# A Constant-Volatility Framework for Managing Tail Risk

ALEXANDRE HOCQUARD, SUNNY NG, AND NICOLAS PAPAGEORGIOU

**ALEXANDRE HOCQUARD**
is the director of quantitative strategies and portfolio manager at Pavilion Advisory Group Ltd. in Montréal, QC, Canada.
**alexandre.hocquard@hec.ca**

**SUNNY NG**
is director of research at Morningstar Asia in Hong Kong.
**sunny.ng@morningstar.com**

**NICOLAS PAPAGEORGIOU**
is associate professor of finance at HEC Montreal and director of quantitative research at Pavilion Advisory Group Ltd. in Montréal, QC, Canada.
**nicolas.papageorgiou@hec.ca**

The recent financial crisis reminded us that investing in financial markets is risky business. It also underlined the limitations of conventional, asset allocation–based risk management strategies. The swift, 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 that are uncorrelated in normal times suddenly become correlated; alternative investments, selected for their ability to generate alpha without beta, suddenly deliver high beta with little alpha. The phase-locking behavior that occurred during the most recent crisis, coupled with the jump in market volatility, resulted in dramatic draw-downs for many investors and put the spotlight on risk management. Increasingly, investors are realizing the importance of mitigating tail risk in order to achieve long-term investment objectives. Most investors can withstand an annual loss of 5% or even 10%, but few are able to absorb another draw-down like the one they suffered in 2008.

Historically, plan sponsors have relied on the fixed-income component of their portfolios to provide protection during equity-market draw-downs. A balanced portfolio has proven to be a prudent and effective approach, a[s] Treasuries have consistently generated positiv[e] returns during periods of market correction[s]. But with long-term yields below 3%, ther[e] is clearly very limited upside potential to b[e] found in the Treasury market. Unfortunatel[y], no other asset has offered the same de-couplin[g] with equities during periods of crisis.[1]

Now managers must learn to effectivel[y] cope with systematic risk, specifically tail ris[k] that cannot be diversified away and is increas-ingly unpredictable.

In this article we present a novel, cost[-]effective portfolio management approach th[at] focuses on delivering returns with constant vola[-]tility and without undue exposure to the risk o[f] fat tails.

## TRADITIONAL TAIL-RISK MANAGEMENT TECHNIQUES: PORTFOLIO INSURANCE

An effective tail-risk hedge should pos[-]sess two important characteristics: it must b[e] negatively correlated to asset returns an[d] exhibit convex behavior to the upside durin[g] periods of market stress.

Typically, implementing tail-risk hedgin[g] has involved using equity put options. Unfo[r]tunately, the cost is often prohibitive, creatin[g] a significant drag on portfolio performanc[e]. As an alternative to purchasing put option[s], investors can also resort to dynamic portfol[io] 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, then delta managing the overall exposure. Other dynamic strategies include the notorious constant proportion portfolio insurance, an arguably more robust approach that Black and Jones [1987] and Black and Perold [1992] have proposed.

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 advantages include no broker premiums, no up-front costs, complete flexibility to change, adjust, or remove the hedge, and no exposure to counterparty risk. The fact remains, however, that all portfolio insurance techniques create a significant drag on portfolio performance.

How can investors protect their portfolios against large draw-downs without relinquishing substantial upside? The answer lies in properly understanding and monitoring market volatility.

## RE-THINKING VOLATILITY: BLACK SWANS VERSUS WHITE SWANS

Many researchers and quantitative strategists (including black swan enthusiast Nassim Taleb and his dedicated followers) have long advocated the importance of giving greater consideration to distribution tails, calling attention to the fact that traditional risk management methods typically underestimate tail events' frequency and/or severity. Although the normality assumption of asset returns certainly makes the mathematics a lot easier, it 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 MPT's key assumptions is that asset returns follow a normal distribution with constant volatility. But Exhibit 1, which plots levels of the S&P 500 index and its implied volatility over the last 20 years, clearly shows that risk measured by implied volatility does not remain constant but changes significantly over time. Over the measurement period, the VIX index ranged from less than 10% to a peak of more than 80%.

The most recent financial crisis illustrates how volatility's tendency to vary over time affects MPT, and in particular the associated tail-risk assumptions. Using the average historical volatility of the S&P 500 as our reference point, the October 2008 monthly decline in U.S. equity markets is close to four standard deviations away from its neighbors. Under the common assumption that returns are normally distributed, a gap of four standard deviations has a nearly 1 in 10,000 chance of occurring, implying that a monthly loss of that magnitude should occur approximately once every 750 years. From a statistical point of view, such a rare event is a black swan.

**EXHIBIT 1**
**S&P 500 Index and VIX (1990–2011)**

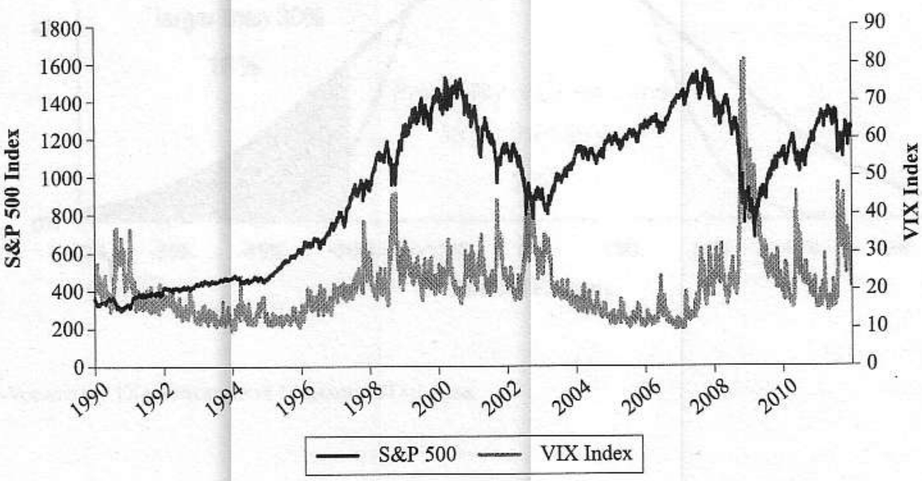

The actual S&P 500 returns over the last 80 years show that October 2008 only ranks ninth in terms of worst monthly performances, implying that such a significant draw-down is much more likely than we imagine. The assumption of normally distributed historical returns clearly underestimates the probability of tail events.

