![CREATES - Center for Research in Econometric Analysis of Time Series]

Option Valuation with Long-run and Short-run Volatility Components

Peter Christoffersen, Kris Jacobs, Chayawat Ornthanalai and Yintian Wang

School of Economics and Management
University of Aarhus
Building 1322, DK-8000 Aarhus C
Denmark

Aarhus School of Business
University of Aarhus
Handelshøjskolen
Aarhus Universitet

UNIVERSITY OF COPENHAGEN

Option Valuation with Long-run and Short-run Volatility Components*

Peter Christoffersen

McGill University and CREATES

Chayawat Ornthanalai

McGill University

Kris Jacobs

McGill University and Tilburg University

Yintian Wang

Tsinghua University

February 7, 2008

Abstract

This paper presents a new model for the valuation of European options, in which the volatility of returns consists of two components. One of these components is a long-run component, and it can be modeled as fully persistent. The other component is short-run and has a zero mean. Our model can be viewed as an affine version of Engle and Lee (1999), allowing for easy valuation of European options. The model substantially outperforms a benchmark single-component volatility model that is well-established in the literature, and it fits options better than a model that combines conditional heteroskedasticity and Poisson-normal jumps. The component model's superior performance is partly due to its improved ability to model the smirk and the path of spot volatility, but its most distinctive feature is its ability to model the volatility term structure. This feature enables the component model to jointly model long-maturity and short-maturity options.

JEL Classification: G12
Keywords: Volatility term structure; GARCH; out-of-sample.

1 Introduction

There is a consensus in the equity options literature that combining time-variation in the conditional variance of asset returns (Engle (1982), Bollerslev (1986)) with a leverage effect (Black (1976)) constitutes a potential solution to well-known biases associated with the Black-Scholes (1973) model, such as the implied volatility smirk. These asymmetric dynamic volatility models generate negative skewness in the distribution of asset returns which in turn generates higher prices for out-of-the-money put options as compared to the Black-Scholes formula. In the continuous-time option valuation literature, the Heston (1993) model addresses some of these biases. This model contains a leverage effect as well as stochastic volatility.[1] In the discrete-time literature, the NGARCH$(1, 1)$ option valuation model proposed by Duan (1995) contains time-variation in the conditional variance as well as a leverage effect. The model by Heston and Nandi (2000) is closely related to Duan's model.

Many existing empirical studies have confirmed the importance of time-varying volatility, the leverage effect and negative skewness in continuous-time and discrete-time setups, using parametric as well as non-parametric techniques.[2] However, it has become clear that while these models help explain the biases of the Black-Scholes model in a qualitative sense, they come up short in a quantitative sense. Using parameters estimated from returns or options data, these models reduce the biases of the Black-Scholes model, but the magnitude of the effects is insufficient to completely resolve the biases. The resulting pricing errors have the same sign as the Black-Scholes pricing errors, but are smaller in magnitude. We therefore need models that possess the same qualitative features as the models in Heston (1993) and Duan (1995), but that contain stronger quantitative effects. These models need to generate more flexible skewness and volatility of volatility dynamics in order to fit observed option prices. Existing studies have attempted to address this by combing stochastic volatility specifications with jump processes, or by using non-normal innovations in heteroskedastic models.[3]

The shortcomings of existing models in modeling the moneyness dimension are compounded by their shortcomings in modeling the term structure of volatility, as well as the path of spot volatility. It has been observed using a variety of diagnostics that it is difficult to fit the dynamics of return volatility using a benchmark model such as a GARCH$(1, 1)$. A similar observation applies to stochastic volatility models such as Heston (1993). The main problem is that volatility autocorrelations are too high at longer lags to be explained by a GARCH$(1, 1)$, unless the process is extremely persistent. This extreme persistence may impact negatively on other aspects of option valuation, such as the valuation of short-maturity options.

In fact, it has been observed in the literature that volatility may be better modeled using a fractionally integrated process, as in Baillie, Bollerslev and Mikkelsen (1996), rather than an exponential GARCH process. Andersen, Bollerslev, Diebold and Labys (2003) confirm this finding

---

[1] The importance of stochastic volatility is also studied in Hull and White (1987), Melino and Turnbull (1990), Scott (1987) and Wiggins (1987).

[2] See for example Ait-Sahalia and Lo (1998), Amin and Ng (1993), Bakshi, Cao and Chen (1997), Bates (2000), Benzoni (1998), Bollerslev and Mikkelsen (1999), Chernov and Ghysels (2000), Duan, Ritchken and Sun (2005, 2006), Engle and Mustafa (1992), Eraker (2004), Heston and Nandi (2000), Jones (2003), Nandi (1998) and Pan (2002).

[3] See for example Bakshi, Cao and Chen (1997), Bates (2000), Broadie, Chernov and Johannes (2007), Christoffersen, Heston and Jacobs (2006), Eraker, Johannes and Polson (2003), Eraker (2004), Huang and Wu (2004) and Pan (2002).

using realized volatility. Bollerslev and Mikkelsen (1996, 1999) and Comte, Coutin and Renault (2001) investigate and discuss some of the implications of long memory for option valuation. Using fractional integration models for option valuation is somewhat cumbersome. Optimization is time-intensive and certain ad-hoc choices have to be made regarding implementation.

This paper attempts to remedy remaining option biases by modeling richer volatility dynamics. We use a model that is relatively easy to implement and that captures the stylized facts addressed by long-memory models at horizons relevant for option valuation. The model builds on Heston and Nandi (2000) and Engle and Lee (1999). The volatility of returns consists of two components. One of these components is a long-run component, and it can be modeled as (fully) persistent. The other component is short-run and mean zero. We study two models: one where the long-run component is constrained to be fully persistent and another where it is not. We refer to these models as the persistent component model and the component model respectively. These models are able to generate autocorrelations that are richer than those of a GARCH(1, 1) model while using just a few additional parameters.

Unobserved component or factor models are very popular in the finance literature. See Fama and French (1988), Poterba and Summers (1988) and Summers (1986) for applications to stock prices. In the option pricing literature, Bates (2000) and Taylor and Xu (1994) investigate two-factor stochastic volatility models. Duffie, Pan and Singleton (2000) provide a general continuous-time framework for the valuation of contingent claims using multifactor affine models. Eraker (2004) suggests the usefulness of a multifactor approach based on his empirical results. Alizadeh, Brandt and Diebold (2002) uncover two factors in stochastic volatility models of exchange rates using range-based estimation. Bollerslev and Zhou (2002), Brandt and Jones (2006), Chacko and Viceira (2003), Chernov, Gallant, Ghysels and Tauchen (2003) and Maheu (2005) also find that two-factor stochastic volatility models outperform single factor models when modeling daily asset return volatility. Adrian and Rosenberg (2005) investigate the relevance of a two-component volatility model for pricing the cross-section of stock returns. Unobserved component models are also very popular in the term structure literature, although in this literature the models are more commonly referred to as multifactor models.[4] There are very interesting parallels between our approach and results and stylized facts in the term structure literature. In the term structure literature, it is customary to model short-run fluctuations around a time-varying long-run mean of the short rate. In our framework we model short-run fluctuations around a time-varying long-run volatility.

Dynamic factor and component models can be implemented in continuous or discrete time.[5] We choose a discrete-time approach because of the ease of implementation. In particular, our model is related to the GARCH class of processes, and therefore volatility filtering and forecasting are relatively straightforward, which is critically important for option valuation.[6] An additional advantage of our model is parsimony: the most general model we investigate has seven parame-

---

[4]See for example Dai and Singleton (2000), Duffee (1999), Duffie and Singleton (1999) and Pearson and Sun (1994).

[5]Duffie, Pan and Singleton (2000) suggest a multifactor continuous-time model that captures the spirit of our approach, but do not investigate the model empirically.

[6]Because the filtering problem is extremely simple in the GARCH framework, we are able to analyze an extensive option sample. See also Heston and Nandi (2000). See among others Bates (2000, 2006), Chernov and Ghysels (2000), Eraker (2004) and Pan (2002) for other empirical studies that estimate model parameters using options data.

ters. We speculate that parsimony may help our model's out-of-sample performance.

