AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

WP 2003-05
February 2003

Working Paper

Department of Applied Economics and Management
Cornell University, Ithaca, New York 14853-7801 USA

Switching Asymmetric GARCH and Options on a Volatility Index

Hazem Daouk and Jie Qun Guo

It is the policy of Cornell University actively to support equality of educational and employment opportunity. No person shall be denied admission to any educational program or activity or be denied employment on the basis of any legally prohibited discrimination involving, but not limited to, such factors as race, color, creed, religion, national or ethnic origin, sex, age or handicap. The University is committed to the maintenance of affirmative action programs which will assure the continuation of such

Switching Asymmetric GARCH and Options on a Volatility Index*

Hazem Daouk[†]

Hazem Daouk is an Assistant Professor of Finance at
Cornell University, Department of Applied Economics
& Management, Ithaca, New York.

Jie Qun Guo

Jie Qun Guo is an Economist at Fannie Mae,
Washington DC.

February 2003

*Our thanks go to Timothy Crack, Partha Deb, Craig Holden, Michael Sullivan, George Wang, Robert Webb (the editor), two anonymous referees, and participants at the 2000 Financial Management Association meeting for helpful discussions and comments. Opinions of the authors are not necessarily those of Fannie Mae. Remaining errors are our own.

[†]Correspondence author, Cornell University, Department of Applied Economics & Management, 446 Warren Hall, Ithaca, NY 14853; e-mail: hd35@cornell.edu.

Switching Asymmetric GARCH and Options on a Volatility Index

Abstract

Few proposed types of derivative securities have attracted as much attention and interest as option contracts on volatility. Grunbichler and Longstaff (1996) is the only study that proposes a model to value options written on a volatility index. Their model, which is based on modeling volatility as a GARCH process, does not take into account the switching regime and asymmetry properties of volatility. We show that the Grunbichler and Longstaff (1996) model underprice a 3-month option by about 10%. A Switching Regime Asymmetric GARCH is used to model the generating process of security returns. The comparison between the switching regime model and the traditional uni-regime model among GARCH, EGARCH, and GJR-GARCH demonstrates that a switching regime EGARCH model fits the data best. Next, the values of European call options written on a volatility index are computed using Monte Carlo integration. When comparing the values of the option based on the Switching Regime Asymmetric GARCH model and the traditional GARCH specification, it is found that the option values obtained from the different processes are very different. This clearly shows that the Grunbichler-Longstaff model is too stylized to be used in pricing derivatives on a volatility index.

Switching Asymmetric GARCH and Options on a Volatility Index

Few proposed types of derivative securities have attracted as much attention and interest as option contracts on volatility. The objective of this paper is to show that pricing options on a volatility index demands extreme care in modeling the underlying process. Grunbichler and Longstaff (1996), or G-L, is the only paper that proposes a model to value options on a volatility index. Their model is based on a GARCH specification for the underlying asset. It is argued that the GARCH specification leads to severe mispricing in the option value because it does not allow for regime switching and volatility asymmetry (leverage effect). No volatility options data is available, since those derivatives are not traded yet. Hence, we cannot use real prices to test the option pricing models. For that reason, we use simulation to obtain our results. Nevertheless, it is argued that it is more useful to explore and test option pricing models before the derivatives are introduced and traded, rather than after they are already priced in the market.

Ever since Mandelbrot (1963) suggested that “large changes tend to be followed by large changes - of either sign - and small changes tend to be followed by small changes”, a number of papers have illustrated that stock variance exhibits temporal dependence (for example, Akgiray, 1989). To capture this feature, Engle (1982) develops a new class of stochastic processes, the AutoRegressive Conditional Heteroscedasticity (ARCH) models, in which the conditional variance is a function of past residuals. It has been shown that allowing the conditional variance to be time-varying generates fat tails for the distribution of returns. Bollerslev (1986) introduces the Generalized ARCH (GARCH) which extends Engle’s ARCH by allowing the conditional variance to be a function of the lagged variance. Bollerslev (1986) shows that the GARCH model allows a better representation of the volatility process while being more parsimonious.

Another important stylized fact of stock returns is that their volatility tends to rise

in response to “bad news” (excess returns lower than expected) and fall in response to “good news” (excess returns higher than expected). This is also known in the literature under the name of “volatility asymmetry” or “leverage effect”, since leverage provides at least a partial explanation. As the price of a stock decreases, the firm leverage will increase, assuming that debt value did not decrease by as much. Since the volatility of equity is an increasing function of leverage, this provides a simple explanation for asymmetry (Black, 1976). Another explanation was given by French, Schwert, and Stambaugh (1987). They argue that the negative relation between volatility and stock prices is due to the fact that an increase in unexpected volatility will increase the expected future volatility (assuming persistence). Then, the risk premium will also increase causing prices to drop. Nelson (1991) introduces the EGARCH model (Exponential GARCH) to account for volatility asymmetry. Also, Glosten, Jagannathan, and Runkle (1993) introduce an alternative specification known as the GJR-GARCH (see Bollerslev, Chou, and Kroner, 1992 for an extensive review of the literature).

Recently, Cai (1994) and Hamilton and Susmel (1994) introduce the Markov-Switching Regime ARCH (SWARCH) model that incorporates the features of both Hamilton’s (1988, 1989) switching regime model and Engle’s ARCH model. Hamilton’s switching regime model can be described as an autoregressive process in which the parameters of the autoregression can change as the result of a regime-shift variable. The regime itself will be described as the outcome of an unobserved Markov chain. The basic idea is to allow the data generating process to be different in different regimes. The existence of multiple regimes is a pervasive phenomenon in financial economics. More recently, Switching-regime models have been applied to stock market returns (Norden and Schaller, 1993), foreign stock markets (Rockinger, 1994, Norden, 1995, and Bollen, Gray and Whaley, 2000), and interest rates (Gray, 1996a, 1996b, Ang and Bekaert, 1998, Dahlquist and Gray, 2000).

In this paper, we extend the SWARCH model to a Markov Switching Regime Asym-

metric GARCH (SW-Asymmetric GARCH). The SW-Asymmetric GARCH allows both regime switching in volatility and asymmetry. It is argued that the additional effort required in estimating the model is justified by the benefit of providing a richer modeling of the volatility dynamics.

The remainder of this paper is organized as follows. The next section describes volatility indices and options written on volatility indices. The following section presents a brief exposition of the Grunbichler and Longstaff (1996) model, and shows that the assumed process for volatility is a continuous time limit of a GARCH model. Next we present the general methodology used in assessing if the volatility process in G-L is appropriate. The next section provides detailed discussion of the switching regime model and the model selection criteria. The following section contains results from fitting the volatility models, and rejects the GARCH specification in favor of those allowing for switching regime and asymmetry. Next we show that option prices computed with the G-L assumed process underprice volatility index options by a considerable amount. Then, we provide further analysis and robustness checks. The final section concludes.

1. Options on a Volatility Index

Few proposed types of derivative securities have attracted as much attention and interest as option contracts on volatility. Grunbichler and Longstaff (1996) cite various sources including Reuters, the Wall Street Journal, Futures, Futures and Options World, Barrons, etc., that discuss those derivatives. The underlying asset for those options is the Market Volatility Index (VIX). VIX is an average of S&P 100 option (OEX) implied volatilities. It was introduced by the Chicago Board Options Exchange (CBOE) to provide an-to-the-minute measure of expected volatility of the US equity market¹.