Two possible strategies might better characterize and model the inherent risk in equity returns. The first involves using complex statistical distributions based, for example, on extreme value theory, to help parameterize true tail risk. This is a highly quantitative approach and represents a significant shift away from traditional thinking.

The second, more appealing approach is to re-think how we measure and interpret volatility within a traditional mean–variance framework. We think that the prevailing market volatility level, not the historical level, is the relevant measure. If we use the prevailing volatility level as a reference point, the draw-down in October 2008 is an event that's closer to one standard deviation. October 2008 becomes much less of a black swan—just an undesirable white one.

## THE VOLATILITY OF VOLATILITY: A STORY OF TWO TAILS

When we consider an asset's historical return distribution, the volatility provides us with a measure of the returns' dispersion around the mean. Exhibit 2 illustrates two normal distributions. The light grey line shows a distribution with volatility (standard deviation) of 15%, while the darker line shows a distribution with volatility of 30%. The mean return for the two distributions is the same, but the probability of a large loss (or gain) is significantly higher for the 30% volatility distribution. In fact, the probability of losing 30% or more is approximately eight times higher for the 30% volatility distribution.

The two distributions do not necessarily represent two different assets; in fact, they can represent the same asset at two different points in time. Market conditions change over time, and assets' risk (volatility) profiles also vary.

As market volatility increases, an asset's return distribution flattens and the tails appear to fatten, relative to their average historical distribution. As volatility

**EXHIBIT 2**

**Comparison of Tail Risk for Normal Distributions with Different Volatilities**

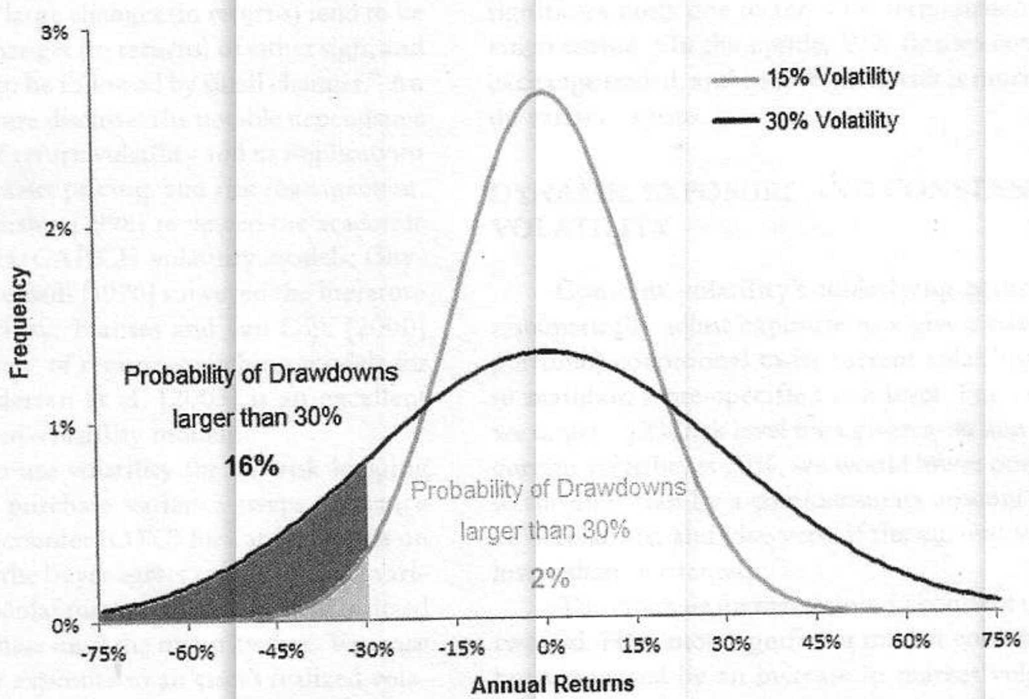

increases, the probability of the asset undergoing large swings also becomes much greater, and historical probabilities no longer represent actual loss potential. The temporal cumulative effect of variable volatility leads to asymmetric tails, especially negative fat tails, in assets' return distributions. Efficient frontier analysis/strategic asset allocation based on a static measure of volatility becomes relatively useless as a risk management tool.

### Using Volatility to Smooth Returns and Manage Tail Risk

Equity volatility is not constant, implying that a portfolio's risk level (and therefore the probability of a large draw-down) is constantly changing. If we can accurately measure the prevailing volatility level and effectively hedge against changes in that volatility, we can greatly reduce tail risk and potentially improve risk-adjusted returns.

Because most assets tend to exhibit volatility clustering, an asset's recent (realized) volatility provides useful information about near-term risks. As Mandelbrot [1963] noted, "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." An abundance of literature discusses the notable dependence and predictability of return volatility and its implications on asset allocation, asset pricing, and risk management. Andersen and Bollerslev [1998] reviewed the academic literature on ARCH/GARCH volatility models; Ghysels, Harvey, and Renault [1996] surveyed the literature on stochastic volatility; Franses and van Dijk [2000] provided an overview of regime-switching models for volatility; and Andersen et al. [2003] is an excellent reference for realized-volatility models.