The component model is a generalization of the GARCH$(1,1)$ model, and because its implementation uses similar techniques, the GARCH$(1,1)$ is a natural benchmark. Moreover, Christoffersen, Jacobs and Mimouni (2005) find that the performance of the GARCH$(1,1)$ model is similar to that of the Heston (1993) model, which is the most commonly used benchmark in the literature. Heston and Nandi (2000) find that the GARCH$(1,1)$ slightly outperforms the ad-hoc implied volatility benchmark model in Dumas, Fleming and Whaley (1998). As there is evidence that Poisson-normal jump processes can alleviate some of the biases associated with the Heston (1993) model and the GARCH$(1,1)$ model, we also include a GARCH$(1,1)$ model augmented with Poisson-normal jumps in our analysis.

We provide two different analyses of the component model. We first estimate the physical model parameters by maximum likelihood estimation (MLE) on historical S&P500 returns for 1962-2001. We compare the component model and the persistent component model with the GARCH$(1,1)$ benchmark as well as with the more general GARCH$(1,1)$-Jump model. Based on the log-likelihood criterion, the GARCH$(1,1)$-Jump model performs the best, followed by the component model, the persistent component model and the GARCH$(1,1)$ model. However, when we compare the models based on option RMSE fit using MLE parameters, the best fit is obtained using the component and persistent component models, followed by the GARCH$(1,1)$-Jump model. The GARCH$(1,1)$ model is again the worst performer. The RMSE in the component and persistent component models is about 23.7% below that of the GARCH(1,1).

We also use the MLE parameters to emphasize differences in important model features, such as the conditional volatility of variance, the correlation between returns and conditional variance, the term structure of conditional skewness and kurtosis, the volatility smirk and the volatility term structure. The improvement in the model's performance is primarily due to its richer dynamics, which result in different modeling of the term structure, and which enable the component model to capture patterns in long-maturity as well as short-maturity options.

In a second empirical investigation, we estimate the models using options data, while filtering the latent volatility from the underlying returns data. The performance of the component model is impressive when compared with a benchmark GARCH$(1,1)$ model. When using option data in estimation, the RMSE of the component model is between 10.3% and 22.7% lower than that of the benchmark GARCH model in-sample and between 20.8% and 23.5% lower out-of-sample. Our out-of-sample results strongly suggest that the in-sample results are not simply due to spurious overfitting. The persistent component model performs better than the benchmark GARCH$(1,1)$ model, but in contrast to the results obtained using MLE parameters, it is clearly inferior to the component model both in- and out-of-sample.

The paper proceeds as follows. Section 2 introduces the model. Section 3 discusses the volatility term structure and Section 4 discusses option valuation. Sections 5 and 6 present the empirical results, and Section 7 concludes.

2   Return Dynamics with Volatility Components

In this section we first present the Heston-Nandi GARCH$(1,1)$ model which will serve as the benchmark model throughout the paper. We then construct the component model as a natural extension of the GARCH$(1,1)$ model. We also present the persistent component model as a

special case of the component model.

2.1   The Heston and Nandi GARCH(1,1) Model

Heston and Nandi (2000) propose a class of GARCH models that allow for a closed-form solution for the price of a European call option. They present an empirical analysis of the GARCH$(1, 1)$ version of this model, which is given by

$$
\begin{aligned}
R_{t+1} & \equiv \ln(S_{t+1}/S_t) = r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1} && (2.1) \\
h_{t+1} & = w + b h_t + a(z_t - c\sqrt{h_t})^2
\end{aligned}
$$

where S_{t+1} denotes the underlying asset price, r the risk free rate, λ the price of risk and h_{t+1} the daily variance on day $t+1$ which is known at the end of day t. The z_{t+1} shock is assumed to be i.i.d. $N(0, 1)$. The Heston-Nandi model captures time variation in the conditional variance as in Engle (1982) and Bollerslev (1986),[7] and the parameter c captures the so-called leverage effect. The leverage effect captures the negative relationship between shocks to returns and volatility (Black (1976)), which results in a negatively skewed distribution of returns.[8] Note that the GARCH$(1, 1)$ dynamic in (2.1) is slightly different from the more conventional NGARCH model used by Engle and Ng (1993) and Hentschel (1995), which is used for option valuation in Duan (1995). The reason is that the dynamic in (2.1) is engineered to yield a closed-form solution for option valuation, whereas a closed-form solution does not obtain for the more conventional GARCH dynamic. Hsieh and Ritchken (2005) provide evidence that the more traditional GARCH model may actually slightly dominate the fit of (2.1). Our main point can be demonstrated using either dynamic. Because of the convenience of the closed-form solution provided by dynamics such as (2.1), we use this as a benchmark in our empirical analysis and we model the richer component structure within the Heston-Nandi framework.[9]

To better appreciate the workings of the component models presented below, note that by using the expression for the unconditional variance

$$
E\left[h_{t+1}\right] \equiv \sigma^2 = \frac{w + a}{1 - b - ac^2}
$$

to substitute out w, the variance process can be rewritten as

$$
h_{t+1} = \sigma^2 + b\left(h_t - \sigma^2\right) + a\left((z_t - c\sqrt{h_t})^2 - (1 + c^2\sigma^2)\right). \tag{2.2}
$$

---

[7]For an early application of GARCH to stock returns, see French, Schwert and Stambaugh (1987).

[8]Its importance for option valuation has been emphasized among others by Benzoni (1998), Chernov and Ghysels (2000), Christoffersen and Jacobs (2004), Eraker (2004), Eraker, Johannes and Polson (2003), Heston (1993), Heston and Nandi (2000) and Nandi (1998).

[9]See Bollerslev and Mikkelsen (1996), Engle and Mustafa (1992), Christoffersen and Jacobs (2004), and Hsieh and Ritchken (2005) for other empirical studies of European option valuation using GARCH dynamics. Ritchken and Trevor (1999) discusses the pricing of American options with GARCH processes.

2.2 Building a Component Volatility Model

The expression for the GARCH$(1,1)$ variance process in (2.2) highlights the role of the parameter σ^2 as the constant unconditional mean of the conditional variance process. A natural generalization is then to specify σ^2 as time-varying. Denoting this time-varying component by q_{t+1}, the expression for the variance in (2.2) can be generalized to

$$h_{t+1} = q_{t+1} + \beta (h_t - q_t) + \alpha \left((z_t - \gamma_1 \sqrt{h_t})^2 - (1 + \gamma_1^2 q_t) \right). \tag{2.3}$$

This model is similar in spirit to the component model of Engle and Lee (1999). The difference between our model and Engle and Lee (1999) is that the functional form of the GARCH dynamic (2.3) allows for a closed-form solution for European option prices. This is similar to the difference between the Heston-Nandi (2000) GARCH$(1,1)$ dynamic and the more traditional NGARCH$(1,1)$ dynamic discussed in the previous subsection. In specification (2.3), the conditional volatility h_{t+1} can most usefully be thought of as having two components. Following Engle and Lee (1999), we refer to the component q_{t+1} as the long-run component, and to $h_{t+1} - q_{t+1}$ as the short-run component. We will discuss this terminology in some more detail below. Note that by construction the unconditional mean of the short-run component $h_{t+1} - q_{t+1}$ is zero.

The model can also be written as

$$
\begin{aligned}
h_{t+1} &= q_{t+1} + (\alpha\gamma_1^2 + \beta)(h_t - q_t) + \alpha \left((z_t - \gamma_1 \sqrt{h_t})^2 - (1 + \gamma_1^2 h_t) \right) \\
&= q_{t+1} + \tilde{\beta}(h_t - q_t) + \alpha \left((z_t - \gamma_1 \sqrt{h_t})^2 - (1 + \gamma_1^2 h_t) \right)
\end{aligned}
\tag{2.4}
$$

where $\tilde{\beta} = \alpha\gamma_1^2 + \beta$. This representation is useful because we can think of

$$
\begin{aligned}
v_{1,t} &\equiv \left(z_t - \gamma_1 \sqrt{h_t} \right)^2 - (1 + \gamma_1^2 h_t) \\
&= (z_t^2 - 1) - 2\gamma_1 \sqrt{h_t} z_t
\end{aligned}
\tag{2.5}
$$

as a mean-zero innovation.