¹Eight near-the-money, nearby, and second nearby OEX call and put options provide the inputs for constructing the volatility index. First, the implied volatilities of the call and the put in each of the four categories of options are averaged. Second, "at-the-money" implied volatilities are created by interpolating between the nearby and second nearby average implied volatilities. In this way, VIX is

According to Duke Chapman, CBOE former chairman, the index is particularly suited for institutional investors interested in hedging future, not past changes in portfolio risk (see Vosti, 1993). Robert Whaley, the developer of the index declared that “the availability of an option contract on the CBOE Market Volatility Index will allow market participants to manage the risks associated with unanticipated changes in volatility.”² The CBOE has applied to the Securities and Exchange Commission to trade European options on the Volatility Index. Yet, this operation was constantly delayed. The reason for that seems to be that market makers are still trying to understand the implications and risks involved in trading volatility derivatives. More recently, many European exchanges have introduced volatility indices in the hope of trading volatility derivatives (two examples are the VX1 and VX6 volatility indices from MONEP in France and the VLEU volatility index in Switzerland). Some exchanges have even started to trade futures on volatility indices (VOLAX futures from the Deutsche Terminborse in Germany, SEK volatility futures for the Swedish market, and FTSE volatility futures in the United Kingdom). A better understanding of volatility options, and the way they should be priced, will facilitate the introduction and adoption of volatility options worldwide. This paper contributes in this direction.

2. The Grunbichler and Longstaff (1996) Model

Grunbichler and Longstaff (1996), or G-L, is the only paper to our knowledge that presents a model to price options on a volatility index. The major assumption of their model is that related to the specification of the dynamics of the underlying process, i.e.

always “at-the-money.” Finally, a constant time to expiration is maintained. The nearby and second nearby at-the-money volatilities are weighted to create a constant 30-calendar day (22-trading day) time to expiration.

²Brenner and Galai (1989) and Whaley (1993) provide excellent discussions of how volatility derivatives can be used to hedge the volatility risk of portfolios containing options or securities with option-like features. See also Fleming, Ostdiek and Whaley (1995) and Whaley (2000) for a thorough analysis of VIX.

volatility. The volatility index, denoted V , is modeled as a mean reverting process

$$dV = (\omega - \kappa V)dt + \sigma\sqrt{V}dZ, \quad (2.1)$$

where ω, κ , and σ are constants, and Z is a standard Wiener process.

Since V is not the price of a traded asset, it is assumed that the expected premium for volatility risk is proportional to the level of volatility, ζV where ζ is a constant parameter.

The G-L model captures many of the observed properties of volatility. In particular, the model allows volatility to be mean reverting and conditionally heteroskedastic.

We argue below that the G-L specification for volatility is indeed the continuous time limit of a GARCH model. Nelson (1990) investigates the convergence of stochastic difference equations for volatility models to stochastic differential equations as the length of the discrete time intervals between observations goes to zero. We present the results from Nelson's (1990) derivation of the continuous time limits of a GARCH(1,1) in our own notation.

A GARCH(1,1) model can be written as follows

$$h_t = \alpha + \delta_1 h_{t-1} + \gamma \varepsilon_{t-1}^2. \quad (2.2)$$

$$\begin{aligned} \text{where } \varepsilon_t &= h_t^{\frac{1}{2}} z_t, \\ z_t \text{ i.i.d.} &\sim N(0, 1). \end{aligned} \quad (2.3)$$

h_t is conditional volatility, and α, δ_1 , and γ are constant parameters.

Nelson (1990) derives diffusion limits of the form

$$dh = (\varphi - \varpi h)dt + \vartheta h dZ, \quad (2.4)$$

where φ, ϖ , and ϑ are constant parameters, and Z is a standard Wiener process.

As can be seen, the derived model in (2.4) is essentially the G-L specification. In fact, GARCH models do imply mean reversion in volatility (the coefficients need to sum to less than one for stationarity) and conditional heteroskedasticity (the change in conditional volatility is larger for larger levels of volatility).

3. Is the Grunbichler and Longstaff (1996) Volatility Process Appropriate?

The GARCH specification for the volatility process suffers from two major drawbacks. The first one is that it does not allow for volatility asymmetry (leverage effect). The second drawback is that it does not account for regime switching (regime shifts). The objective of this paper is not to criticize the specific form of the volatility model in G-L. Rather, we want to show that any model not accounting for the switching regime and asymmetry properties will severely misprice (undervalue) volatility options. To that end, we compare option prices computed using a simple GARCH process to those computed using switching and asymmetric models³.

In the literature, many studies have derived closed form solutions for stock options with stochastic volatility (Hull and White, 1987, Stein and Stein, 1991, Heston, 1993). Those models came in response to the empirical evidence related to time-varying volatility. Also, stock option pricing models where stock volatility follows a GARCH process have been derived by Duan (1995) and Heston and Nandi (1999). More recently, studies like Ritchey (1990), Billio and Pelizzon (1996), Bollen (1998), Duan, Popova and Ritchken (1999), and Campbell and Li (1999) develop models obtained when the stock dynamic is modeled by a Markov regime switching process. The regimes are character-

³The options we are interested in are written on a volatility index implied from options. The volatility processes we simulate to compute the option prices represent the instantaneous (or a discrete approximation to the instantaneous) volatility of the S&P 100 index. As Grunbichler and Longstaff (1996) note, the volatility can be either the instantaneous volatility or the volatility implied from some option pricing model - the distinction does not affect the form of the resulting valuation procedure. For example, Taylor and Xu (1994) show that when volatility is stochastic, the volatility estimate implied by inverting the Black-Scholes model is nearly a linear function of the actual instantaneous volatility.

ized by different volatility levels. Those models differ from the previous ones in that rather than permitting volatilities to follow a continuous time, continuous state process, their focus is on cases where volatilities can take on a finite set of values, and can only switch regimes at finite times. There is still very little work on volatility asymmetry as applied to option pricing. Along those lines, Schobel and Zhou (1998) use Fourier inversion techniques to price stock options allowing for correlation between instantaneous volatilities and the underlying stock returns. This is a first step towards accounting for volatility asymmetry⁴.

Many of the above studies are able to show that their model present an improvement over Black and Scholes (1973) that is statistically significant. However, this improvement is not economically large. On the other hand, in our case, the model in Grunbichler and Longstaff (1996) will be shown to misprice option on volatility by a very large amount. There are reasons to believe that misspecifying the volatility process when the option is written on a *stock index* is much less serious than in the case where the option is written on a *volatility index*. In the first case, the volatility process is a major determinant of the underlying asset process. In the second case, the volatility process *is* the underlying asset. Ignoring regime switching in the volatility process might not be too bad of an approximation when pricing derivatives on a stock index, as long as volatility is allowed to be stochastic. However, It is argued that ignoring regime switching and asymmetry in volatility when pricing derivatives on a volatility index could be disastrous.

None of the above models can be used to price options on volatility. The only model that can do that is that of G-L. Their model assumes that volatility follows a stochastic process that is the continuous time analogue of GARCH. As argued before, this process does not allow volatility asymmetry. Also, it does not allow the *parameters* of the autoregressive process to be different for different regimes. We note that this

⁴Also, Walsh (1999) examines the effect of ignoring conditional volatility on the price of stock options. He considers many GARCH models including asymmetric ones. He finds that asymmetric GARCH effects exacerbate the bias in the price of the options.

is a shortcoming that is common to all the option models reviewed above. Given the large empirical evidence related to volatility asymmetry and regime switching, there are reasons to suspect that volatility options priced under the GARCH or other “simple” stochastic volatility assumptions will be significantly mispriced. It will be demonstrated that the mispricing in the case of options on a volatility index is very large economically.

There are no analytic solutions for the price of options on a volatility index under the assumption of volatility following a switching asymmetric autoregressive process. Therefore, Monte Carlo integration is used to compute the option values. Also option values for the G-L model are computed using Monte Carlo integration instead of using analytic solutions. There are two reasons for that choice.

First, we want to conduct hypothesis testing from fitting the different processes to the S&P 100 data. By using a GARCH process for the G-L model, we are able to compare its fit to the data with that of other more rich volatility processes. As will be shown in the next section, the GARCH model can be nested in the switching asymmetric GARCH models, allowing for hypothesis testing for nested models (likelihood ratio). This will allow us to test the hypothesis that the GARCH process is an adequate representation of the data as compared to the more general processes.