Investors who use volatility for tail-risk hedging purposes typically purchase 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 purchase until the maturity date. Variance swaps provide pure exposure to an asset's realized volatility. Investors use them to take views on future volatility, capture the spread between realized and implied volatility, and to hedge asset volatility exposure.

A tail-risk hedging strategy would involve purchasing a basket of variance swaps on the market you're hedging. The drawbacks: variance swaps are relatively illiquid, offer limited capacity, are subject to counterparty risk, are priced based on the prevailing implied-volatility level, and involve paying a generous brokerage premium. Most recently, these instruments have garnered much attention, so overwhelming demand on the long side for variance swaps has made them extremely expensive.

Another means of purchasing volatility exposure is to directly trade an implied-volatility index-linked instrument. The VIX Index is the often-quoted implied volatility of traded options on the S&P 500 index. An investor with a long position in the VIX will profit if the level of implied volatility increases over the holding period.

Unlike variance swap returns, a VIX strategy's return is not a function of the underlying asset's realized volatility, but simply the change in implied volatility. The main difficulty in trading implied volatility is that the spot VIX index is not investable, and so investors must acquire VIX exposure through futures contracts (also packaged in VIX ETFs). These contracts have a fixed maturity and therefore must be rolled, resulting in significant costs due to the VIX term-structure's contango nature. On the upside, VIX futures contracts are exchange traded, and counterparty risk is much less than on variance swaps.

## DYNAMIC EXPOSURE AND CONSTANT VOLATILITY

Constant volatility's underlying principle is to systematically adjust exposure to a given asset (or asset portfolio) conditional to its current volatility, in order to maintain a pre-specified risk level. For example, if we target a 12% risk level for a given asset and the asset's current volatility is 20%, we would lower our exposure to the asset class by a commensurate amount to yield a 12% volatility, and vice versa if the current volatility is lower than our target.

The rationale for maintaining a constant volatility is twofold. First, most significant market corrections have been preceded by an increase in market volatility. By conditioning their exposure market volatility, investors can dampen the impact of a market correction. Second, 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, relatively smooth pattern. During these periods, investors should maximize asset exposure, taking advantage of a favorable risk–reward tradeoff. As volatility increases, decrease asset exposure to maintain the desired risk level.

Other authors have proven that using volatility as a risk-conditioning, portfolio-optimizing strategy is extremely efficient. Fleming, Kirby, and Ostdiek [2001, 2003] 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] defined the volatility of volatility "vovo" and identified trading strategies using leveraged ETFs to target the desired risk exposure. The author concluded that constant volatility strategies are able to profit from the upside of leverage, without all the downside. In effect, Cooper [2010] founds that risk smoothing can generate alpha, due to volatility's predictability.

Although these results clearly support the predictability and use of volatility as a conditioning variable, we are still confronted with the problem of translating the prevailing volatility level 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 the distribution's higher moments.

## THE PAYOFF DISTRIBUTION MODEL

Dybvig [1988] introduced the PDM to price and evaluate the distribution of consumption for a given portfolio. The author proposed 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 article we extend the PDM to a more general portfolio- and risk-management methodology. The PDM lets us derive and price any contingent claim on an underlying asset or asset pool.

We use the PDM to solve for the payoff function that provides us with the target return density conditional to the underlying asset's distributional properties. In effect, it provides us with the distortion that must be applied to the underlying asset's distribution 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 monthly payoffs will deliver the specified target density over the long term.

To begin the PDM approach, we derive the monthly payoff structure for the target distribution. For a constant-volatility fund, we target a normal distribution and volatility level. Once the monthly payoff structure is determined, we dynamically adjust the portfolio exposure to the underlying asset to achieve two key objectives: 1) a constant volatility level, regardless of the prevailing volatility level 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 showed 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 function that integrates the asset's entire risk profile.

## MODEL IMPLEMENTATION

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 that implementing the target volatility strategy requires.

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

1. *Define the underlying asset or fund and (if required) its tradable proxies.* 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 volatility level to illustrate the strategy's benefits.

3. *Estimate the daily process of the underlying asset return 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 lets us capture two specific volatility features: 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:

$$\log\frac{S_{Under,t}}{S_{Under,t-1}} = R_t = \mu + \sigma_t \in_t, \quad \in \sim i.i.d.N(0,1)$$

$$\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha(R_{t-1} - \mu)^2, \quad with\ \alpha + \beta < 1$$

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

We estimate the parameters using standard maximum-likelihood maximization. We perform the estimation every month, using all available data.

4. *Derive the targeted distribution's monthly payoff.* The payoff—the function $g$—can be written in closed form, such as:

$$g(x) = \mu_{Target} + \sigma_{Target} * \Phi^{-1}\left[\Phi\left(\frac{x - \mu_{Under}}{\sigma_{Under}}\right)\right],$$

with $x$ the monthly underlying return

where

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

$\mu_{Under}$ is the monthly expected return of the underlying asset, computed as the historical expected return.

$\sigma_{Under}$ is the monthly volatility of the underlying asset, set as the possible levels of forecasted volatility of the underlying asset at the end of the month.

$\sigma_{Target}$ is the targeted monthly volatility that will allow the constant-volatility property.

$\Phi$ is the standard normal cumulative distribution function and $\Phi^{-1}$ its inverse.

5. *Derive the hedging strategy throughout the month.* In essence, the dynamic trading strategy distorts the underlying asset's distribution 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. As a discrete time-hedging strategy, we compute delta surfaces for every trading day during the month. The required exposure is conditional to the underlying asset's GARCH forecasted volatility and cumulative month-to-date performance. For more details on the hedging methodology, see Papageorgiou et al. [2008].

## IMPLEMENTING A CONSTANT VOLATILITY OVERLAY ON A PORTFOLIO

The constant volatility framework can also be implemented on top of an existing asset portfolio. To do so, we must define tradable benchmarks representing the assets in the portfolio and use an overlay of long and short futures contracts to adjust the portfolio's market exposures to target a pre-specified distribution and volatility level.

