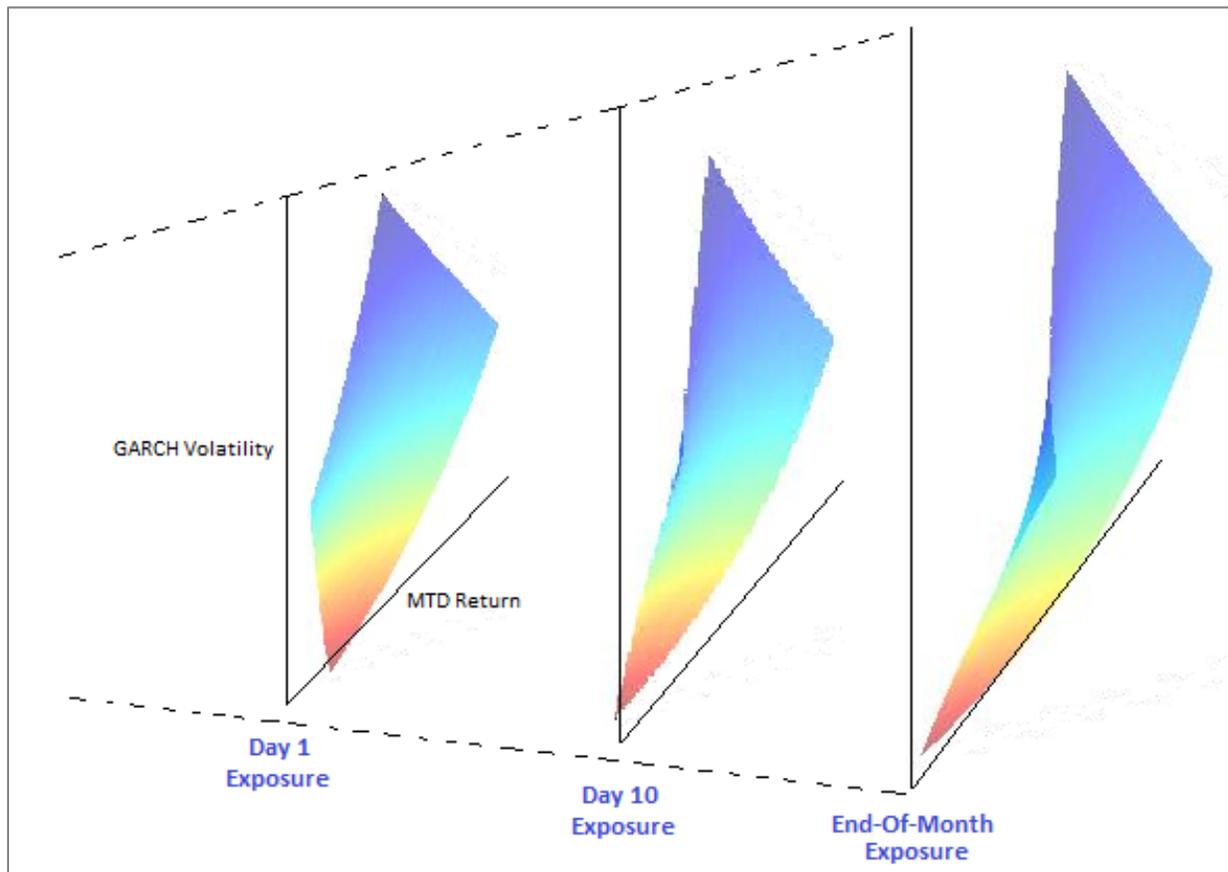
σ_{Target} is the targeted monthly volatility, that will allow the constant volatility property.

Φ is the standard normal cumulative distribution function and Φ^{-1} the inverse.

5. Derive the hedging strategy throughout the month. In essence, the dynamic trading strategy distorts the distribution of the underlying asset so as to generate the desired payoff. We price and derive the replication strategy by minimizing the root mean square hedging error using a Monte Carlo approach under the real probability measure. For more details on the hedging methodology one could refer to Papageorgiou et al. (2008).