The second reason why option prices for the G-L model are computed using Monte Carlo integration is that the GARCH specification is widely used in the literature and has been found to be flexible enough to fit the data well. The point is that using the simple GARCH implies accepting the restrictive assumptions of no switching or asymmetric properties in volatility. Those assumptions might be convenient to work with in some contexts. However, we want to show that those assumptions are not appropriate to use to price options on volatility.

4. Switching Regime Asymmetric GARCH Models

4.1. Model Specification

The Switching Regime Asymmetric GARCH models used in this paper are part of the family of Markov switching regime models (See chapter 22 of Hamilton, 1994). A common finding using high-frequency financial data concerns the apparent persistence implied by the estimates for the conditional variance functions. This has prompted Bollerslev and Engle (1986) to introduce the class of integrated GARCH or IGARCH. However, the intuitive appeal and fit of this model has raised questions. Cai (1994) and Hamilton and Susmel (1994) have addressed the issue by proposing Markov switching volatility models. They note that the persistence implied by the model in Bollerslev and Engle (1986) is difficult to reconcile with the poor forecasting performance. Further, it is argued that the poor forecasting performance and spuriously high persistence might be related to structural change in the volatility process. This is related to Perron's (1989) observation that changes in regime may give the spurious impression of unit roots in characterizations of the level of a series. For example, Cai (1994) and Hamilton and Susmel (1994) find that volatility of returns appears to be much less persistent when one models changes in the parameter through a Markov switching process. The SWARCH model allows occasional shifts in the asymptotic variance of the ARCH process by having the intercept term "switch" between different values depending on what volatility regime the process is in. The traditional ARCH-type models rely on the assumption of parametric homogeneity. This means that the parameters of the model remain the same whether the market is in a calm regime or a volatile regime. What distinguishes periods of high volatility from periods of low volatility is the size of the residual. On the other hand, the SWARCH model allows the relation between the conditional volatility and the explanatory variables (lagged residuals) to be different between periods of high volatility and periods of low volatility. Kim and Kon (1999) provide evidence that the time-series

properties of stock returns include both structural change and time dependence in the conditional variance. The estimation of a SWARCH model determines the parameters of the ARCH terms (which include different parameters for different regimes) along with transition probabilities between the different regimes. Cai (1994) and Hamilton and Susmel (1994) show that their model is better able to capture the behavior of returns volatility than previous ARCH-type models. Since then, the SWARCH model has been applied to international equity markets (Susmel, 1994, 1998a, and Fornari and Mele, 1997), and inflation time series (Susmel, 1998b). We follow Cai (1994) and Hamilton and Susmel (1994) in adopting a Markov switching regime model. We believe that this will account for volatility persistence in an appropriate manner. Cai (1994) and Hamilton and Susmel (1994) propose and estimate a Switching-ARCH. The dynamic lag structure, was restricted to an ARCH specification. This is a drawback since the GARCH specification has been shown to be a better fit to financial data. We generalize the SWARCH model to a GARCH specification. We are able to account for volatility asymmetry by using an Asymmetric GARCH. Furthermore, we allow the persistence of volatility and the asymmetry to be different between different regimes. This means that when there is a change in the regime, the model that is used to generate conditional volatility is replaced by a different model with different parameters.

Formally, let R_t be a vector of observed variables and let s_t denote an unobserved random variable that can take on the values $1, 2, \dots, C$. Suppose that s_t can be described by a Markov chain,

$$\text{Prob}(s_t = j \mid s_{t-1} = i, s_{t-2} = k, \dots, \Omega_t) \quad (4.1)$$

$$= \text{Prob}(s_t = j \mid s_{t-1} = i) = \pi_{ij},$$

for $i, j = 1, 2, \dots, C$. It is sometimes convenient to collect the transition probabilities

in a $(C \times C)$ matrix:

$$\Pi = \begin{bmatrix} \pi_{11} & \pi_{21} & \dots & \pi_{C1} \\ \pi_{12} & \pi_{22} & \dots & \pi_{C2} \\ \vdots & \vdots & \dots & \vdots \\ \pi_{1C} & \pi_{2C} & \dots & \pi_{CC} \end{bmatrix} \quad (4.2)$$

Note that each column of Π sums to unity.

A useful representation for a Markov chain is obtained by letting ξ_t denote a random $(C \times 1)$ vector whose j th element is equal to unity if $s_t = j$ and whose j th element equals zero otherwise:

$$\xi_t = \begin{cases} (1, 0, 0, \dots, 0)' & \text{when } s_t = 1 \\ (0, 1, 0, \dots, 0)' & \text{when } s_t = 2 \\ \vdots & \\ \vdots & \\ (0, 0, 0, \dots, 1)' & \text{when } s_t = C. \end{cases} \quad (4.3)$$

The variable s_t is regarded as the state or regime that the process is in at date t . By this we mean that s_t governs that parameters of the conditional distribution of R_t . The density of R_t conditional on its own lagged values as well as on the current and previous q values for the state is of the form,

$$f(R_t | s_t, s_{t-1}, \dots, s_{t-q}, \Omega_t),$$

The methods developed in Hamilton (1989) are used to evaluate the likelihood function for the observed data and make inferences about the unobserved regimes. In this case, the returns R_t follow a process in the ARCH family whose parameters depend on the unobserved realization of $s_t, s_{t-1}, \dots, s_{t-q}$. The objective is to select a parsimonious representation for the different possible regimes.

For a 2-regime switching model, C is equal to 2. In other words, the volatility of security returns are assumed to be generated from two regimes which are different in describing the underlying behavior.

Bera and Higgins (1993) find that with a parametric specification of the conditional variance function, there does not appear to be a density function that is suitable for all security data. It is of interest to see if after allowing for regime switching, a normal distribution for R_t will be adequate to fit the data⁵.

Since we focus on the influence of volatility on the price of options, we use the simplest specification for the mean equation

$$R_t = \mu + \varepsilon_t, \quad (4.4)$$

where $\varepsilon_t = h_t^{\frac{1}{2}} z_t$, z_t i.i.d. $\sim N(0, 1)$. We use different specifications for the conditional variance equation.

For GARCH(1,1),

$$h_t = \alpha + \delta_1 h_{t-1} + \gamma \varepsilon_{t-1}^2. \quad (4.5)$$

For EGARCH(1,1),

$$\ln(h_t) = \alpha + \delta_1 h_{t-1} + \gamma_1 |z_{t-1}| + \gamma_2 z_{t-1}. \quad (4.6)$$

For GJR-GARCH(1,1),

$$h_t = \alpha + \delta_1 h_{t-1} + \gamma_1 \varepsilon_{t-1}^2 + \gamma_2 S_{t-1} \varepsilon_{t-1}^2, \quad (4.7)$$

where $S_t = 1$ if $\varepsilon_{t-1} < 0$ and 0 otherwise.

As can be seen, h_t is a function of Ω_t . That is, the conditional variance is time dependent and relies on past information. The (1,1) formulation ($p = q = 1$) is used because it has consistently been found to model variance dynamics adequately (e.g. Bollerslev, 1986, Akgiray, 1989). Moreover, a natural specification is needed to avoid overfitting the data. The focus is on highlighting the switching and asymmetry properties of the

⁵One of the benefits of the normality assumption is that the parameter estimates are asymptotically unbiased, even if the true distribution is non-normal, provided the first two conditional moments are correctly specified (Weiss, 1986). On the other hand, the t-distribution does not have this robust property.

data. The EGARCH and GJR-GARCH specifications allow for volatility asymmetry (leverage effect). If there exists volatility asymmetry consistent with previous literature, then we should find that $\gamma_2 < 0$ in EGARCH(1,1) and $\gamma_2 > 0$ in GJR-GARCH(1,1).