The overlay does not in any way impact the strategic asset- or manager-allocation decisions or affect the portfolio's alpha component. It simply aims to smooth exposure to market (beta) risk. The strategy can be implemented using exchange-traded futures contracts, eliminating any potential liquidity constraints, capacity constraints, or counterparty risk and offering full transparency with minimal transaction costs.

## RESULTS

We analyzed a globally diversified equity portfolio[2] (40% S&P 500 Index, 35% MSCI EAFE Index, 10% Russell 2000 Index, 10% MSCI EM Index, 5% S&P/TSX 60 Index), one with a 12% target volatility level, selected because it is close to the median volatility level over the sample period. We provide robustness tests with respect to both the specified volatility level and the composition of the underlying portfolio.

Exhibit 3 illustrates the advantages of implementing a 12% constant-volatility strategy on a global diversified equity portfolio. The shaded area displays the cumulative return of the benchmark equity portfolio, while the darker line represents the cumulative return of a 12% constant-volatility portfolio. To better illustrate the evolution of the excess return generated by the constant-volatility strategy, Exhibit 4 displays the cumulative difference between the returns on the equity portfolio and the constant-volatility strategy over the sample period.

Exhibit 5 illustrates the out-of-sample GARCH volatility modeled on the monthly returns of the globally diversified equity portfolio returns and the

## Exhibit 3

**Cumulative Performance of the Equity Portfolio and Constant-Volatility Strategy (1990–2011)**

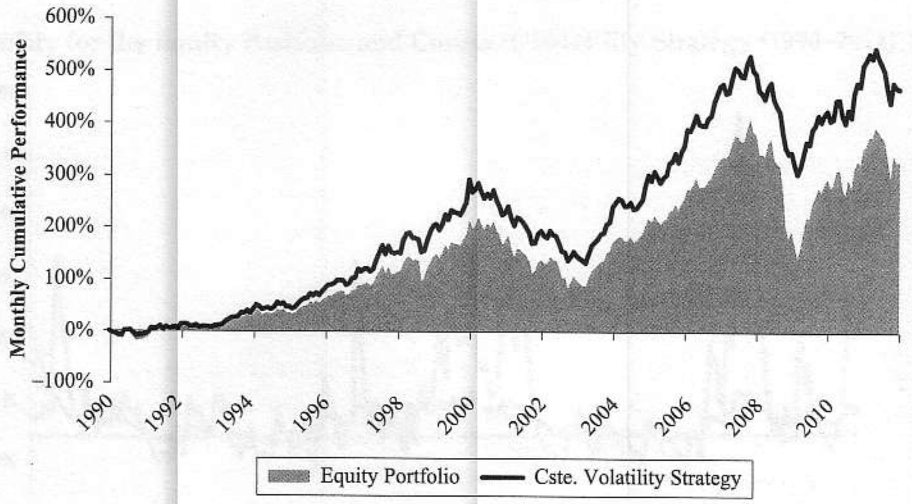


## Exhibit 4

**Cumulative Difference between the Equity Portfolio and Constant-Volatility Strategy (1990–2011)**

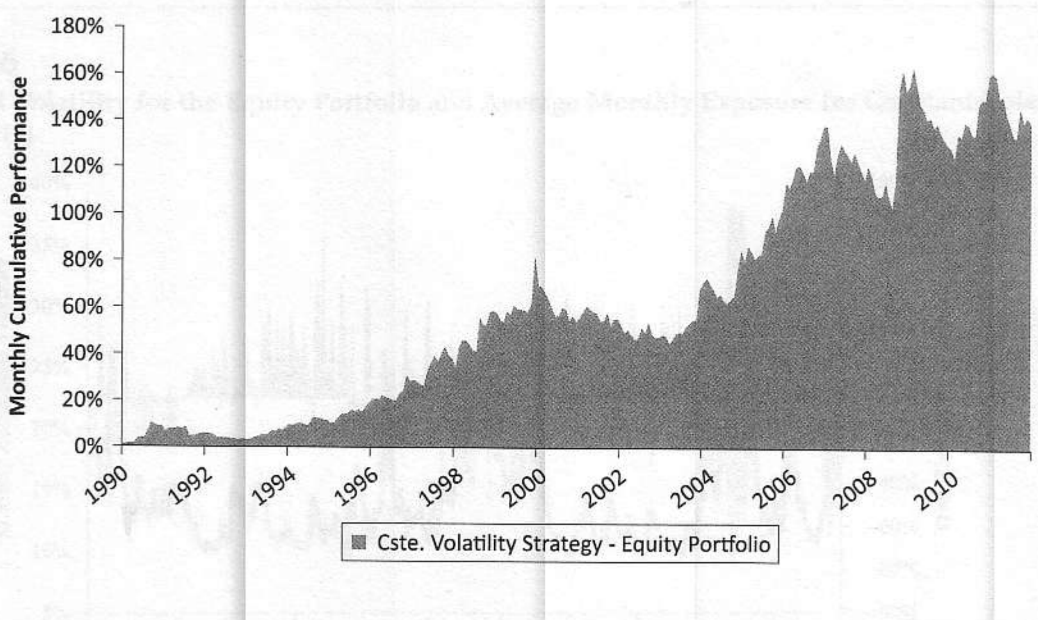

constant-volatility strategy. The constant-volatility strategy's realized volatility is on average slightly above the 12% target shown on the graph but is nowhere near the underlying portfolio's realized volatility.

Exhibit 6 displays the evolution of the constant-volatility strategy's average monthly exposure over the sample period. We overlaid a graph of the volatility to provide some additional intuition into strategy dynamics. During periods of low volatility, the strategy offered an additional 50% of exposure to the underlying portfolio. During the highly volatile months of late 2008, the model removed nearly 80% of the exposure.