The model is completed by specifying the functional form of the long-run volatility component. In a first step, we assume that q_{t+1} follows the process

$$q_{t+1} = \omega + \rho q_t + \varphi \left((z_t^2 - 1) - 2\gamma_2 \sqrt{h_t} z_t \right). \tag{2.6}$$

Note that $E[q_{t+1}] = E[h_{t+1}] = \sigma^2 = \dfrac{\omega}{1-\rho}$ as long as $\rho < 1$. We can therefore write the component volatility model as

$$
\begin{aligned}
h_{t+1} &= q_{t+1} + \tilde{\beta}(h_t - q_t) + \alpha v_{1,t} \\
q_{t+1} &= \omega + \rho q_t + \varphi v_{2,t} \\
&= \sigma^2 + \rho(q_t - \sigma^2) + \varphi v_{2,t}
\end{aligned}
\tag{2.7}
$$

with

$$v_{i,t} = \left(z_t^2 - 1\right) - 2\gamma_i \sqrt{h_t} z_t, \text{ for } i = 1, 2. \tag{2.8}$$

and $E_{t-1}\left[v_{i,t}\right] = 0$, $i = 1, 2$. Also note that in addition to the price of risk, λ, the model contains seven parameters: α, $\tilde{\beta}$, γ_1, γ_2, ω, ρ and φ. In comparison, the GARCH(1,1) model above has four variance parameters.

2.3 A Fully Persistent Special Case

In our empirical work, we also investigate a special case of the model in (2.7). Notice that in (2.7) the long-run component of volatility will be a mean reverting process for $\rho < 1$. We also estimate a version of the model which imposes $\rho = 1$. The resulting process is

$$\begin{aligned}
h_{t+1} &= q_{t+1} + \tilde{\beta}\left(h_t - q_t\right) + \alpha v_{1,t} \\
q_{t+1} &= \omega + q_t + \varphi v_{2,t}
\end{aligned} \tag{2.9}$$

and $v_{i,t}$, $i = 1, 2$ are as in (2.8). In addition to the price of risk, λ, the model now contains six parameters: $\alpha, \tilde{\beta}, \gamma_1, \gamma_2, \omega$ and φ.

In this case the process for long-run volatility contains a unit root and shocks to the long-run volatility never die out: they have a "permanent" effect. Recall that following Engle and Lee (1999) in (2.7) we refer to q_{t+1} as the long-run component and to $h_{t+1} - q_{t+1}$ as the short-run component. In the special case (2.9) we can also refer to q_{t+1} as the "permanent" component, because innovations to q_{t+1} are truly "permanent" and do not die out. It is then customary to refer to $h_{t+1} - q_{t+1}$ as the "transitory" component, which reverts to zero. It is in fact this permanent-effects version of the model that is most closely related to models which have been studied more extensively in the finance and economics literature, rather than the more general model in (2.7).[10] We will refer to this model as the persistent component model.

It is clear that (2.9) is nested by (2.7). It is therefore to be expected that the in-sample fit of (2.7) is superior. However, out-of-sample this may not necessarily be the case. It is often the case that more parsimonious models perform better out-of-sample if the restriction imposed by the model is a sufficiently adequate representation of reality. The persistent component model may also be better able to capture structural breaks in volatility out-of-sample, because a unit root in the process allows it to adjust to a structural break, which is not possible for a mean-reverting process. It will therefore be of interest to verify how close ρ is to one when estimating the more general model (2.7).

3 Variance Term Structures

In order to intuitively understand the shortcomings of existing models such as the GARCH$(1, 1)$ model in (2.1) and the improvements provided by our model (2.7), it is instructive to graphically show some of the models' statistical properties that are key for option valuation. In this section we therefore illustrate the models' variance term structures and impulse response functions.

---

[10]See Fama and French (1988), Poterba and Summers (1988) and Summers (1986) for applications to stock prices. See Beveridge and Nelson (1981) for an application to macroeconomics.

3.1 The Variance Term Structure for the GARCH(1,1) Model

Following the logic used for the component model in (2.7), we can rewrite the GARCH$(1,1)$ variance dynamic in (2.2). We have

$$h_{t+1} = \sigma^2 + \tilde{b}\left(h_t - \sigma^2\right) + a\left((z_t^2 - 1) - 2c\sqrt{h_t}z_t\right) \tag{3.1}$$

where $\tilde{b} = b + ac^2$ and where the innovation term has a zero conditional mean. From (3.1) the multi-step forecast of the conditional variance is

$$E_t\left[h_{t+k}\right] = \sigma^2 + \tilde{b}^{k-1}(h_{t+1} - \sigma^2)$$

where the conditional expectation is taken at the end of day t. Notice that $\tilde{b}$ is directly interpretable as the variance persistence in this representation of the model.

We can now define a convenient measure of the variance term structure for maturity T as

$$h_{t+1:t+T} \equiv \frac{1}{T}\sum_{k=1}^{T} E_t\left[h_{t+k}\right] = \frac{1}{T}\sum_{k=1}^{T}\sigma^2 + \tilde{b}^{k-1}(h_{t+1} - \sigma^2) = \sigma^2 + \frac{1 - \tilde{b}^T}{1 - \tilde{b}}\frac{(h_{t+1} - \sigma^2)}{T}.$$

This variance term structure measure succinctly captures important information about the model's potential to explain the variation of option values across maturities.[11] To compare different models, it is convenient to set the current variance, h_{t+1}, to a simple m multiple of the long run variance. In this case the variance term structure relative to the unconditional variance is given by

$$h_{t+1:t+T}/\sigma^2 \equiv 1 + \frac{1 - \tilde{b}^T}{1 - \tilde{b}}\frac{(m - 1)}{T}.$$

The bottom-left panel in Figure 1 show the term structure of variance for the GARCH$(1,1)$ model for low (dashed line) and high (solid line) initial conditional variances. We use parameter values estimated via MLE on daily S&P500 returns (the estimation details are in Table 1 and will be discussed further below). We set $m = \frac{1}{2}$ to investigate low initial variance and $m = 2$ to investigate high initial variance. The figures present the variance term structure for up to 250 days, which corresponds approximately to the number of trading days in a year and therefore captures the empirically relevant term structure for option valuation. It can be clearly seen that for the GARCH$(1,1)$ model, the conditional variance converges to the long-run variance rather fast.

We can also learn about the dynamics of the variance term structure through impulse response functions. For the GARCH$(1,1)$ model, the effect of a shock at time t, z_t, on the expected k-day ahead variance is

$$\partial(E_t\left[h_{t+k}\right])/\partial z_t^2 = \tilde{b}^{k-1}a\left(1 - c\sqrt{h_t}/z_t\right)$$

---

[11]Notice that due to the price of risk term in the conditional mean of returns, the term structure of variance as defined here is not exactly equal to the conditional variance of cumulative returns over T days.

and thus the effect on the variance term structure is[12]

$$\partial E_t\left[h_{t+1:t+T}\right]/\partial z_t^2 = \frac{1-\tilde{b}^T}{1-\tilde{b}}\frac{a}{T}\left(1-c\sqrt{h_t}/z_t\right).$$

The bottom-left panel of Figure 2 plots the impulse responses to the term structure of variance for $h_t = \sigma^2$ and $z_t = 2$ (dashed line) and $z_t = -2$ (solid line) respectively, again using the parameter estimates from Table 1. The impulse responses are normalized by the unconditional variance. Notice that the effect of a shock dies out rather quickly for the GARCH$(1,1)$ model. Figure 2 also illustrates the asymmetric response of the variance term structure from a positive versus negative shock to returns. Due to the presence of a positive c, a positive shock has less impact than a negative shock along the entire term structure of variance.