4.2. Model Estimation and Selection

The switching regime model is based on the discrete heterogeneity assumption. The volatility of security returns has been shown to be persistent. However, this relationship behaves differently between different regimes (see Cai, 1994). The underlying reasons that causes this different behavior are unobservable but are highly related to regime conditions. Our model is able to classify heterogeneity into a small finite number of relatively homogeneous regimes (a calm regime and a volatile regime)⁶.

One of the difficulties in using switching regime models is deciding on how many regimes are needed in the model. In practice, one can start with a small number of regimes and add more provided the fit of the model is significantly improved. In this paper, we use a two regime switching model because of the extreme difficulty in estimating a model with a larger number of regimes⁷.

Let θ denote all unknown parameters, $\theta = (\Pi, \mu, \delta, \gamma)$. Let T be the sample size. Let $\text{Prob}(s_t = j | \Omega_t; \theta)$ denote the inference about the value of s_t based on data obtained through date t and based on a "guess" of the population parameters θ (starting values). Collect these conditional probabilities $\text{Prob}(s_t = j | \Omega_t; \theta)$ for $j = 1, 2, \dots, C$ in a $(C \times 1)$ vector denoted $\hat{\xi}_{t|t}$.

Also, forecasts of how likely the process is to be in regime j in period $t+1$ are formed given observations obtained through date t . These forecasts are collected in a $(C \times 1)$

⁶The switching regime model is considered to be semi-parametric because knowledge of the distribution of the unobserved mixing variable is not required. Hence, we need not know which factors determine Π , the matrix of transition probabilities.

⁷However, we are comfortable with that choice because in an earlier version of this paper, we compared a more restrictive version of our model (i.i.d. mixtures) with three regimes and two regimes respectively. We found that the model with two regimes was a better fit to the data.

vector $\widehat{\xi}_{t+1|t}$, which is a vector whose j th element represents $\text{Prob}(s_{t+1} = j | \Omega_t; \boldsymbol{\theta})$.

The optimal inference and forecast for each date t in the sample can be found by iterating on the following pair of equations:

$$\widehat{\xi}_{t|t} = \frac{(\widehat{\xi}_{t|t-1} \odot \eta_t)}{1'(\widehat{\xi}_{t|t-1} \odot \eta_t)}, \quad (4.8)$$

$$\widehat{\xi}_{t|t+1} = \Pi \cdot \widehat{\xi}_{t|t}. \quad (4.9)$$

Here $1'$ represents a $(1 \times C)$ vector of 1s, and the symbol $\odot$ denotes element-by-element multiplication. Given a starting value $\widehat{\xi}_{1|0}$ and an assumed value for the population parameter vector $\boldsymbol{\theta}$, one can iterate on the above two equations for $t = 1, 2, \dots, T$ to calculate the values of $\widehat{\xi}_{t|t}$ and $\widehat{\xi}_{t+1|t}$ for each date t in the sample. The log likelihood function $\mathcal{L}(\boldsymbol{\theta})$ for the observed data Ω_T evaluated at the value of $\boldsymbol{\theta}$ that was used to perform the iterations can also be calculated as a by-product of this algorithm from

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(R_t | \Omega_{t-1}; \boldsymbol{\theta}), \quad (4.10)$$

where

$$f(R_t | \Omega_{t-1}; \boldsymbol{\theta}) = 1'(\widehat{\xi}_{t|t-1} \odot \eta_t). \quad (4.11)$$

Robust standard errors are computed to account for heteroskedasticity and serial correlation. The likelihood ratio (LR) test, penalized AIC and BIC are used as selection vehicles to compare models.

The LR test is a popular test for multiple parameter situations (see Kon, 1984). The following models were tested against each others: GARCH, Switching Variance, GJR-GARCH, EGARCH, SW-GARCH, SW-GJR-GARCH, and SW-EGARCH.

The model corresponding to the null hypothesis (H_0) is nested within the model corresponding to the alternative (H_1). If H_0 is true,

$$-2\mathcal{L}(\theta) = -2[\mathcal{L}(\theta_{H_0}) - \mathcal{L}(\theta_{H_1})], \quad (4.12)$$

has an asymptotic chi-square distribution with degrees of freedom equal to p (the difference in the number of parameters between H_0 and H_1). However, the method may not be appropriate for the switching model because the hypothesis is on the boundary of the parameter space and this violates the standard regularity condition for Maximum Likelihood. Böhning, Dietz, Schlattmann and Lindsay (1994) show that using LR test is likely to underreject the false null when the boundary condition is involved. Hence, the test tends to choose a small number of components. As alternative methods, we use penalized Akaike Information Criterion ($AIC = -2\mathcal{L}(\theta) + 2K$) and Bayesian Information Criterion ($BIC = -2\mathcal{L}(\theta) + K\ln(N)$). These methods are robust in a sense that they are valid even when the model is misspecified. The model with the lowest value of AIC or BIC is considered the best⁸.

5. Estimation Results and Hypotheses Testing

We use daily continuously compounded S&P 100 index (SP100) returns from the period January 3, 1980 to March 26, 1999. Summary statistics are reported in Table I.

INSERT TABLE I ABOUT HERE

The t-test for skewness equal to zero is rejected at the 5% significance level, showing that the distribution is significantly skewed. The table also demonstrates the existence of excess kurtosis. The Ljung-Box statistic is used to check for serial dependence in the conditional moments. The Ljung-Box statistic and first order autocorrelation tests show significant serial correlations between squared residuals⁹.

⁸Sometimes, the conclusions drawn from AIC and BIC are not the same. This can happen because BIC is more severe in penalizing non-parsimonious specifications than AIC.

⁹We test the standard deviation series of the S&P 100 index for non-stationarity. Using an Augmented Dickey-Fuller test with 6 lags, we obtain a p-value of 0.000. Using a Phillips-Perron test with 6 covariance lags, we get a p-value of 0.000. These results indicate that the hypothesis of non-stationarity is strongly

Model selection tests for uni-regime and switching regime models among GARCH, EGARCH, GJR-GARCH and "Switching Variance" are conducted. The "Switching Variance" model is one where the variance is fixed (not autoregressive) in each of the regimes. In other words, the variance switches between two endogenously estimated values depending on which regime is in place. The information criteria are implemented to choose the model with the best fit. The results from the tests are listed in Table II.

***INSERT TABLE II ABOUT HERE**

According to panel A (LR test), the results unanimously reject the uni-regime hypothesis in favor of the multi-regime model in all cases. Also, the results reject the GARCH specification in favor of the asymmetric models. Moreover, the simple Switching Variance model is rejected in favor of the switching regime autoregressive variance models (GARCH and asymmetric GARCH). The penalized AIC and BIC for all models are presented in table II panel B. The results confirm that the uni-regime models and the symmetric models are not an adequate representation of the data generating process for volatility. The Markov Switching Regime EGARCH or SW-EGARCH represents the best fit among the seven models. Also, the switching regime property seems to be more important than asymmetry. This is because the model that has switching regime but no asymmetry (SW-GARCH) is found superior to those with leverage but no switching regime (EGARCH and GJR-GARCH).

The parameter estimates for the preferred SW-EGARCH model are shown in Table III¹⁰.

rejected.

¹⁰To check if the estimated model is stable, simulated data (conditional volatility) from the model are generated and tested for non-stationarity. The tests are made over 1000 replications. Using an Augmented Dickey-Fuller with 6 lags, the largest p-value over the 1000 replications is 0.000. Using

***INSERT TABLE III ABOUT HERE**

Bollerslev, Chou and Kroner (1992) suggest that the estimated volatility asymmetry may be attributed to a few outliers. One advantage of the switching regime specification is that it allows one component to adjust to these outliers, with the remaining “normal” observations being estimated by another component. In this case, the influence of the outliers is isolated. Table III shows that γ_2 is significantly negative in the tranquil regime. The coefficient γ_2 is large in the volatile regime, albeit not statistically significant because of the large standard error. This seems to indicate that volatility asymmetry is present in both the tranquil and volatile regimes. δ_1 is larger in the tranquil regime than in the volatile regime, which seems to indicate that the tranquil regime exhibits more persistence than the volatile regime. This is consistent with literature that shows that removing exceptionally volatile observations (the 1987 crash for example) increases the persistence parameters of ARCH models (see for example Blair, Poon, and Taylor, 1998).