## EXHIBIT 5

**Monthly GARCH Volatility for the Equity Portfolio and Constant-Volatility Strategy (1990–2011)**

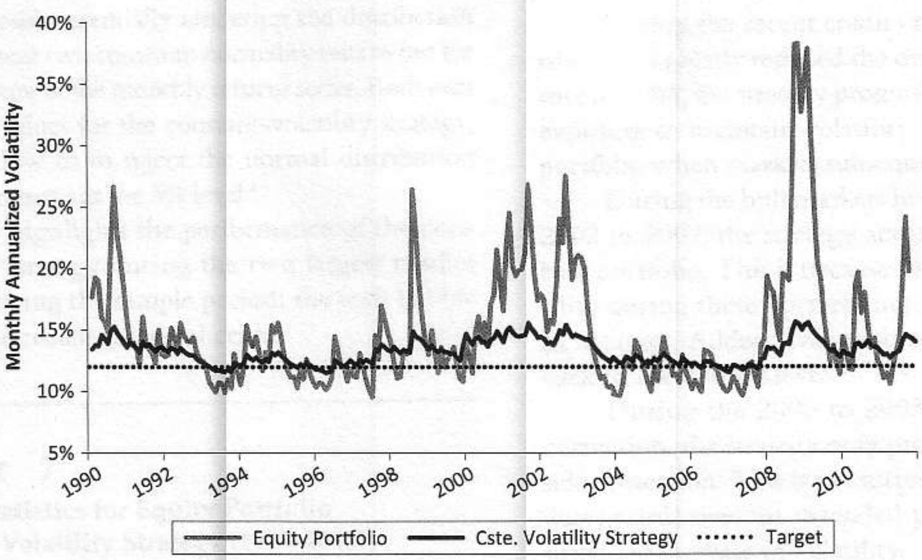


## EXHIBIT 6

**Monthly GARCH Volatility for the Equity Portfolio and Average Monthly Exposure for Constant-Volatility Strategy (1990–2011)**

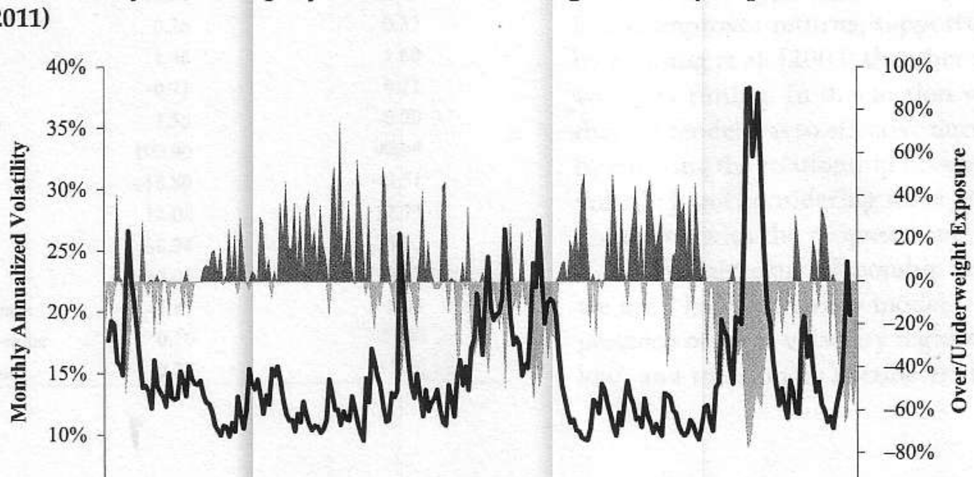

Exhibit 7 summarizes the two strategies' performance. We compute several performance and risk measures, including the Sharpe ratio, omega ratio, and 95% one-month value at risk. In contrast to the Sharpe ratio, the omega ratio introduced by Keating and Shadwick [2002] relaxes the hypothesis that returns follow a Gaussian distribution. This measure leads to a full characterization of the distribution's risk–reward properties by measuring the overall impact of all moments.[3]

The numbers confirm the constant volatility fund's superior risk-adjusted return: the Sharpe ratio increases from 0.26 to 0.37 and the omega ratio goes from 1.46 to 1.60. The worst draw-downs are much smaller, on both a monthly and annual basis. The constant volatility approach essentially eliminates the return distribution's higher moments (skew and excess kurtosis), essentially rendering the distribution normal. We present two common normality tests to test for the Gaussian nature of the monthly returns series. Both tests exhibit high P-values for the constant-volatility strategy, which won't allow us to reject the normal distribution assumption for returns at the 5% level.[4]

Exhibit 8 highlights the performance of the constant-volatility strategy during the two largest market draw-downs during the sample period: the tech bubble collapse and the recent financial crisis.

## EXHIBIT 7
**Descriptive Statistics for Equity Portfolio and Constant-Volatility Strategy (1990–2011)**

| | Base Equity Portfolio | 12% Const. Vol. Strategy |
|---|---|---|
| Ann. Return | 6.84 | 8.22 |
| Ann. Volatility | 15.65 | 13.34 |
| Sharpe Ratio | 0.26 | 0.37 |
| Omega Ratio | 1.46 | 1.60 |
| Skew | –0.71 | 0.21 |
| Excess Kurtoss | 1.55 | –0.20 |
| Correlation | 100.00 | 92.94 |
| Worst Month | –18.80 | –8.51 |
| Best Month | 12.02 | 12.79 |
| Worst Year | –38.94 | –25.07 |
| Best Year | 37.01 | 38.33 |
| 95%VaR 1-Month | –7.85 | –5.43 |
| Jarque-Bera P-value | 0.10 | 25.40 |
| Liliefors P-value | 0.27 | 45.90 |

## EXHIBIT 8
**Performance During Draw-Downs**

| | Tech Bubble (August 2000 to March 2003) | |
|---|---|---|
| | Maximum Draw-down | Months to Recovery |
| Base Equity Portfolio | 41.00% | 27 |
| Cons. Vol. Strategy | 37.32% | 20 |
| | **Financial Crisis (October 2007 to March 2009)** | |
| Base Equity Portfolio | 52.36% | – |
| Const. Vol. Strategy | 35.76% | 24 |

During the recent credit crisis, the constant-volatility fund greatly reduced the draw-down. As volatility rose in 2008, the strategy progressively decreased market exposure to maintain volatility at 12%, protecting the portfolio when markets subsequently plunged.