3.2   The Variance Term Structure for the Component Model

In the component model we have

$$
\begin{aligned}
h_{t+1} &= q_{t+1} + \tilde{\beta}\left(h_t - q_t\right) + \alpha v_{1,t}\\
q_{t+1} &= \sigma^2 + \rho(q_t - \sigma^2) + \varphi v_{2,t}.
\end{aligned}
$$

The multi-day forecast of the two components are

$$
\begin{aligned}
E_t\left[h_{t+k} - q_{t+k}\right] &= \tilde{\beta}^{k-1}(h_{t+1} - q_{t+1})\\
E_t\left[q_{t+k}\right] &= \sigma^2 + \rho^{k-1}\left(q_{t+1} - \sigma^2\right).
\end{aligned}
$$

The simplicity of these multi-day forecasts is a key advantage of the component model. The multi-day variance forecast is a simple sum of the forecasts for the two variance components. Notice that $\tilde{\beta}$ and ρ correspond directly to the persistence of the short-run and long-run components respectively.

We can now calculate the variance term structure in the component model for maturity T as

$$
\begin{aligned}
h_{t+1:t+T} &\equiv \frac{1}{T}\sum_{k=1}^{T}E_t\left[q_{t+k}\right] + E_t\left[h_{t+k} - q_{t+k}\right]\\
&= \frac{1}{T}\sum_{k=1}^{T}\sigma^2 + \rho^{k-1}\left(q_{t+1} - \sigma^2\right) + \tilde{\beta}^{k-1}(h_{t+1} - q_{t+1})\\
&= \sigma^2 + \frac{1-\rho^T}{1-\rho}\frac{q_{t+1} - \sigma^2}{T} + \frac{1-\tilde{\beta}^T}{1-\tilde{\beta}}\frac{h_{t+1} - q_{t+1}}{T}.
\end{aligned}
$$

If we set q_{t+1} and h_{t+1} equal to m_1 and m_2 multiples of the long run variance respectively, then

---

[12]Engle and Patton (2001) refer to the effect of a shock on the expected k-day ahead variance and on the variance term structure as the forward persistence and the cumulative forward persistence respectively.

we get the variance term structure relative to the unconditional variance simply as

$$h_{t+1:t+T}/\sigma^2 = 1 + \frac{1-\rho^T}{1-\rho}\frac{m_1-1}{T} + \frac{1-\tilde{\beta}^T}{1-\tilde{\beta}}\frac{m_2-m_1}{T}. \tag{3.2}$$

The top-left panel in Figure 1 shows the term structure of variance for the component model using parameters estimated via MLE on daily S&P500 returns from Table 1. We set $m_1 = \frac{3}{4}$, $m_2 = \frac{1}{2}$ to investigate low initial variance (dashed line) and $m_1 = \frac{7}{4}$, $m_2 = 2$ to investigate high initial variance (solid line). By picking m_2 equal to the m used for the GARCH$(1,1)$ model, we ensure comparability across models within each figure because the spot variances relative to their long-run variances are identical.[13] The main conclusion from Figure 1 is that compared to the GARCH$(1,1)$, the conditional variance converges more slowly to the unconditional variance in the component model. This is particularly so on days with a high spot variance. The right-side panels show the contribution to the total variance from each component. Notice the strong persistence in the long-run component.

We can also calculate impulse response functions in the component model. The effects of a shock at time t, z_t on the expected k-day ahead variance components are

$$\partial E_t\left[q_{t+k}\right]/\partial z_t^2 = \rho^{k-1}\varphi\left(1-\gamma_2\sqrt{h_t}/z_t\right)$$

$$\partial E_t\left[h_{t+k}-q_{t+k}\right]/\partial z_t^2 = \tilde{\beta}^{k-1}\alpha\left(1-\gamma_1\sqrt{h_t}/z_t\right)$$

$$\partial E_t\left[h_{t+k}\right]/\partial z_t^2 = \tilde{\beta}^{k-1}\alpha\left(1-\gamma_1\sqrt{h_t}/z_t\right)+\rho^{k-1}\varphi\left(1-\gamma_2\sqrt{h_t}/z_t\right).$$

Notice again the simplicity due to the component structure. The impulse response on the term structure of variance is then

$$\partial E_t\left[h_{t+1:t+T}\right]/\partial z_t^2 = \frac{1-\tilde{\beta}^T}{1-\tilde{\beta}}\frac{\alpha}{T}\left(1-\gamma_1\sqrt{h_t}/z_t\right)+\frac{1-\rho^T}{1-\rho}\frac{\varphi}{T}\left(1-\gamma_2\sqrt{h_t}/z_t\right).$$

The top-left panel of Figure 2 plots the impulse responses to the term structure of variance for $h_t = \sigma^2$ and $z_t = 2$ (dashed line) and $z_t = -2$ (solid line) respectively. This figure reinforces the message from Figure 1 that when using parameters estimated from the data, the component model is quite different from the GARCH$(1,1)$ model. The effects of shocks are much longer-lasting in the component model using estimated parameter values because of the parameterization of the long-run component. Figure 2 also indicates that the term structure of the volatility asymmetry is more flexible in the component model. As a result current shocks and the current state of the economy potentially have a much more profound impact on the pricing of options across maturities in the component model.

It has been argued in the literature that the hyperbolic rate of decay displayed by long memory processes may be a more adequate representation for the conditional variance of returns.[14] We do

---

[13]Note that we need $m_1 \neq m_2$ in this numerical experiment to generate a "short-term" effect in (3.2). Changing m_1 will change the picture but the main conclusions stay the same.

[14]See Bollerslev and Mikkelsen (1996,1999), Baillie, Bollerslev and Mikkelsen (1996) and Ding, Granger and Engle (1993).

not disagree with these findings. Instead, we argue that Figures 1 and 2 demonstrate that in the component model the combination of two variance components with exponential decay gives rise to a slower decay pattern that sufficiently adequately captures the hyperbolic decay pattern of long memory processes for the horizons relevant for option valuation. This is of interest because although the long memory model may be a more adequate representation of the data, it is harder to implement.

4   Option Valuation

We now turn to the ultimate purpose of this paper, namely the valuation of derivatives on an underlying asset with dynamic variance components. For the purpose of option valuation we first derive the conditional moment generating function for the return process and then present the risk-neutral return dynamics.

4.1   The Moment Generating Function

For the return dynamics in this paper we can characterize the moment generating function (MGF) of the log stock price with a set of difference equations, using the techniques in Heston and Nandi (2000). Appendix A demonstrates that for the component GARCH model, the MGF defined by

$$f(t, T; \phi) \equiv E_t \left[\exp \left(\phi \ln \left(S_T \right) \right) \right]$$

can be written as

$$f(t, T; \phi) = S_t^\phi \exp[A_t + B_{1,t}(h_{t+1} - q_{t+1}) + B_{2,t} q_{t+1}] \tag{4.1}$$

with coefficients

$$A_t = A_{t+1} + r\phi - (\alpha B_{1,t+1} + \varphi B_{2,t+1}) - 1/2 \ln \left(1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1} \right) + B_{2,t+1} \omega$$

$$B_{1,t} = B_{1,t+1} \tilde{\beta} + \lambda \phi + 2 \frac{(\alpha \gamma_1 B_{1,t+1} + \varphi \gamma_2 B_{2,t+1} - 0.5\phi)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}$$

$$B_{2,t} = B_{2,t+1} \rho + \lambda \phi + 2 \frac{(\alpha \gamma_1 B_{1,t+1} + \varphi \gamma_2 B_{2,t+1} - 0.5\phi)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}$$

and terminal conditions

$$A_T = B_{1,T} = B_{2,T} = 0.$$

For the moment generating function in the GARCH(1, 1) case we refer to Heston and Nandi (2000).

4.2 The Risk-Neutral GARCH(1,1) Dynamic

The risk-neutral dynamics for the GARCH$(1,1)$ model are given in Heston and Nandi (2000)[15] as

$$
\begin{aligned}
R_{t+1} &= r - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z^*_{t+1} \\
h_{t+1} &= w + bh_t + a(z^*_t - c^*\sqrt{h_t})^2
\end{aligned}
\tag{4.2}
$$

with $c^* = c + \lambda + 0.5$ and $z^*_t \sim N(0,1)$.