The standardized residual ($\hat{\varepsilon}_t/\sqrt{\hat{h}_t}$) diagnostics are reported at the bottom of Table III. If the model is correctly specified, the null hypothesis of zero autocorrelation of standardized residuals and squared standardized residuals should not be rejected. We indeed find that $Q(20)$ and $Q^2(20)$ are not significantly different from zero. The results demonstrate that there is no evidence of autocorrelation in the squared residuals, which suggests that the SW-EGARCH model provides an adequate description of the data.

The Jarque-Bera normality tests still reject the null of normality of standardized residuals. However, comparing with the raw residuals in Table I, we can see that there

a Phillips-Perron with 6 covariance lags, the largest p-value over the 1000 replications is also 0.000. Therefore, the tests show that the simulated data is indeed very stationary, which confirms that the model is stable.

is a dramatic reduction in the skewness and kurtosis.

The mean of the conditional variance of the first regime is found to be 0.000083. That of the second regime is 0.000568. This clearly shows that the second regime is much more volatile than the first regime. The ratio of the variances is about 6.8. The unconditional probability of being in the tranquil regime is 0.783, which indicates that during the period of interest (4861 days), roughly 1053 days were in the volatile regime.

6. Monte Carlo Integration of Option Prices

At this juncture, it has been shown that the SW-EGARCH specification for volatility is a better fit to security returns than a simple GARCH. In this section, values for options on volatility are computed using different specifications for the volatility process. The objective is to assess the economic significance of the impact of the volatility index properties on the value of the options. More specifically, we want to quantify the difference in the options value between the different volatility specifications analyzed in the previous section.

The focus is on the switching and asymmetry properties of volatility. However, the issue of risk premium has to be addressed. As stated before, volatility is not a traded asset. Therefore, a volatility option pricing model needs to account for that by making assumptions about the expected premium for volatility risk. G-L assume that it is proportional to the level of volatility. This paper focuses on the assumption for the underlying process of volatility. Therefore, we want to show that regardless of the choice of risk premium, the process in G-L leads to a severe mispricing in the options. To attain this objective, option values that are discounted by the risk-free rate alone are computed. This is done for simplicity. This is appropriate because we are not proposing a new pricing model, but rather highlighting the impact of assumptions about the underlying asset on the price of the options. All what is needed is to compare values computed using different underlying processes irrespective of the discount factor. In fact, if those

values are discounted using the same discount factor, the relative difference between the values based on different underlying processes will remain the same. In the next section, the results are replicated using a discount factor that incorporates a risk premium. It is shown that in this case, the conclusions remain the same.

The procedure to compute option prices is as follows:

- 1- The parameters from the volatility models estimated previously are used to generate a simulated volatility index¹¹.
- 2- The payoff at maturity of a European call option on the volatility process is computed.
- 3- The payoff is discounted using the appropriate discount rate.
- 4- The procedure is repeated for 100,000 independent replications.
- 5- The value of the option is then computed as the average of all option discounted payoffs at maturity over the 100,000 replications¹².

Call option values are computed for different time to maturity (1 month to 12 months), different exercise prices (5 to 30), and different volatility index - underlying asset - values (5 to 40). Figures 1, 3, and 5 present 3-Dimensional graphs depicting option values from the SW-EGARCH for different sets of characteristics. Figures 2, 4, and 6 provide 3-Dimensional graphs presenting the difference in option values between

¹¹We use a technique known as “synchronization” (see Law and Kelton, 1991). This implies two things. First, we adopt the same seed for the random number generator for each model. Second, we make sure to adopt the same sequence of generation of random variates for each model. This is important since we need to generate more random variates for the switching models to allow the random choice of different regimes. For that reason, for the uni-regime models, we generate the same sequence of random variates as for the switching models, but we simply discard those that are meant to determine the regime. The result from “synchronization” is that the random variates used by each model are identical. This ensures that the differences in option values will arise uniquely from the difference in assumed volatility processes (as opposed to the random variates generated).

¹²Standard errors are also computed for each option value. Given that a large number of replications (100,000) are made, the standard errors are extremely small compared to the difference between values computed from different processes. For that reason, statistical significance holds for all our results. As an indication, the largest standard error of any option value in our study is 0.0457. As will be seen, this is dwarfed by the difference between option values computed with different volatility processes. For that reason, the issue of statistical significance will be omitted throughout the exposition of the results.

those based on the SW-EGARCH and those based on the GARCH process. This will allow us to assess if the difference in option values between the two processes is large enough as compared to the option values under SW-EGARCH. This will tell us if the assumptions in G-L are acceptable approximations or not.

Figure 1 gives option values for different exercise prices and time to maturity.

***INSERT FIGURE 1 ABOUT HERE**

The volatility index value is fixed at 15. As expected, the option value is increasing in time to maturity and decreasing in the exercise price. The more interesting figure is figure 2. Figure 2 shows the difference in option values between the SW-EGARCH specification and the GARCH model. The option characteristics (exercise prices and time to maturity) are the same as in figure 1.

***INSERT FIGURE 2 ABOUT HERE**

A comparison of figures 1 and 2 shows that the differences in option values depicted by figure 2 are large and economically significant for a wide range of characteristics (exercise prices and time to maturity). The process in G-L (GARCH) always underpredicts the option value.

Figure 3 is similar to figure 1 in that it presents option values from the SW-EGARCH process for different characteristics. However, the difference is that the characteristics of interest are time to maturity and volatility index values (underlying asset). The exercise price is fixed at 15.

***INSERT FIGURE 3 ABOUT HERE**

The 3-Dimensional graph has an intriguing shape. This is due to the fact that the underlying asset (volatility index) has the property of mean-reversion. For low values of volatility index, the option value is increasing in time to maturity because the longer the time to maturity the higher the chance that the volatility index will “mean-revert” (increase) to its long run mean. On the other hand, for high values of volatility index, the option value is decreasing in time to maturity because the longer the time to maturity the higher the chance that the volatility index will “mean-revert” (decrease) to its long run mean. Also, as expected, the option value is increasing in the volatility index values for all maturities. Figure 4 shows the difference in option values between the SW-EGARCH specification and the GARCH model. The option characteristics (volatility index values and time to maturity) are the same as in figure 3.

***INSERT FIGURE 4 ABOUT HERE**

A comparison of figures 3 and 4 reveals that the differences in option values depicted by figure 4 are large and economically significant for a wide range of characteristics (volatility index values and time to maturity). The process in G-L (GARCH) almost always underpredicts the option value. Also, the difference is larger for longer time to maturity. This is an intuitive result since there is a higher likelihood of regime switching in the volatility index the longer the time to maturity. This regime switching behavior differentiates the SW-EGARCH from the GARCH process.

Figure 5 is similar to figures 1 and 3 in that it shows option values from the SW-EGARCH process for different characteristics. However, the difference is that the characteristics of interest are exercise price and volatility index values (underlying asset). Time to maturity is fixed at 3 months.

***INSERT FIGURE 5 ABOUT HERE**

Not surprisingly, the option value is increasing in exercise price and volatility index values. Figure 6 shows the difference in option values between the SW-EGARCH model and the GARCH specification. The option characteristics (volatility index values and exercise price) are the same as in figure 5.

***INSERT FIGURE 6 ABOUT HERE**

Again, a comparison between figures 5 and 6 demonstrates that the differences in option values depicted by figure 6 are large and economically significant for a wide range of characteristics (volatility index values and exercise prices).