During the bull markets in the late 1990s and from 2002 to 2007, the strategy actually overperformed the base portfolio. This is because the level of realized volatility during these up-trending markets was below the 12% target. Added leverage brought the risk exposure back to the desired level.

During the 2000 to 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.

## VOLATILITY REGIMES AND ASSET RETURNS

The strategy not only controls risk but also reduces it and improves returns, supporting previous findings by Fleming et al. [2003] that there is economic value to volatility timing. In this section we discuss the reasons that the model was so effective during the sample period, by studying the relationship between market returns and volatility and considering some of the potential risks of associated with the proposed approach.

To analyze the relationship between risk and return, we use a hidden Markov model (HMM) to identify the presence of three volatility regimes (high, medium, and low) and to estimate parameters for the three regimes.

EXHIBIT 9
**Volatility Regimes for the Equity Portfolio (1990–2011)**

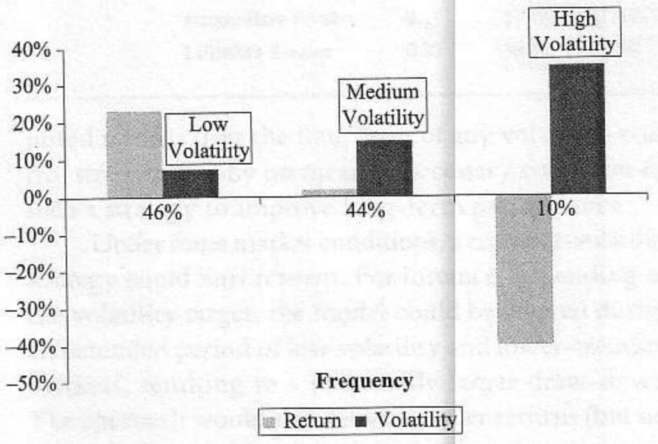


Exhibit 9 presents the annualized average volatility and the corresponding annualized return for each regime.

There is a clear trend across the three regimes. The high-volatility regime, which occurs 10% of the time, produces an average volatility of 35% and an average annualized return of −40%. Markets find themselves in the medium volatility regime 44% of the time, where the average volatility and average annualized return are 14% and 2%, respectively. The low-volatility regime (46% probability) offers by far the best risk–reward tradeoff, with an average volatility of 7% and an average annualized return of 23%.

At first, this relationship between returns and volatility might seem counterintuitive, but it is not inconsistent with financial theory. In modern portfolio theory, expected returns—not actual returns—are related to risk. When risk increases, prices should decline, to offer investors higher expected returns. Actual returns should be low when risk goes up, so that expected returns are higher. Exhibit 10 demonstrates that 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 fact, most of the risk premium provided by equity markets is extracted during these periods of low volatility. In contrast, most major market declines are of short duration and develop rapidly, as fear takes hold of market participants. The sharp draw-downs are therefore much more dramatic, with markedly higher volatility.

Although this relationship between returns and volatility is robust and undoubtedly improves risk-ad-

EXHIBIT 10
**Cumulative Returns and Volatility Regimes for the Equity Portfolio**

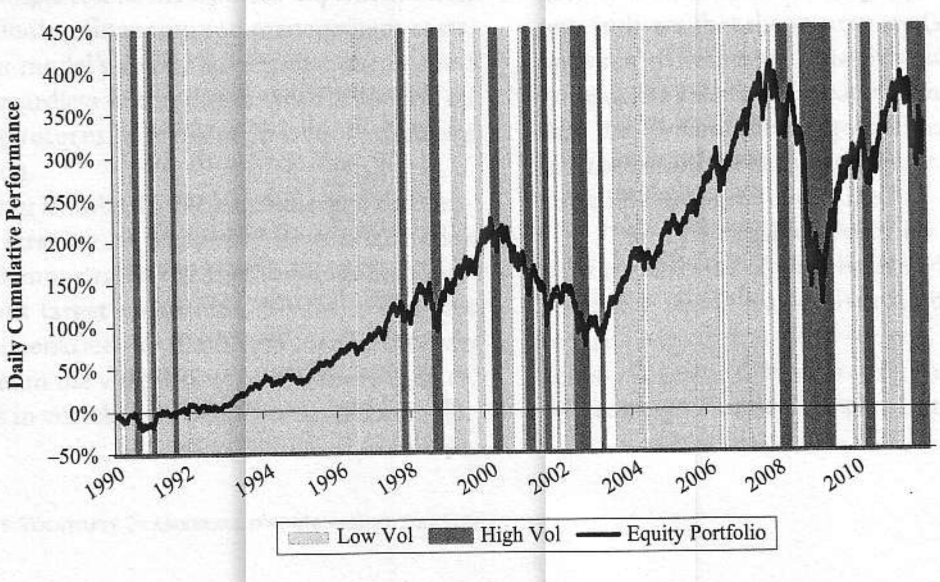

# Exhibit 11

**Constant–Volatility Strategies for Different Target Volatilities (1990–2011)**

| | Base Equity Portfolio | Targeted Volatility | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | 6% | 8% | 10% | 12% | 14% | 16% | 18% | 20% |
| Ann. Return | 6.84 | 5.30 | 6.73 | 7.41 | 8.22 | 9.45 | 9.71 | 11.12 | 11.05 |
| Ann. Volatility | 15.65 | 6.44 | 8.86 | 11.12 | 13.34 | 15.37 | 17.71 | 19.90 | 22.12 |
| 95% VaR 1-Month | −7.85 | −2.53 | −3.48 | −4.51 | −5.43 | −6.24 | −7.23 | −7.97 | −8.99 |
| Jarque-Bera P-value | 0.10 | 31.53 | 21.33 | 26.41 | 25.40 | 21.49 | 24.46 | 13.17 | 12.95 |
| Lilliefors P-value | 0.27 | 38.38 | 37.40 | 50.00 | 45.90 | 50.00 | 50.00 | 29.00 | 50.00 |

justed returns over the long term of any volatility-control strategy, it is by no means a necessary condition for such a strategy to improve long-term performance.