4.3 The Risk-Neutral Component GARCH Dynamic

Appendix B demonstrates that the risk-neutral component GARCH dynamic is given by

$$
\begin{aligned}
h_{t+1} &= q_{t+1} + \tilde{\beta}^*\left(h_t - q_t\right) + \alpha\left(\left(z^*_t - \gamma^*_1\sqrt{h_t}\right)^2 - \left(1 + \gamma^{*2}_1 h_t\right)\right) \\
q_{t+1} &= \omega + \rho^* q_t + \varphi\left(\left(z^*_t - \gamma^*_2\sqrt{h_t}\right)^2 - \left(1 + \gamma^{*2}_2 h_t\right)\right)
\end{aligned}
\tag{4.3}
$$

where the risk neutral parameters are defined as follows

$$
\begin{aligned}
\tilde{\beta}^* &= \tilde{\beta} + \alpha\left(\gamma^{*2}_1 - \gamma^2_1\right) + \varphi\left(\gamma^{*2}_2 - \gamma^2_2\right) \\
\rho^* &= \rho + \alpha\left(\gamma^{*2}_1 - \gamma^2_1\right) + \varphi\left(\gamma^{*2}_2 - \gamma^2_2\right) \\
\gamma^*_i &= \gamma_i + \lambda + 0.5, \ \ i = 1, 2.
\end{aligned}
$$

The moment generating function for the risk-neutral component GARCH process is therefore equal to the one for the physical component GARCH process, setting $\lambda = -0.5$ and using the risk neutral parameters $\gamma^*_1, \gamma^*_2, \rho^*, \tilde{\beta}^*$ as well as ω, α and φ.

4.4 The Option Valuation Formula

Given the moment generating function and the risk-neutral dynamics, option valuation is relatively straightforward. We use the result of Heston and Nandi (2000) that at time t, a European call option with strike price K that expires at time T is worth

$$
\begin{aligned}
\text{Call Price} &= e^{-r(T-t)}E^*_t[Max(S_T - K, 0)] \\
&= S_t\left(\frac{1}{2} + \frac{1}{\pi}\int_0^\infty \text{Re}\left[\frac{K^{-i\phi}f^*(t,T;i\phi+1)}{i\phi S_t e^{r(T-t)}}\right]d\phi\right) \\
&\quad - Ke^{-r(T-t)}\left(\frac{1}{2} + \frac{1}{\pi}\int_0^\infty \text{Re}\left[\frac{K^{-i\phi}f^*(t,T;i\phi)}{i\phi}\right]d\phi\right)
\end{aligned}
\tag{4.4}
$$

where $f^*(t,T;i\phi)$ is the conditional characteristic function of the logarithm of the spot price under the risk neutral measure.

---

[15] For the underlying theory on risk neutral distributions in discrete time option valuation see Rubinstein (1976), Brennan (1979), Amin and Ng (1993), Duan (1995), Camara (2003), and Schroder (2004).

5 Empirical Results from Return Estimation

This section presents the first set of empirical results. We study the model estimated on a long time series of S&P500 index returns and emphasize a number of important stylized facts. We also report the differences between the models when using these parameters for option valuation and compare the models to a GARCH$(1,1)$-Jump model that allows for jumps in returns and variance. Finally, we analyze the differences in option valuation along various dimensions.

5.1 Properties of the Physical Return Process

Table 1 presents maximum likelihood estimates (MLE) of the physical model parameters obtained using returns data for 1962-2001. We use a long sample of returns on the S&P 500 because it is well-known that it is difficult to estimate GARCH parameters precisely using relatively short samples of return data. We compare the models using goodness-of-fit statistics, and we discuss differences in model properties. We present results for three models: the GARCH$(1,1)$ model (2.1), the component model (2.7) and the persistent component model (2.9). Almost all parameters are estimated significantly different from zero at conventional significance levels.[16] The price of risk, λ, is marginally significant in the case of the GARCH$(1,1)$ and not significant in the persistent component model. The log likelihood values indicate that the fit of the component model is much better than that of the persistent component model, which in turn fits much better than the GARCH$(1,1)$ model.

The dynamic variance models can be compared by assessing their persistence properties. The variance persistence in the GARCH$(1,1)$ model is defined by $\tilde{b} = b + ac^2$ from (3.1). In the component model, the total variance persistence is a confluence of the persistence in the two factors. If we substitute out q_{t+1} and q_t from the h_{t+1} equation in (2.7), then persistence can be computed as the sum of the coefficients on h_t and h_{t-1}. In this way, the component persistence formula can be derived to be $\rho + \tilde{\beta}(1 - \rho)$.

The improvement in fit for the component GARCH model over the persistent component GARCH model is perhaps somewhat surprising when inspecting the persistence of the component GARCH model. The persistence is equal to 0.9963. It therefore would appear that equating this persistence to 1, as is done in the persistent component model, is an interesting hypothesis, but apparently modeling these small differences from one is important. It must of course be noted that while the persistence of the long-run component (ρ) is 0.9896 for the component model as opposed to 1 for the persistent component model, the persistence of the short-run component ($\tilde{\beta}$) is 0.6437 versus 0.8822 and this may account for the differences in likelihood. Note that the persistence of the GARCH$(1,1)$ model is estimated at 0.9553, which is consistent with earlier literature. It is slightly lower than the estimate in Christoffersen, Heston and Jacobs (2006) and a bit higher than the average of the estimates in Heston and Nandi (2000).

In Figure 1, we illustrated differences across models in terms of variance term structures that are key for option valuation. Following Das and Sundaram (1999) and Bates (2000), we can use (4.1) to further investigate the conditional term structure of higher moments. The conditional

---

[16]One exception is the parameter w. The reason for this is probably that the point estimate for w is very small. This in turn is due to the enforcement of the constraint that w > 0, which is a sufficient condition for positive conditional variances. All standard errors are computed using the outer product of the gradient at the optimal parameter values.

skewness and excess kurtosis for maturity T are easily defined from the conditional cumulant generating function, $\ln(f(t, T; \phi))$, as follows

$$Skewness(t, T) = \frac{\partial^3 \ln f(t, T; \phi)/\partial\phi^3\big|_{\phi=0}}{\left(\partial^2 \ln f(t, T; \phi)/\partial\phi^2\big|_{\phi=0}\right)^{3/2}}, \quad Kurtosis(t, T) = \frac{\partial^4 \ln f(t, T; \phi)/\partial\phi^4\big|_{\phi=0}}{\left(\partial^2 \ln f(t, T; \phi)/\partial\phi^2\big|_{\phi=0}\right)^2}$$

We compute these moments by taking numerical derivatives of $\ln(f(t, T; \phi))$ in (4.1).

In Figure 3 we plot the term structure of skewness and kurtosis in the three GARCH models. The initial volatility is set to its long run value in the GARCH$(1, 1)$ and component GARCH models. In the persistent component model the initial volatility is set to the unconditional volatility from the component model. The parameter estimates are again taken from Table 1.

Figure 3 reveals important differences between the term structures of these moments for the GARCH$(1, 1)$ model, the component model and the persistent component model. While the term structures of skewness and kurtosis are hump-shaped for the GARCH$(1, 1)$ model over the maturities relevant for pricing the options in our sample, they are downward sloping for skewness and upward sloping for kurtosis in the persistent component model. For the component model, the minimum and the maximum respectively for the conditional skewness and kurtosis occur for options with approximately a six month maturity, but the skewness and kurtosis for longer maturities are very close to these extrema. These fundamental differences in higher moment term structures may have important implications for the option valuation properties across models.

Figure 4 further analyzes the component model's improvement in performance over the benchmark GARCH$(1, 1)$ model. It presents the 1990-1995 sample path for the spot variance in the GARCH$(1, 1)$ model and the component model, as well as the sample path for volatility components for the component models.[17] Figure 4 indicates that the mean zero short run component in the top-right panel adds high-frequency noise to the long-run component in the bottom-right panel. This results in a volatility dynamic for the component model in the top-left panel that is more noisy than the volatility dynamic for the GARCH$(1, 1)$ model in the bottom-left panel. This more noisy sample path suggests a higher value for the volatility of variance in the component model. An analysis of the persistent component model (not reported) supports similar conclusions, even though the sample paths are somewhat different from those in Figure 4.