The evidence from figures 1 to 6 shows that the difference in option values between the process used in G-L (GARCH) and a model that allows for volatility switching and asymmetry (SW-EGARCH) is economically large and significant for a wide range of characteristics (volatility index values, exercise price, and time to maturity). To confirm this evidence in a more formal way, the difference in the option value between the two processes is computed as a percent of the option values under the SW-EGARCH. This will show the percentage point the model in G-L misprice the option as compared to the SW-EGARCH. The percentage difference described above is computed for different

maturities (1 to 12 months) by taking the average over all exercise prices (5 to 30) and all volatility index values (5 to 40) within that maturity. The results from this computation are graphed in figure 7. Another curve that uses the median (instead of average) over all exercise prices and volatility index values is also graphed in figure 7 for robustness. The curves in figure 7 plot the percentage difference (mispricing or undervaluation in this case) for different time to maturity.

***INSERT FIGURE 7 ABOUT HERE**

The percentage differences in figure 7 are all large ranging from 7.0% for a 1-month option to 21.1% for a 12-month option (average). The numbers computed using the median are even larger confirming our results. Again, the mispricing induced by the process in G-L (GARCH) is shown to be most pronounced for long lived options. A 3-month option is mispriced (undervalued) by about 10% , which is large enough for derivative traders to care about. This clearly shows that the G-L model is too stylized to be used in pricing derivatives on volatility indices.

7. Further Analysis and Robustness

It is of interest to further analyze why the GARCH process is inappropriate. Which restrictions that are implicit in the GARCH process are those that make it inappropriate? In other words, it is interesting to see the independent contribution of the asymmetry and switching regime effects. Thus far, it is not clear which (or lack of which) contributes most to the economically large and significant mispricing found previously. To do that, the difference in option prices between the SW-EGARCH models and other more restrictive ones is computed. The restrictive models used for comparison are the SW-GARCH, EGARCH, GARCH and Switching Variance. The results are graphed in

figure 8. The figure plots the difference in percent between the option prices computed with SW-EGARCH and those with the competing model in question. As in figure 7, the percentage difference described above is computed for different maturities (1 to 12 months) by taking the average over all exercise prices (5 to 30) and all volatility index values (5 to 40) within that maturity.

***INSERT FIGURE 8 ABOUT HERE**

As can be seen from figure 8, the difference between prices from the Switching Variance and the SW-EGARCH is very large (the underpricing is between 50% and 60%). This is not surprising given that the Switching Variance model is very restrictive (disallows GARCH effects and asymmetry). The difference between the EGARCH and the SW-EGARCH is also large (the underpricing is between 20% and 30%). This indicates that the switching regime feature is very important. On the other hand, the difference between the SW-GARCH and the SW-EGARCH is rather small (between 2% and 6%). This shows that the volatility asymmetry feature is much less important than the regime switching in terms of option pricing. Finally, the underpricing from the GARCH process (as compared to the SW-EGARCH) is less than that of the EGARCH. This indicates that accounting for asymmetry without regime switching can make things worse (especially for short maturities).

Next, the consequences on our previous results of relaxing the assumption that volatility risk is not priced, is examined. G-L follow Wiggins (1987), Stein and Stein (1991), and others in assuming that the expected premium for volatility risk is proportional to the level of volatility, ζV where ζ is a constant parameter. This assumption is similar to the implications of general equilibrium models such as Cox, Ingersoll and Ross (1985), Hemler and Longstaff (1991), and Longstaff and Schwartz (1992), in which

risk premia in security returns are proportional to the level of volatility. We replicate the results in figure 7 for the case where volatility risk is priced. The only difference in option prices computed in this section from those presented before is in the discount factor. Previously, the discount factor was assumed to be the risk-free rate. In this section however, the discount factor is defined to be the risk-free rate plus a risk premium. The risk premium is a linear function of the level of volatility (ζV). ζ is chosen to be equal to 0.25. The choice of this coefficient is somewhat arbitrary. However, it seemed to deliver realistic risk premia for different levels of volatility. Moreover, we checked the robustness of our results by trying a wide range of values for ζ . The results were extremely close. Figure 9 replicates the results from figure 7 with the addition of the risk premium described above.

***INSERT FIGURE 9 ABOUT HERE**

A comparison between the two figures reveals that they are very similar. Essentially, the mispricing between the GARCH and the SW-EGARCH processes are as described before. There is a very small decrease in the mispricing when the risk premium is taken into account (less than 1% in all cases). Therefore, assuming a risk premium as opposed to just using the risk-free rate does not alter our results.

Finally, we analyze Volatility Index data (VIX). We want to make sure that the SW-EGARCH is also a better fit than GARCH in the case of VIX. Since the options are written on the Volatility Index itself, we need to prove that VIX is more consistent with SW-EGARCH than it is with GARCH. We obtain VIX daily data from the internet site of the Chicago Board Options Exchange. Our data is from the period January 2, 1986 to December 6, 2002. We estimate two models. The first one is the analog to a GARCH process:

$$VIX_t = \alpha + \delta_1 VIX_{t-1} + \varepsilon_t. \quad (7.1)$$

The second model we estimate is the analog to a SW-EGARCH:

$$\ln(VIX_t) = \alpha + \delta_1 VIX_{t-1} + \gamma_1 \left| \frac{r\varepsilon_{t-1}}{\sqrt{VIX_{t-1}}} \right| + \gamma_2 \frac{r\varepsilon_{t-1}}{\sqrt{VIX_{t-1}}} + \varepsilon_t. \quad (7.2)$$

where VIX_t is the Volatility Index, and $r\varepsilon_{t-1}$ is the residual from $R_t = \mu + r\varepsilon_{t-1}$. This residual term introduces asymmetry in the model analogously to SW-EGARCH, by allowing the sign (and magnitude) of the return on the S&P 100 to affect the Volatility Index. Also, note that the second model is a regime switching model.

The models are estimated using a normal likelihood for the error terms. The δ_1 coefficient in the model analog to GARCH is estimated to be 0.95. As expected, this indicates that the Volatility Index is highly persistent. The coefficient estimates for the second model are qualitatively similar to the SW-EGARCH model, and are available upon request from the authors. The key here is to figure out which model fits the data best. Since the two models are not nested, we use the penalized Akaike Information Criterion and the Bayesian Information Criterion. The AIC is 20079 for the GARCH analog and 13395 for the SW-EGARCH analog, respectively. As can be seen, the AIC for the SW-EGARCH analog is much smaller than that of the GARCH analog. This indicates that the SW-EGARCH representation is a much better fit to the VIX data than the GARCH representation. The BIC is 20098 for the GARCH analog and 13471 for the SW-EGARCH analog, respectively. This confirms the AIC results. To conclude, we are now comfortable that our results related to the conditional volatility carry through to the observed Volatility Index.

8. Conclusion

In this paper, it is shown that volatility index option pricing models that do not take into account the regime switching and asymmetry properties of volatility, undervalue a 3-month option by about 10%. Such a model was proposed in Grunbichler and Longstaff (1996). Their model is based on modeling volatility as a GARCH process. For that reason, their model does not take into account the asymmetric relation between volatility and returns. More importantly, their model does not account for the fact that volatility might be subject to different regimes. Switching Regime Asymmetric GARCH is used to model the generating process of security returns. The model specifies the conditional variance of security returns as following two regimes possessing different characteristics (tranquil regime and volatile regime). The comparison between the switching regime models and the traditional uni-regime models among GARCH, EGARCH, and GJR-GARCH demonstrates that the SW-EGARCH model fits the data best. Next, the values of European call options written on a volatility index are computed using Monte Carlo integration. Option values based on the SW-EGARCH model are compared with those based on the traditional GARCH specification. It is found that the option values obtained from the different processes are very different. This clearly shows that the G-L model is too stylized to be used in pricing derivatives on volatility indices. Therefore, uni-regime models for volatility should be used with caution because they are based on the assumption of parametric homogeneity and do not allow regime switching. There is an urgent need for research to develop volatility derivatives pricing models that account for volatility asymmetry and volatility regime switching. This will need to be done before derivatives on volatility indices are introduced by the major futures and options exchanges of the world. Finally, we propose three promising avenues for future research. An alternative approach to address the persistence in volatility (or spurious persistence) issue is that of Bollerslev and Mikkelsen (1996) and Baillie, Bollerslev, and Mikkelsen

(1996). They propose and use a Fractionally Integrated GARCH, which implies a slow hyperbolic rate of decay for the influence of lagged squared innovations. It might be worthwhile to explore volatility option pricing using the FIGARCH model. This study has focused on call options. It might be a worthwhile extension to check that the results still hold with put options.