Under some market conditions, a constant-volatility strategy could hurt returns. For instance, depending on the volatility target, the model could be levered during an extended period of low volatility and lower-trending markets, resulting in a potentially larger draw-down. The approach would also deliver smaller returns (but not necessarily on a risk-adjusted basis) if markets increased with high volatility, as the model would not be fully invested.

## ROBUSTNESS TO TARGET VOLATILITY

Exhibit 11 shows the results of a target volatility strategy that targets normal distributions and volatilities, ranging from 6% to 20%, for the global diversified equity portfolio.

The results are robust to different target volatility values, although realized fund volatilities are slightly higher than targeted values. This is because the strategy's monthly profits and losses are reinvested in the fund. These out-of-sample results incorporate implementation constraints, including financing and management costs, and support the model's ability to generate the desired risk profile. Regardless of the target volatility level, all funds' monthly returns pass normality tests when we use this approach.

In analyzing Exhibit 11, we implemented the constant-volatility strategy on a globally diversified index portfolio and demonstrated the methodology's robustness for different target volatilities. Volatility changes for a globally diversified portfolio can come from two sources: changes in the volatility of one or more indices, and/or changes in correlations between the indices. The proposed approach models the portfolio's overall realized volatility, it captures both these sources of volatility shifts. The approach does not seek to disentangle these two factors. Its focus is uniquely on measuring and adjusting exposure in response to changes in portfolio-level risk. This makes the model both parsimonious and robust.

## ROBUSTNESS TO INDICES

To illustrate the effectiveness of the constant-volatility approach to 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 December 2011. Exhibit 12 shows the results. For all assets we target a normal distribution with a 14% monthly annualized volatility.

These results strongly support the benefits of a constant-volatility framework. Return normalization is the most notable transformation to assets' statistical properties when we implement the payoff distribution model. In all cases, both skew and excess kurtosis are greatly reduced; both the Jarque–Bera and Lilliefors tests indicate that the returns are Gaussian. 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.

The 14% constant-volatility funds tend to underperform the underlying indices due to lower realized volatility levels but deliver a superior risk-adjusted return. Maximum draw-downs are greatly reduced across all assets, demonstrating the important role that controlling volatility can play in reducing tail risk.

## Exhibit 12

**Constant–Volatility Strategies across Indices (1990–2011)**

| 14% Volatility Target | S&P 500 Index | | Russell 2000 Index | | TSX 60 Index | | MSCI EAFE Index | | MSCI EEM Index | | NIKKEI 225 Index | | GSCI Index | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Ann. Return | 8.22 | 7.60 | 8.09 | 7.50 | 7.12 | 7.64 | 2.83 | 2.54 | 7.98 | 10.72 | −5.94 | −4.09 | 5.30 | 5.55 |
| Ann. Volatility | 15.22 | 14.93 | 19.88 | 13.81 | 15.70 | 14.20 | 17.82 | 13.43 | 24.26 | 15.39 | 22.15 | 13.81 | 21.36 | 14.47 |
| Sharpe Ratio | 0.34 | 0.31 | 0.30 | 0.31 | 0.24 | 0.28 | 0.03 | −0.04 | 0.28 | 0.49 | −0.34 | −0.52 | 0.17 | 0.18 |
| Omega Ratio | 1.56 | 1.50 | 1.43 | 1.53 | 1.48 | 1.52 | 1.20 | 1.20 | 1.40 | 1.73 | 0.89 | 0.84 | 1.31 | 1.38 |
| Skew | −0.56 | 0.18 | −0.50 | 0.32 | −0.78 | −0.15 | −0.41 | 0.03 | −0.69 | 0.03 | −0.16 | 0.72 | −0.13 | 0.14 |
| Excess Kurtosis | 0.99 | 0.08 | 0.86 | 0.19 | 2.24 | −0.04 | 0.90 | −0.57 | 1.62 | 0.16 | 0.59 | 0.83 | 1.85 | −0.01 |
| Correlation | 100.00 | 92.01 | 100.00 | 91.90 | 100.00 | 93.57 | 100.00 | 93.51 | 100.00 | 95.38 | 100.00 | 92.64 | 100.00 | 93.54 |
| Worst Month | −16.80 | −10.98 | −20.80 | −8.97 | −20.41 | −13.72 | −20.18 | −8.08 | −29.29 | −12.76 | −23.83 | −8.78 | −27.77 | −8.92 |
| Best Month | 11.44 | 14.82 | 16.53 | 15.54 | 12.09 | 11.69 | 15.40 | 9.50 | 17.14 | 14.19 | 20.07 | 14.41 | 21.10 | 13.40 |
| Worst Year | −37.00 | −25.04 | −33.79 | −18.02 | −31.17 | −15.22 | −42.99 | −25.46 | −53.18 | −27.75 | −41.11 | −23.75 | −42.80 | −23.08 |
| Best Year | 37.58 | 63.56 | 47.29 | 38.00 | 34.18 | 36.37 | 39.29 | 34.59 | 78.58 | 61.19 | 41.51 | 42.68 | 50.31 | 32.26 |
| 95% VaR 1-Month | −7.27 | −6.38 | −9.54 | −5.63 | −8.08 | −6.10 | −8.80 | −6.15 | −12.15 | −6.25 | −11.07 | −5.84 | −9.24 | −6.36 |
| Jarque-Bera P-value | 0.12 | 42.91 | 0.25 | 6.76 | 0.10 | 50.00 | 0.39 | 12.61 | 0.10 | 50.00 | 6.63 | 0.10 | 0.10 | 50.00 |
| Lilliefors P-value | 0.63 | 50.00 | 0.20 | 2.38 | 1.82 | 39.58 | 4.48 | 50.00 | 0.10 | 50.00 | 50.00 | 0.14 | 17.52 | 50.00 |