Figure 5: *Delta Surface evolution from Day 1 to end-of-month*



As a discrete time hedging strategy, delta surfaces are computed for every day during the month. In figure 5 above, we illustrate the evolution of the exposure grid through time, with representations of day 1, day 10 and end-of-month exposure. The exposure is conditional to the GARCH forecasted volatility of the underlying asset and its cumulative monthly performance.

IMPLEMENTING A CONSTANT VOLATILITY AS AN INDEXED FUND

To illustrate the effectiveness of the constant volatility approach in managing tail risk, and normalizing return distributions, we implement the strategy on various equity indices and the GCSI commodity index on an out of sample basis from January 1990 to August 2010. 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 management fees, margin and financing costs and implementation costs (25bps).



IMPLEMENTING A CONSTANT VOLATILITY OVERLAY ON A PORTFOLIO

The constant volatility framework could also be implemented on the top of an existing portfolio of assets. To do so, 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. The strategy can be implemented using exchange traded futures contracts, thereby eliminating any potential liquidity or capacity constraints, counterparty risk and offering full transparency with minimal transaction costs.

RESULTS

The 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. We provide robustness tests with respect to both the specified level of volatility and the composition of the underlying portfolio in the following section.

Figure 6 illustrates the advantages of implementing a 12% constant volatility strategy on the Canadian pension fund proxy. 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

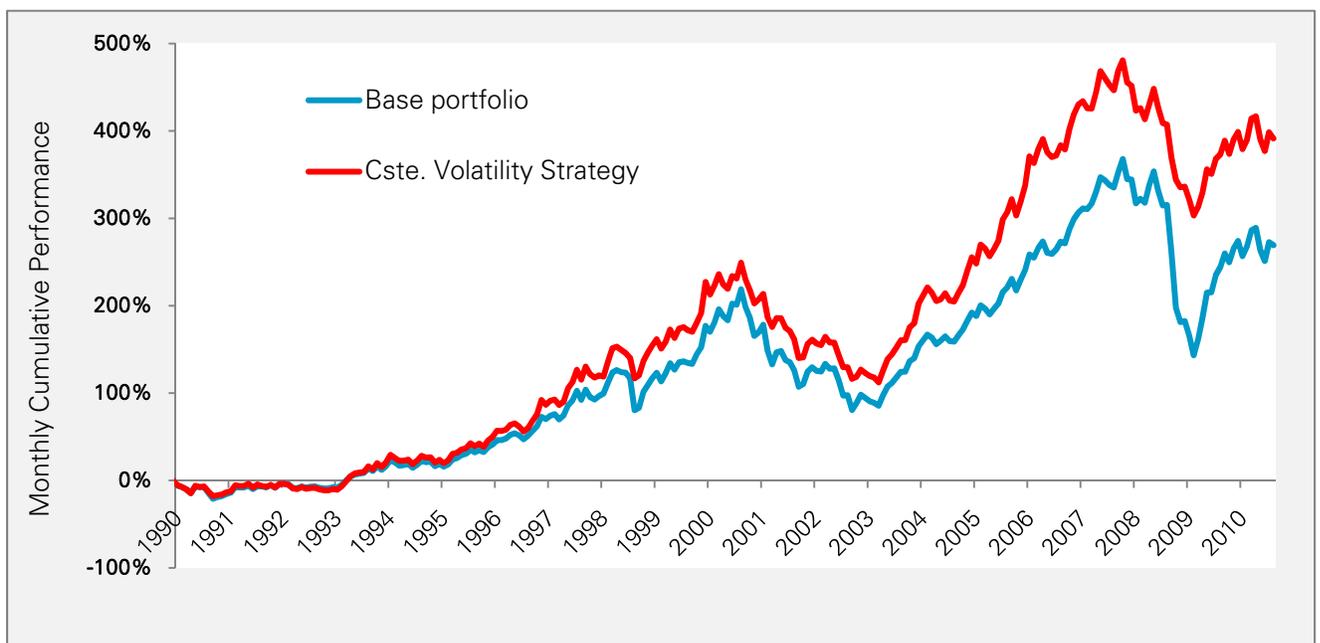
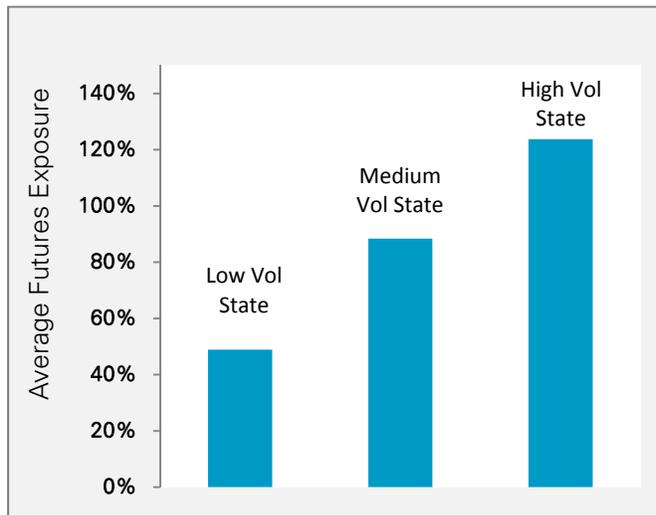