BIBLIOGRAPHY

- Akgiray, V. (1989). Conditional heteroskedasticity in time series of stock returns: evidence and forecasts. *Journal of Business*, 62, 55-80.
- Ang, A., & Bekaert, G. (1998). Regime switches in interest rates. Working paper, Columbia University.
- Baillie, R.T., Bollerslev, T., & Mikkelsen, H.O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74, 3-30.
- Bera, A. K., & Higgins, M. L. (1993). ARCH models: Properties, estimation and testing. *Journal of Economic Surveys*, 7, 305-366.
- Billio, M., & Pelizzon, L. (1996). Pricing option with switching regime volatility. Working paper, University of Venice.
- Böhning, D., Dietz, E., Schaub, R., Schlattmann, P., & Lindsay, B.G. (1994). The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family. *Annals of the Institute of Statistical Mathematics*, 46, 373-388.
- Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the Business and Economic Statistics*, American Statistical, 177-181.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637-654.
- Blair, B., Poon, S., & Taylor, S. (1998). Asymmetric and crash effects in stock volatility for the S&P 100 index and its constituents. Working paper, Lancaster University.
- Bollen, N. (1998). Valuing options in regime-switching models. *Journal of Derivatives*, 6, 38-49.
- Bollen, N., Gray, S., & Whaley, R. (2000). Regime-switching in foreign exchange rates: Evidence from currency options. *Journal of Econometrics*, 94, 239-276.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal*

- of *Econometrics*, 31, 307-327.
- Bollerslev, T., Chou, R.Y., & Kroner, K.F. (1992). ARCH modeling in finance. *Journal of Econometrics*, 52, 5-59.
- Bollerslev, T., & Engle, R.F. (1986). Modelling the persistence of conditional variances. *Econometric Reviews*, 5, 1-50.
- Bollerslev, T., & Mikkelsen, H.O. (1996). Modeling and pricing long memory in stock market volatility. *Journal of Econometrics*, 73, 151-184.
- Brenner, M., & Galai, D. (1989). New financial instruments for hedging changes in volatility. *Financial Analysts Journal*, July/August, 61-65.
- Cai, J. (1994). A Markov model of switching-regime ARCH. *Journal of Business and Economic Statistics*, 12, 309-316.
- Campbell, S., & Li C. (1999). Option pricing with unobserved and regime-switching volatility. Working paper, University of Pennsylvania.
- Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica*, 53, 385-408.
- Dahlquist, M., & Gray, S. (2000). Regime-switching and interest rates in the European Monetary System. *Journal of International Economics*, forthcoming.
- Duan, J. (1995). The GARCH option pricing model. *Mathematical Finance*, 5, 13-32.
- Duan, J., Popova, I., & Ritchken, P. (1999). Option pricing under regime switching. Working paper, Hong Kong University of Science and Technology.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 50, 987-1008.
- Fleming, J., Ostdiek, B., & Whaley, R. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15, 265-302.
- Fornari, F., & Mele, A. (1997). Sign- and volatility-switching ARCH models: Theory and application to international stock markets. *Journal of Applied Econometrics*, 12, 49-67.

- French, K., Schwert, W., & Stambaugh, R. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19, 3-29.
- Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 68, 1779-1801.
- Gray, S. (1996a). Modeling the conditional distribution of interest rates as a regime switching process. *Journal of Financial Economics*, 42, 27-62.
- Gray, S. (1996b). Regime-switching in Australian interest rates. *Accounting and Finance*, 36, 65-88.
- Grunbichler, A., & Longstaff, F. (1996). Valuing futures and options on volatility. *Journal of Banking and Finance*, 20, 985-1001.
- Hamilton, J. (1994). Time series analysis. Princeton University Press, Princeton, N.J.
- Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57, 357-384.
- Hamilton, J. (1988). Rational-expectations econometric analysis of changes in regime: An investigation of the term structure of interest rates. *Journal of Economic Dynamics and Control*, 12, 385-423.
- Hamilton, J., & Susmel, R. (1994). Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64, 307-333.
- Hemler, M., & Longstaff, F. (1991). General equilibrium stock index futures prices: Theory and empirical evidence. *Journal of Financial and Quantitative Analysis*, 26, 287-308.
- Heston, S. (1993). A closed form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6, 327-343.
- Heston, S., & Nandi, S. (1999). A closed-form GARCH option pricing model. *Review of Financial Studies*, forthcoming.

- Hull, J., & White, A. (1987). Pricing options with stochastic volatilities. *Journal of Finance*, 42, 281-300.
- Kon, S. J. (1984). Models of stocks returns - A comparison. *Journal of Finance*, 39, 147-165.
- Kim, D. & Kon, S. J. (1999). Structural change and time dependence in models of stock returns. *Journal of Empirical Finance*, 6, 283-308.
- Law, A. M., & Kelton, W. D. (1991). Simulation modeling & analysis (2nd ed.). McGraw-Hill.
- Longstaff, F., & Schwartz, E. (1992). Interest rate volatility and the term structure: A two-factor general equilibrium model. *Journal of Finance*, 47, 1259-1282.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36, 394-419.
- Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A New Approach. *Econometrica*, 59(2), 347-370.
- Nelson, D. (1990). ARCH models as diffusion approximations. *Journal of Econometrics*, 45, 7-38.
- Norden, S. (1995). Regime switching as a test for exchange rate bubbles. Working paper, Bank of Canada.
- Norden, S., & Schaller, H. (1993). Regime switching in stock market returns. Manuscript, Bank of Canada and Carleton University.
- Perron, P. (1989). The great crash, the oil price shock, and the unit root hypothesis. *Econometrica*, 57, 1361-1401.
- Ritchey, R. (1990). Call option valuation for discrete normal mixtures. *Journal of Financial Research*, 13, 285-296.
- Rockinger M. (1994). Regime switching: Evidence for the French stock market. Mimeo, HEC.
- Schobel, R., & Zhu, J. (1998). Stochastic volatility with an Ornstein-Uhlenbeck Process:

- An extension. Working paper, University of Tuebingen.
- Stein, E., & Stein, C. (1991). Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies*, 4, 727-752.
- Susmel, R. (1998a). Switching volatility in Latin American emerging markets. *Emerging Markets Quarterly*, 2.
- Susmel, R. (1998b). A switching volatility model for inflation: International evidence. Working paper, University of Houston.
- Susmel, R. (1994). Switching volatility in international equity markets. Working paper, University of Houston.
- Taylor, S., & Xu, X. (1994). Implied volatility shapes when price and volatility shocks are correlated. Working paper, Lancaster University.
- Vosti, C. (1993). CBOE creates new index to deal with market volatility. *Pensions & Investments*, 8 February, 1993.
- Walsh, D. (1999). Some exotic options under symmetric and asymmetric conditional volatility of returns. *Journal of Multinational Financial Management*, 9, 403-417.
- Weiss, A. (1986). Asymptotic theory for ARCH models: Estimation and testing. *Econometric Theory*, 2, 107-131.
- Whaley, R. (2000). The investor fear gauge. Working paper, Duke University.
- Whaley, R. (1993). Derivatives on market volatility: Hedging tools long overdue. *Journal of Derivatives*, Fall, 71-84.
- Wiggins, J. (1987). Option values under stochastic volatility: Theory and empirical estimates. *Journal of Financial Economics*, 19, 351-372.