We now investigate in more detail differences between the models in the modeling of the volatility (standard deviation) of the conditional variance, as well as differences in the modeling of the covariance and correlation between returns and variance. For option valuation, the conditional versions of these quantities and their variation through time are just as important as the unconditional versions.

For the GARCH$(1, 1)$ model the conditional variance of variance is

$$\begin{aligned} Var_t(h_{t+2}) &= E_t\left[h_{t+2} - E_t\left[h_{t+2}\right]\right]^2 \\ &= 2a^2 + 4a^2c^2 h_{t+1} \end{aligned} \tag{5.1}$$

---

[17]We plot results for the 1990-1995 subsample here because this is the sample that we will subsequently use for option valuation.

and for the component and persistent component models, the conditional variance of variance is

$$Var_t(h_{t+2}) = 2(\alpha + \varphi)^2 + 4(\gamma_1\alpha + \gamma_2\varphi)^2 h_{t+1}. \tag{5.2}$$

The three left panels in Figure 5 use the parameters from Table 1 to plot the standard deviation of variance in the three GARCH models. Notice that the standard deviation of variance in the component model is in general much higher than in the GARCH$(1,1)$ model and it is also more volatile. The average level of the conditional standard deviation of variance in the persistent component is in between that of the other two models. Table 1 reports the average volatility of variance during 1990-1995 in the row labeled "Average Vol of Var".

If we think of the option price as being a function of the spot variance, then we can view variation in the option prices as being driven by the volatility of variance. The volatility of variance is also related to kurtosis. Figure 5 thus shows that the component model is able to generate richer time-varying kurtosis dynamics than the GARCH$(1,1)$ model and thus potentially richer option price dynamics.

The conditional covariance between return and variance in the GARCH$(1,1)$ model is given by

$$
\begin{aligned}
Cov_t(R_{t+1}, h_{t+2}) &= E_t\left[(R_{t+1} - E_t[R_{t+1}])(h_{t+2} - E_t[h_{t+2}])\right] \tag{5.3} \\
&= E_t\left[\sqrt{h_{t+1}}z_{t+1}\left(az_{t+1}^2 - 2acz_{t+1}\sqrt{h_{t+1}} - a\right)\right] \\
&= -2ach_{t+1}.
\end{aligned}
$$

Conditional correlation is easier to interpret than conditional covariance. The conditional correlation in the GARCH$(1,1)$ model is

$$Corr_t(R_{t+1}, h_{t+2}) = \frac{-2c\sqrt{h_{t+1}}}{\sqrt{2 + 4c^2 h_{t+1}}} \tag{5.4}$$

where we have used the conditional variance of variance from (5.1).

The conditional covariance in the component model is

$$Cov_t(R_{t+1}, h_{t+2}) = -2(\gamma_1\alpha + \gamma_2\varphi)h_{t+1} \tag{5.5}$$

and the conditional correlation in the component model is thus given by

$$Corr_t(R_{t+1}, h_{t+2}) = \frac{-2(\gamma_1\alpha + \gamma_2\varphi)\sqrt{h_{t+1}}}{\sqrt{2(\alpha + \varphi)^2 + 4(\gamma_1\alpha + \gamma_2\varphi)^2 h_{t+1}}}. \tag{5.6}$$

The three right panels in Figure 5 plot the conditional correlation for the three models. The conditional correlation is clearly more negative in the component models than in the GARCH$(1,1)$ model, and furthermore the component correlation paths are much more volatile. Table 1 gives the average correlations during 1990-1995 as -79.40% (GARCH$(1,1)$), -88.49% (component), and -90.14% (persistent component).

For the component and persistent component models, we can also compute the conditional

correlations between the return and each volatility component separately

$$Corr_t(R_{t+1}, h_{t+2} - q_{t+2}) = \frac{-2\gamma_1\sqrt{h_{t+1}}}{\sqrt{2 + 4\gamma_1^2 h_{t+1}}} \quad (5.7)$$

$$Corr_t(R_{t+1}, q_{t+2}) = \frac{-2\gamma_2\sqrt{h_{t+1}}}{\sqrt{2 + 4\gamma_2^2 h_{t+1}}}.$$

We found that the conditional correlation of the return with the short-run variance component is on average more negative than the conditional correlation between the return and the long-run variance component (not reported). This difference can be traced to Table 1 where $\gamma_1 > \gamma_2$ in both models. The correlations with the long-run factor are relatively more negative in the persistent component model whereas the correlations with the short-run factor are relatively more negative in the component model. This can also be traced back to Table 1 where γ_1 is larger in the component model than in the persistent component model; whereas γ_2 is larger in the persistent component model.

Figure 6 shows the correlation between returns and variances from a different perspective. We plot the correlations from (5.4) and (5.6) against levels of the annualized conditional volatility. Notice that for all three models the relationship between the level of volatility and the correlation is negative. This is shown by Jones (2003) to be a desirable feature for option valuation and it is a feature missing in the standard Heston (1993) SV model where the correlation is constant. Interestingly, the Heston and Nandi (2000) GARCH(1, 1) model does have this negative relationship as Figure 6 shows. Figure 6 also shows that when fitted on the more general component model, the return data wants a correlation which is more negative than the simple GARCH(1, 1) model for all levels of volatility. The differences in correlation are quite large for the most common levels of volatility.

Figures 1-3 above showed that the component models are able to generate more flexible variance, skewness and kurtosis term structures than the simple GARCH(1,1) model. Figure 4 showed that the component models generate richer spot variance dynamics. We conclude from Figure 5 that the more flexible component models are capable of generating more flexible dynamics for the conditional correlation between returns and variance, and the conditional variance of variance. These dynamics are critically important for skewness and kurtosis dynamics which in turn are key for explaining the variation in index options prices. This is the topic to which we now turn.

5.2   Option Valuation Results

We use a rich sample of S&P 500 call options from the period 1990-1995. Following Bakshi, Cao and Chen (1997), we apply standard filters to the data. We only use Wednesday options data. Wednesday is the day of the week least likely to be a holiday. It is also less likely than other days such as Monday and Friday to be affected by day-of-the-week effects. For those weeks where Wednesday is a holiday, we use the next trading day. The decision to pick one day every week is to some extent motivated by computational constraints. Using only Wednesday data allows us to study a fairly long time-series, which is useful considering the highly persistent volatility processes. An additional motivation for using Wednesday data is that following the work of

Dumas, Fleming and Whaley (1998), several studies have used this setup.[18]

Table 2 presents descriptive statistics for the options data for 1990-1995 by moneyness and maturity. Panels A and B indicate that the data are standard. We can observe the implied volatility smirk from Panel C and it is clear that the slope of the smirk differs across maturities. Descriptive statistics for different sub-periods (not reported here) demonstrate that the slope changes over time, but that the smirk is present throughout the sample. There is substantial clustering in implied volatilities, and volatility is higher in the early part of the sample.

The last two rows of Table 1 compare the performance of the four models for option valuation. We use the MLE parameters to compute root mean squared errors (RMSEs) for the 1990-1995 option sample.[19] As was the case in the log likelihood comparison, the GARCH(1, 1) model is also the worst performer based on the option RMSE. But whereas the component model had a considerably better log likelihood than the persistent component model, the persistent component model now achieves a marginally better RMSE than the component model. Both component models have much better RMSEs than the GARCH(1,1) model.

Table 3 provides additional evidence on the option fit of the three models. We report option RMSE by moneyness and maturity. The top panel reports the RMSE for the GARCH(1, 1) model, while the two other panels report the ratio of the RMSE for the two other models to that of the GARCH(1, 1) The improvements of the component models over the GARCH(1, 1) model are fairly robust across maturity and moneyness. Importantly, the component models are never much worse than the GARCH(1, 1) model and they fail to improve on the GARCH(1, 1) model only for short term deep-in-the-money call options. This finding leads us to consider jumps in returns, which by way of adding non-normality to the conditional density may lead to improvements in the valuation of short-term options.