Figure 7: Average exposure of the strategy for each volatility regime



As shown in Figure 7, during periods of “normal” volatility (near the 12 per cent target) the strategy is almost fully invested in the underlying assets through futures contracts, and hence tracks the underlying portfolio quite closely. Only when volatility diverges from the target level, does the constant volatility strategy come into play.

Figure 8 illustrates the modeled GARCH volatility on the out-of-sample generated returns of the constant volatility strategy and the base portfolio. The realized volatility evolves slightly above the targeted 12% value but does not suffer anywhere near the same variability as the base portfolio.

Figure 8: Monthly GARCH volatility for the base portfolio and the 12% constant volatility strategy





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 1 summarizes the performance of the two strategies. We compute several performance and risk measures, such as Sharpe Ratio, Omega Ratio and 95% 1-month value at risk. In contrast to the Sharpe Ratio, the Omega Ratio introduced by Keating and Shadwick (2002) relaxes the hypothesis that returns follow a Gaussian distribution. This measure leads to a full characterization of the risk reward properties of the distribution by measuring the overall impact of all moments. The numbers confirm the superior risk-adjusted return for the constant volatility fund: the Sharpe Ratio increases from 0.23 to 0.36 and Omega Ratio goes from 2.28 to 2.70. 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 approach essentially eliminates the higher moments (skewness and excess kurtosis) of the return distribution, essentially rendering the distribution Normal. We present two common normality tests to test for the Gaussian nature of the monthly returns series. We note that both tests exhibit high p-values for the constant volatility strategy according to which we cannot reject the Normal distribution assumption of the returns at the 5% level.

Table 2 highlights the performance of the constant volatility 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)		Financial Crisis (Oct 2007 to March 2009)	
	Maximum Drawdown	Months to Recovery	Maximum Drawdown	Months to Recovery
Base Equity Portfolio	43.32%	35	48.10%	--
Cons. Vol. Strategy	38.35%	21	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.

ROBUSTNESS TO TARGET VOLATILITY

In Table 3 below we illustrate the results of a target volatility strategy which targets Normal distributions and volatilities ranging from 6% to 20% for the Canadian Pension fund proxy equity portfolio.

Table 3: Constant Volatility Strategies for different target volatilities

Targeted Volatility <14%

	Equity Portfolio	6%	7%	8%	9%	10%	11%	12%	13%
Ann. Return	7.4%	5.6%	6.2%	6.7%	7.2%	7.5%	8.1%	8.6%	9.1%
Ann. Volatility	14.5%	6.6%	7.6%	8.8%	9.8%	10.9%	11.9%	13.1%	14.1%
95% VaR 1-Month	-7.5%	-2.6%	-3.1%	-3.6%	-4.1%	-4.5%	-4.9%	-5.5%	-5.9%
Jarque-Berap-value	0.1%	50.0%	50.0%	50.0%	50.0%	50.0%	50.0%	50.0%	50.0%
Lilliefors p-value	1.2%	18.6%	12.7%	19.3%	26.8%	14.7%	22.3%	18.3%	7.1%