Table I. Summary Statistics

Variable	SP100
number of observations	4861
number of observations=0	52
number of observations>0	2530
mean	0.00051
t-statistic: mean=0	3.327**
standard deviation	0.01066
skewness	-2.397**
excess kurtosis	55.265**
Q(20)	11.208
Q ² (20)	350.29**
rho(1)	-0.013
rho ² (1)	0.159**

"Q(20)" is the Ljung-Box statistic of residuals, which follows a $\chi^2(20)$ distribution.

"Q²(20)" is the Ljung-Box statistic of squared residuals.

"rho(1)" is the first order autocorrelation of residuals.

"rho²(1)" is the first order autocorrelation of squared residuals.

** statistically significant at the 5% level.

Table II. Model Selection Tests

Panel A: Likelihood Ratio Tests		
Null	Alternative	SP100
GARCH (4)	GJR-GARCH (5)	51.84**
GARCH (4)	SW-GARCH (9)	373.84**
GARCH (4)	SW-GJR-GARCH (11)	288.12**
GJR-GARCH (5)	SW-GJR-GARCH (11)	236.28**
EGARCH (5)	SW-EGARCH (11)	345.94**
SWITCH VAR (5)	SW-GARCH (9)	376.90**
SWITCH VAR (5)	SW-GJR-GARCH (11)	291.18**
SWITCH VAR (5)	SW-EGARCH (11)	422.34**

Panel B: Maximum Penalized Likelihood Tests (AIC & BIC)		
	SP100	
	AIC	BIC
GARCH	-31603.90	-31577.94
GJR-GARCH	-31653.74	-31621.29
EGARCH	-31675.23	-31642.79
SW-GARCH	-31967.74	-31909.34
SW-GJR-GARCH	-31878.02	-31806.64
SW-EGARCH	-32009.20 ^A	-31937.80 ^B
SWITCH VAR	-31598.84	-31566.40

The number in parenthesis is the number of parameters in the corresponding model.

The Likelihood Ratio tests follow a $\chi^2(K_A - K_N)$ distribution, where K_N and K_A are the number of parameters corresponding to the Null and Alternative respectively.

$AIC = -2\ln(L) + 2K$; $BIC = -2\ln(L) + K\ln(N)$; Where $\ln(L)$, K , and N are the maximized log-likelihood, number of parameters, and number of observations, respectively.

** statistically significant at the 5% level.

^A model preferred by the AIC.

^B model preferred by the BIC.

Table III. Parameter Estimates of SW-EGARCH

Variable	SP100	
	Regime 1	Regime 2
α	-0.1267** (0.0203)	-0.3376 (0.7453)
δ_1	0.9928** (0.0018)	0.9676** (0.0785)
γ_1	0.0695** (0.0126)	0.2670 (0.3753)
γ_2	-0.0197** (0.0100)	-0.3113 (0.2070)
μ	0.0005** (0.0001)	
π_{11}	0.9675** (0.0393)	
π_{22}	0.2675 (0.2513)	
log -likelihood	16015.585	
Q(20)	20.916	
Q ² (20)	7.953	
rho(1)	0.0138	
rho ² (1)	-0.0133	
skewness	-0.4158**	
excess kurtosis	3.6155**	

The EGARCH(1,1) specification is defined as follows

$$\ln(h_t) = \alpha + \delta_1 h_{t-1} + \gamma_1 |z_{t-1}| + \gamma_2 z_{t-1}.$$

where h_t is the conditional variance, z_t i.i.d. $\sim N(0, 1)$, and $\alpha, \delta_1, \gamma_1$, and γ_2 are constant parameters.

The numbers in parenthesis are standard errors.

"Q(20)" is the Ljung-Box statistic of residuals, which follows a $\chi^2(20)$ distribution.

"Q²(20)" is the Ljung-Box statistic of squared residuals.

"rho(1)" is the first order autocorrelation of residuals.

"rho²(1)" is the first order autocorrelation of squared residuals.

** statistically significant at the 5% level.

Figure 1. SW-EGARCH Option Values for Different Exercise Prices and Time to Maturity

This figure is a 3-Dimensional graph depicting call option values from the SW-EGARCH for different time to maturity (1 month to 12 months) and different exercise prices (5 to 30). The volatility index value is fixed at 15.

Figure 2. Difference between SW-EGARCH and GARCH Based Option Values for Different Exercise Prices and Time to Maturity

This figure is a 3-Dimensional graph presenting the difference in call option values between those based on the SW-EGARCH and those based on the GARCH processes for different time to maturity (1 month to 12 months) and different exercise prices (5 to 30). The volatility index value is fixed at 15.

Figure 3. SW-EGARCH Option Values for Different Volatility Index Values and Time to Maturity

This figure is a 3-Dimensional graph depicting call option values from the SW-EGARCH for different time to maturity (1 month to 12 months) and different volatility index values (5 to 40). The exercise price is fixed at 15.

Figure 4. Difference between SW-EGARCH and GARCH Based Option Values for Different Volatility Index Values and Time to Maturity

This figure is a 3-Dimensional graph presenting the difference in call option values between those based on the SW-EGARCH and those based on the GARCH processes for different time to maturity (1 month to 12 months) and different volatility index values (5 to 40). The exercise price is fixed at 15.

Figure 5. SW-EGARCH Option Values for Different Volatility Index Values and Exercise Prices

This figure is a 3-Dimensional graph depicting call option values from the SW-EGARCH for different exercise prices (5 to 30) and different volatility index values (5 to 40). Time to maturity is fixed at 3 months.

Figure 6. Difference between SW-EGARCH and GARCH Based Option Values for Different Volatility Index Values and Exercise Prices

This figure is a 3-Dimensional graph presenting the difference in call option values between those based on the SW-EGARCH and those based on the GARCH processes for different exercise prices (5 to 30) and different volatility index values (5 to 40). Time to maturity is fixed at 3 months.

Figure 7. Percent Mispricing (Undervaluation) in Option Values from using a GARCH Process Instead of a SW-EGARCH

The curves in this figure plot the percentage difference (mispricing or undervaluation in this case) for different time to maturity. The difference in the call option value between the two processes is computed as a percent of the option values under the SW-EGARCH. The percentage difference described above is computed for different maturities (1 to 12 months) by taking the average/median over all exercise prices (5 to 30) and all volatility index values (5 to 40) within that maturity.

Figure 8. Percent Undervaluation in Option Values from Different Processes as Compared to SW-EGARCH

The curves in this figure plot the percentage difference (undervaluation) for different time to maturity. The difference in the call option value between the process in question and the SW-EGARCH is computed as a percent of the option values under the SW-EGARCH. The percentage difference described above is computed for different maturities (1 to 12 months) by taking the average over all exercise prices (5 to 30) and all volatility index values (5 to 40) within that maturity. “SV” stands for the Switching Variance model.

Figure 9. Percent Undervaluation in Option Values from using a GARCH Process Instead of a SW-EGARCH (Accounting for Volatility Risk Premium)

The curves in this figure plot the percentage difference (undervaluation) for different time to maturity. The difference in the call option value between the two processes is computed as a percent of the option values under the SW-EGARCH. The percentage difference described above is computed for different maturities (1 to 12 months) by taking the average/median over all exercise prices (5 to 30) and all volatility index values (5 to 40) within that maturity. The discount factor is defined to be the risk-free rate plus a risk premium factor. The risk premium factor is a linear function of the level of volatility (ζV). ζ is chosen to be equal to 0.25