5.3 Comparison with a GARCH(1,1)-Jump Model

The Heston-Nandi GARCH(1, 1) model is a useful first benchmark, but it has well-known empirical biases. These biases are similar to those displayed by the Heston (1993) model. The continuous-time literature has attempted to improve the performance of the Heston (1993) model by adding to it (potentially correlated) jumps in returns and volatility, and this strategy has been partly successful. Poisson-normal jumps in returns and volatility lead to improved modeling of the historical time series of returns. When model parameters are estimated using the cross-section of option prices, Poisson-normal jumps usually do not lead to improved model fit, but Broadie, Chernov and Johannes (2007) find evidence of the importance of jumps for option pricing when imposing consistency between the physical and risk-neutral parameters.[20]

In the discrete-time literature, some studies have attempted to address these model biases by combining conditional heteroskedasticity with non-normal innovations. This strategy may seem very different from including jumps in the return process, but both approaches essentially intro-

---

[18]See for instance Heston and Nandi (2000).

[19]As the price of risk parameter, λ, is poorly estimated in Table 1, and in order to keep the persistence at unity under both measures in the persistent model, we simply set $\lambda = -0.5$ across models when computing option prices.

[20]For evidence on the importance of Poisson-normal jumps, see for example Andersen, Benzoni and Lund (2002), Bakshi, Cao and Chen (1997), Bates (1996, 2000), Chernov, Gallant, Ghysels and Tauchen (2003), Eraker, Johannes and Polson (2003), Eraker (2004) and Pan (2002).

duce conditional nonnormalities in the daily return distribution. However, Christoffersen, Heston and Jacobs (2006) find that inverse Gaussian innovations do not improve out-of-sample model fit. To provide a more challenging benchmark for the component model, we therefore combine the GARCH(1, 1) dynamic with a Poisson jump process similar to the one used in the continuous-time option valuation literature.[21] We refer to the resulting model as the GARCH(1, 1)-Jump model. Consider the return and variance dynamics

$$
\begin{aligned}
R_{t+1} &= r + \lambda_z h_{t+1} + \lambda_y \chi + \sqrt{h_{t+1}} z_{t+1} + y_{t+1} - \mu \chi \qquad (5.8) \\
h_{t+1} &= w + b h_t + a \left(z_t + \frac{y_t}{\sqrt{h_t}} - c \sqrt{h_t} \right)^2
\end{aligned}
$$

The innovation to returns consists of a diffusive component z_t and a jump component y_t. The return dynamic has been designed to ensure that the expected log excess return equals the sum of the diffusive and jump risk premium. The diffusive innovation z_t is assumed to be i.i.d. $N(0,1)$. We consider a simple jump structure where y_t is a compound Poisson process defined as

$$
y_t = \sum_{j=1}^{N_t} X_t^{(j)}
$$

with

$$
X_t^{(j)} \sim N\left(\mu, \tau^2\right) \quad \text{for } j = 1, 2, ..., N_t.
$$

and N_t is a Poisson random variable with constant intensity χ. We therefore have

$$
y_t \sim CPoisson(\chi, \mu, \tau^2)
$$

For the purpose of comparison with the Heston-Nandi GARCH(1,1) model and our component model, we have allowed for time-varying variance only in the diffusive term, and we have specified an affine variance dynamic h_{t+1}. The model nests the Heston-Nandi GARCH(1,1) model as a special case for $\chi = 0$.

It is relatively straightforward to estimate this model by maximum likelihood. Using the Esscher transform,[22] it can be shown that the resulting risk-neutral dynamic has the same functional form with parameterization

$$
\begin{aligned}
R_{t+1} &= r - \frac{1}{2} h_{t+1} + \lambda_y^* \chi^* + \sqrt{h_{t+1}} z_{t+1}^* + y_{t+1}^* - \mu^* \chi^* \qquad (5.9) \\
h_{t+1} &= w + b h_t + a \left(z_t^* + \frac{y_t^*}{\sqrt{h_t}} - c^* \sqrt{h_t} \right)^2
\end{aligned}
$$

where $z_t^* \sim N(0,1)$ and $y_t^* \sim CPoisson(\chi^*, \mu^*, \tau^2)$, with $c^* = c + \frac{1}{2} + \lambda$, $\lambda_y^* = (\lambda_y - \mu)\frac{\chi}{\chi^*} + \mu^*$, $\mu^* = \mu + \pi_y \tau^2$, $\chi^* = \chi \exp(\frac{\pi_y^2 \tau^2}{2} + \pi_y \mu)$, and π_y is the market price of jump risk which can

---

[21] See Maheu and McCurdy (2004) and Duan, Ritchken and Sun (2005, 2006) for alternative specifications of the jump process in a discrete-time setup.

[22] See Huang and Wu (2004) and Carr and Wu (2004) for applications of the Esscher transform to option valuation. See Dai, Le and Singleton (2006) for applications to fixed income pricing.

be solved for numerically from the physical MLE estimates. The term $(\lambda_y^* - \mu^*)\chi^*$ makes the discounted price process a martingale. No closed-form solution exists for the GARCH$(1,1)$-Jump model so that option prices must be computed by Monte Carlo simulation.

The structure of the jump component in the GARCH$(1,1)$-Jump model in (5.9) is similar to that of several papers in the continuous-time option valuation literature, and it is known in this literature as the SVCJ process with perfectly correlated return and volatility jumps. See for example Broadie, Chernov and Johannes (2007), Eraker (2004) and Eraker, Johannes and Polson (2003).[23]

Table 4 reports the empirical results for the GARCH$(1,1)$-Jump model. Panel A reports the parameter estimates from maximum likelihood estimation on the sample of daily S&P500 returns used in Table 1. Most parameters are rather precisely estimated, except for λ_z and especially λ_y. In order to prevent these imprecise estimates from influencing the comparison with other models when valuing options, we value options using a λ_y which yields a zero jump risk premium. The risk-neutral parameters are therefore equal to the physical parameter estimates in Table 4, except for λ_z^*, which is equal to -0.5. This approach ensures comparability with the GARCH models, where we impose $\lambda = -0.5$ to impose a zero risk premium. Parameter estimates are generally in line with the literature. The average jump size is -1.2% and the jump intensity is equal to 0.0098, which can be annualized by multiplying by 252, indicating approximately 2.5 jumps per year. These estimates are intuitively plausible, and largely in line with the literature.[24] The standard deviation of the jump τ is almost identical to the estimate in Eraker, Johannes and Polson (2003). The estimated mean jump size μ is similar to the estimate in Chernov, Gallant, Ghysels and Tauchen (2003), but about half of the estimate in Eraker, Johannes and Polson (2003) and Eraker (2004). Our estimated jump intensity χ is higher than the estimate in those papers. Presumably our smaller mean jump and higher jump intensity estimates are due to the fact that compared to these papers we use a more extended sample period, and that consequently the 1987 crash is relatively less important in determining the parameters.

Notice that the log likelihood value for the GARCH$(1,1)$-Jump model is considerably larger than for the three models in Table 1. The GARCH$(1,1)$-Jump model thus provides an excellent description of the conditional density for daily S&P500 returns. The Option RMSE for the GARCH$(1,1)$-Jump model is $1.939, which is much better than the $2.236 for the GARCH$(1,1)$ in Table 1, but somewhat worse than the $1.706 and $1.705 for the component model and persistent model respectively.

Panel B in Table 4 shows the ratio of the RMSE of the GARCH$(1,1)$-Jump to the GARCH$(1,1)$ model. The GARCH$(1,1)$-Jump model in general performs better than the GARCH$(1,1)$, and its performance is especially strong for deep out-of-the money and for long-maturity options. Somewhat surprisingly, the GARCH$(1,1)$-Jump model does not perform well for short-maturity options. Given that the component model also does not perform well in this dimension, the component structure and this particular jump structure may not be complementary in addressing the biases of the GARCH$(1,1)$ model. It is of course possible that the GARCH$(1,1)$-Jump model would perform better for short-maturity options if it was estimated using option data

---

[23]Our setup also allows us to estimate a process with one extra parameter that captures less than perfect correlation, but this parameter is often estimated imprecisely in the continuous-time literature, and so we do not estimate it here.