Targeted Volatility >14%

	Equity Portfolio	14%	15%	16%	17%	18%	19%	20%
Ann. Return	7.4%	9.6%	9.9%	10.4%	10.8%	11.4%	12.0%	12.2%
Ann. Volatility	14.5%	15.2%	16.4%	17.5%	18.6%	19.7%	20.8%	22.0%
95% VaR 1-Month	-7.5%	-6.3%	-6.9%	-7.4%	-7.9%	-8.3%	-8.7%	-9.3%
Jarque-Bera p-value	0.1%	50.0%	50.0%	50.0%	50.0%	50.0%	48.9%	50.0%
Lilliefors p-value	1.2%	16.7%	13.7%	13.1%	10.3%	15.8%	15.5%	22.6%

The results are robust to the different target values of volatility although the realized volatilities of the funds are slightly higher than the targeted values. This is due to the fact that the monthly profit/loss of the strategy is reinvested in the fund. These out-of-sample results, that incorporate implementation constraints such as financing and management costs, support the ability of the model to generate the desired risk profile. Regardless of the target volatility level, the monthly returns for all the funds generated using the approach pass the tests for Normality.

ROBUSTNESS TO INDICES

In Table 4 below we show the performance of the constant volatility strategy on different assets; in particular different equity indices and the GSCI commodity index. For all the assets we target a Normal distribution with a 12% monthly annualized volatility.



Table 4: *Constant Volatility Strategies across indices*

12% Volatility Target	TSX 60		SP500		RUSSELL 2000		MSCI EAFE	
Ann. Return	8.55%	8.08%	8.59%	6.47%	9.17%	7.25%	4.39%	3.34%
Ann. Volatility	15.91%	12.47%	15.16%	11.23%	19.65%	13.45%	17.62%	13.22%
Sharpe Ratio	0.28	0.32	0.29	0.21	0.26	0.23	0.01	-0.06
Omega Ratio	2.12	3.03	2.22	2.96	1.62	2.02	1.35	1.39
Skewness	-0.79	-0.16	-0.63	0.12	-0.57	0.39	-0.46	0.05
Excess Kurtosis	2.25	-0.10	1.06	0.05	0.92	0.23	1.11	-0.46
Correlation	100.00%	93.76%	100.00%	92.25%	100.00%	92.36%	100.00%	94.16%
Worst Month	-20.41%	-11.52%	-16.79%	-8.94%	-20.80%	-8.71%	-20.18%	-8.15%
Best Month	12.09%	10.09%	11.44%	11.02%	16.53%	14.87%	15.40%	10.13%
Worst Year	-31.17%	-13.48%	-37.00%	-20.32%	-33.79%	-19.38%	-42.99%	-26.55%
Best Year	34.18%	31.02%	37.58%	45.37%	47.29%	34.59%	39.29%	34.92%
95% VaR 1-Month	-8.20%	-5.34%	-7.36%	-4.80%	-9.54%	-5.39%	-8.85%	-6.04%
Jarque-Bera p-value	0.10%	50.00%	0.10%	50.00%	0.16%	3.37%	0.18%	26.68%
Lilliefors p-value	3.66%	36.38%	0.14%	50.00%	0.10%	4.90%	1.32%	50.00%