*Note: For each index, the left-hand column shows the result of a direct investment in the asset over the period 1990–2011. The right-hand column shows the output of the 14% constant volatility target fund on the asset.*

## 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. When volatility varies over time, asset returns have been shown to behave in a non-normal fashion, which increases the likelihood of negative tail events for portfolios that maintain static asset allocation. A more cost-effective, prudent approach to managing risk involves actively managing portfolio exposure, according to the prevailing volatility levels within the underlying assets, to maintain a constant risk exposure. Our robust methodology is based on Dybvig's [1988] payoff distribution model and targets a constant volatility level, normalizing monthly returns. This approach to portfolio and risk management can help investors obtain desired risk exposures over the short and long term, reduce tail-risk exposure, and (in general) increase the portfolio's risk-adjusted performance.

## ENDNOTES

The authors wish to acknowledge Pavilion Advisory Group Ltd. for its support as well as the Pavilion Advisory Group Ltd. Quantitative Research and Institutional Consulting teams for their input and comments.

[1]Other asset classes, such as hedge funds, private equity, infrastructure and commodities, have, on occasion, provided some protection against market draw-down, but no evidence supports their inclusion in a portfolio as a effective tail-risk hedge.

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

[3]The omega measure, $\Omega$, involves partitioning returns into loss and gain, above and below a given threshold. The ratio is calculated as:

$$\Omega(r) = \frac{\int_r^{\infty}(1-F(x))dx}{\int_{-\infty}^{r}F(x)dx}$$

where $F$ is the cumulative distribution function and $r$ the threshold that divides the gain from the loss. In our case, $r$ is equal to zero.

[4]For a complete description of the normality tests we used, refer to Jarque and Bera [1980] and Lilliefors [1967].

## REFERENCES


Amin, G., and H. Kat. "Hedge Fund Performance 1990-2000: Do the "Money Machines" Really Add Value?" *Journal of Financial and Quantitative Analysis*, Vol. 38, No. 2 (2003), pp. 251-275.

Andersen, T.G., and T. Bollerslev. "ARCH and GARCH Models." In *Encyclopedia of Statistical Sciences*, edited by S. Kotz, C.B. Read and D.L. Banks, Vol. 2. New York: John Wiley and Sons, 1998.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys. "Modeling and Forecasting Realized Volatility." *Econometrica*, 71 (2003), pp. 579-625.

Black, F., and R. Jones. "Simplifying Portfolio Insurance." *The Journal of Portfolio Management*, 14 (1987), pp. 48-51.

Black, F., and A. Perold. "Theory of Constant Proportion Portfolio Insurance." *Journal of Economic Dynamics and Control*, 16 (1992), pp. 403-426.

Brennan, M.J., and E.S. Schwartz. "Portfolio Insurance and Financial Market Equilibrium." *Journal of Business*, Vol. 62, No. 4 (1989), pp. 455-472.

Cooper, C. "Alpha Generation and Risk Smoothing Using Volatility of Volatility." Working paper, Double Digit Numerics, 2010.

Dybvig, P. "Distributional Analysis of Portfolio Choice." *Journal of Business*, Vol. 61, No. 3 (1988), pp. 369-393.

Fleming, J., C. Kirby, and B. Ostdiek. "The Economic Value of Volatility Timing." *Journal of Finance*, Vol. 56, No. 1 (2001), pp. 329-352.

Fleming, J., C. Kirby, and B. Ostdiek. "The Economic Value of Volatility Timing Using 'Realized' Volatility." *Journal of Financial Economics*, 67 (2003), pp. 474-509.

Hans Franses, P., and D. van Dijk. *Non-Linear Time Series Models in Empirical Finance*. Cambridge, UK: Cambridge University Press, 2000.

Ghysels, E., A. Harvey, and E. Renault. "Stochastic Volatility." In *Handbook of Statistics: Statistical Methods in Finance*, Vol. 14, edited by G.S. Maddala and C.R. Rao, pp. 119-191. Amsterdam: Elsevier Science, 1996.

Jarque, C.M., and A.K. Bera. "Efficient Tests for Normality, Homoskedasticity and Serial Independence of Regression Residuals." *Economic Letters*, 6 (1980), pp. 255-259.

Keating, C., and W.F. Shadwick. "A Universal Performance Measure." Working paper, The Finance Development Centre, London, 2002.

Lilliefors, H.W. "On the Kolmogorov-Smirnov Test for Normality with Mean and Variance Unknown." *Journal of the American Statistical Association*, 62 (1967), pp. 399-402.

Mandelbrot, B. "The Variation of Certain Speculative Prices." *Journal of Business*, Vol. 36, No. 4 (1963), pp. 394-419.

Papageorgiou, N., B. Rémillard, and A. Hocquard. "Replicating the Properties of Hedge Fund Returns." *The Journal of Alternative Investment*, Vol. 11, No. 2 (2008), pp. 8-39.

Rubinstein, M., and H. Leland. "Replicating Option with Positions in Stock and Cash." *Financial Analysts Journal*, 46 (1981), pp. 61-70.


*To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.*