12% Volatility Target	MSCI EEM		NIKKEI 225		GSCI	
Ann. Return	11.36%	10.10%	-3.86%	-2.96%	6.55%	5.62%
Ann. Volatility	24.23%	16.37%	22.46%	12.98%	21.37%	12.50%
Sharpe Ratio	0.30	0.37	-0.36	-0.55	0.11	0.12
Omega Ratio	1.56	2.10	0.85	0.76	1.36	2.06
Skewness	-0.74	-0.30	-0.16	0.00	-0.09	0.20
Excess Kurtosis	1.77	0.42	0.57	-0.16	2.00	0.04
Correlation	100.00%	95.53%	100.00%	94.98%	100.00%	93.09%
Worst Month	-29.29%	-14.00%	-23.83%	-10.08%	-27.77%	-7.61%
Best Month	17.14%	13.19%	20.07%	9.87%	21.10%	11.85%
Worst Year	-53.18%	-25.30%	-41.11%	-24.20%	-42.80%	-20.50%
Best Year	78.58%	56.20%	41.51%	32.51%	50.31%	28.51%
95% VaR 1-Month	-12.25%	-7.47%	-11.25%	-6.52%	-9.25%	-5.44%
Jarque-Bera p-value	0.10%	5.31%	8.72%	50.00%	0.10%	41.20%
Lilliefors p-value	0.19%	37.37%	50.00%	41.36%	8.61%	50.00%

For each index, the left column shows the result of a direct investment in the asset over the period 1990:2010, whereas the right hand column is the output of the 12% constant volatility target find on the asset.



The results in table 4 lend strong support to the benefits of implementing a constant volatility framework. The most notable transformation that occurs to the statistical properties of the different assets when we implement the payoff distribution model is the “Normalization” of the returns. In all cases, both the skewness and excess kurtosis are greatly reduced, and both the Jarque-Bera and Lilliefors tests indicate that the returns are Gaussian. Furthermore, correlations between the constant volatility funds and their underlying assets are always greater than 90%, demonstrating that the dynamic leverage does not dramatically alter the nature of the return series; it simply smoothes the volatility exposure over time. In term of performance, the 12% constant volatility funds tend to underperform the underlying indices due to lower level of realized volatility , but deliver a superior risk-adjusted return. Finally, maximum drawdowns are greatly reduced across all assets, demonstrating the important role that controlling volatility can play in reducing tail risk.

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. We implement a robust methodology based on Dybvig’s (1988) payoff distribution model to target a constant level of volatility, and “Normalize” the monthly returns. 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.

DISCLAIMER

This document was prepared for institutional and sophisticated investors only and without regard to any individual’s circumstances. It is not to be construed as a solicitation, an offer, or an investment recommendation to buy, sell or hold any securities. Any returns discussed represent past performance and are not necessarily representative of future returns, which will vary. Any views expressed herein are those of the author(s), are based on available information, and are subject to change without notice. This material and/or its contents are current at the time of writing and may not be reproduced or distributed in whole or in part, for any purpose, without the express written consent of Brockhouse Cooper Asset Management Inc. We accept no liability for any errors or omissions which may be contained herein and accept no liability whatsoever for any loss arising from any use of or reliance on this document or its contents.

© 2010 Brockhouse Cooper Asset Management Inc. All rights reserved.



REFERENCES

- Amin, G. and Kat, H. (2003). Hedge fund performance 1990-2000: Do the "money machines" really add value. *Journal of Financial and Quantitative Analysis*, 38(2):251-275.
- Black, F. and Jones, R. (1987). Simplifying portfolio insurance. *Journal of Portfolio Management*, 14:48-51.
- Black, F. and Perold, A. (1992). Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics and Control*, 16:403-426.
- Brennan, M. J. and Schwartz, E. S. (1989). Portfolio insurance and financial market equilibrium. *Journal of Business*, 62(4):455-472.
- Cooper, C. (2010), Alpha Generation and Risk Smoothing using Volatility of Volatility, *Working paper*, Double Digit Numerics.
- Dybvig, P. (1988). Distributional analysis of portfolio choice. *The Journal of Business*, 61(3): 369-393.
- 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.
- Keating, C. and Shadwick, W. F. (2002). A universal performance measure. Technical report, *The Finance Development Centre, London*.
- Mandelbrot, B. (1963), The Variation of Certain Speculative Prices", 1963, *Journal of Business*, 36(4), 394-419.
- Papageorgiou, N., Rémillard, B., and Hocquard, A. (2008). Replicating the properties of hedge fund returns. *Journal of Alternative Investment*, Fall.
- Rubinstein, M. and Leland, H. (1981). Replicating option with positions in stock and cash. *Financial Analyst Journal*, 46:61-70.