[24]See Table 1 in Broadie, Chernov and Johannes (2007) for a good overview of estimates available in the literature.

as the jump parameters may be difficult to estimate precisely from returns only. The strong performance by the jump model for long-maturity options suggest that the jump-in-volatility feature incorporated in this model is useful. Jumps in volatility, which in turn is persistent, will yield non-normal behavior in the conditional distribution at long horizons. This long-horizon non-normality will help fit an option smile that persists at longer horizons.

The performance of the GARCH(1,1)-Jump model is encouraging. We also found that alternative specifications of the jump risk premium can lead to improvements in performance, but it is not clear how to calibrate this risk premium and how to ensure a meaningful comparison with the component models. Different specifications for the jump process may provide an even better option pricing fit, but such an investigation is beyond the scope of this paper. In summary, we conclude that while the GARCH(1, 1)-Jump model significantly outperforms the GARCH(1,1) model for the purpose of option valuation, the component model compares very favorably with the GARCH(1, 1)-Jump model. It is important to keep in mind that the component model does so with eight parameters, and the persistent component model with seven, whereas the GARCH(1,1)-Jump model contains nine parameters. Nevertheless, because the GARCH(1,1)-Jump model provides a more challenging benchmark for the component model, we analyze it in some more detail below.

5.4 Analyzing Option Valuation Performance

It must be emphasized that the component models' performance is remarkable and to some extent surprising. First, the GARCH(1, 1) model is a good benchmark which itself has a very solid empirical performance (see Heston and Nandi (2000)). The model captures important stylized facts about option prices such as volatility clustering and the leverage effect (or equivalently negative skewness). When estimating models from option prices, Christoffersen and Jacobs (2004) find that GARCH models with richer news-impact parametrization do not improve the model fit out-of-sample. Christoffersen, Heston and Jacobs (2006) find that a GARCH model with non-normal innovations improves the model's fit in-sample and for short out-of-sample horizons, but not for long out-of-sample horizons. Furthermore, the analysis in Table 4 demonstrates that the component model also provides a better option fit than the GARCH(1, 1) model augmented with Poisson-normal jumps.

We now provide some more insight into the improved performance of the component models by analyzing the differences across models along three critical dimensions: the (spot) volatility level, the volatility term structure and the modeling of the smirk. First, component models may better match the volatility patterns over time. We investigate this by comparing the differences in the time paths between implied volatilities from the data and the models. Second, it may be the case that the component models more adequately capture the term structure of volatility than the GARCH(1, 1) model. We investigate this by comparing the models' term structures of implied volatility for at-the-money options. Third, it may be the case that the component models better capture the implied volatility smirk at various maturities. We study the differences between the models in this dimension in different volatility regimes.

Figure 7 presents the average weekly at-the-money[25] implied volatility bias (average observed market implied volatility less average model implied volatility) over the 1990-1995 option sam-

---

[25]Here and below we define at-the-money to be options with strike price within 2.5% of the underlying index.

ple, using the MLE estimates from Table 1. Clearly the component models outperform the GARCH$(1,1)$ model in this dimension: The GARCH$(1,1)$ model shows significant underpricing (positive bias) during the high volatility episode in 1990-1991 and extended periods of overpricing (negative bias) during the low volatility period in 1993-1995. In comparison, the component model has smaller (positive) bias in 1990-1991 and also smaller (negative) bias in 1993-1995, suggesting that it is much better able to capture the dynamics of market volatility. The persistent model has the smallest (positive) bias in late 1990 but instead has significant (negative) bias in early 1990 and in late 1991. The GARCH$(1,1)$-Jump model performs the best during the 1993-1995 low volatility period. Somewhat surprisingly however, it has a rather large bias during the high volatility 1990-1991 episode.

Broadie, Johannes and Chernov (2007) report that the implied volatility from at-the-money options on average is higher than the realized volatility from returns. It is therefore interesting to ask whether this feature of the data is captured by the GARCH option pricing models. We compute the average implied volatility for at-the-money options during 1990-1995 to be 12.53%. The average at-the-money implied volatility in the GARCH$(1,1)$ model is 13.11% compared with 12.68% for the component and 13.49% for the persistent component model. The component model thus comes very close to matching the average implied volatility in the at-the-money options. Compare these implied volatility numbers with the average annual vols in Table 1, which were 12.06% for the GARCH$(1,1)$, 11.74% for the component model, and 12.39% for the persistent component model. All three models thus produce average implied at-the-money volatility that is roughly 1% higher than the average realized volatility in returns during the 1990-1995 period.

Figure 8 studies the implied volatility term structure for at-the-money options in the four models. We again use the MLE parameters from Table 1 to compute option prices. To study the importance of different volatility regimes, we study at-the-money options in three different sample periods. The first sample period is from September 1, 1995 until December 31, 1995. The average value of the VIX over this period is 12.70%, and we therefore think of it as a low volatility period. The second sample period is from August 1, 1991 until November 31, 1991, with an average VIX value of 16.66%, and it represents a medium volatility period. The third sample period is from August 1, 1990 until November 31, 1990, which is a high volatility period with an average VIX value of 27.97%. The main conclusion is that the differences between the models are sometimes very pronounced, but that these differences depend on the volatility regime. Whereas the fit of the GARCH(1,1) model and the component GARCH model do not differ much in the medium volatility regime, the component GARCH model significantly outperforms the GARCH(1,1) model in the low and high volatility periods. The GARCH-Jump model captures some of the same biases as the component model, but perhaps does a bit poorer in the high volatility period. The persistent component model clearly performs very well in the high volatility regime.

These results can be explained as follows: in the component model, the initial volatility is much more important for the valuation of longer maturity options than in the GARCH$(1,1)$ model, and even more so in the persistent component model. Put differently, in the GARCH$(1,1)$ model, today's level of volatility has virtually no impact on the implied volatility for 1-year to maturity options. For the component model, the initial volatility has an effect on the implied volatility for 1-year maturity options, and in the fully persistent model the effect of initial volatil-

ity is as large at the 1-year maturity as it is at short maturities. The more persistent component therefore allows the component model to capture the variation in longer-maturity implied volatilities. Moreover, the transitory component allows the component model to also adequately capture short-maturity implied volatilities in different volatility regimes.

Figure 9 builds on these findings, and also analyzes a third source of differences in fit between the models. For each of the models, we plot the implied volatility smirks across moneyness for four different maturity buckets: 15-30, 30-60, 90-120 and 250-300 days to maturity. Following the exposition in Figure 8, we repeat the analysis for three different volatility periods. The main conclusion from Figure 9 is that whereas the models greatly differ with respect to the level of the smirk, the slope of the smirk does not seem to differ much between models, although it is steeper for the GARCH(1,1)-Jump model in some cases.[26] Also note that all models tend to undervalue in-the-money call options across volatility regimes for the shortest maturity (top row panels) and for most maturities in the high-volatility regimes (right-side panels). Even richer volatility and jump specifications may thus be needed.

We conclude that there are important differences between the $GARCH(1,1)$, $GARCH(1,1)$-Jump, component and persistent component models in terms of the path of spot volatility, the term structure of volatility, and the level of the volatility smirk. Importantly, the $GARCH(1,1)$-Jump, component and persistent component models outperform the benchmark $GARCH(1,1)$ model in virtually every dimension. The $GARCH(1,1)$-Jump model outperforms the component model in some dimensions, such as the fit of the volatility path in the less volatile part of the sample periods, but it underperforms the component model in the more volatile sample period. The performance of the component models for short-maturity and long-maturity options in high volatility regimes seems to be the main reason for their success. The model is able to successfully address both these dimensions thanks to the different time series properties of the transitory and persistent components.

6 Estimation Using Option Price Information

So far we have used the option price information only to evaluate the different models. However, it stands to reason that the observed option prices should be helpful in estimating the models as well. In this section we therefore implement the $GARCH(1,1)$, component, and persistent component model by minimizing the mean squared option valuation error rather than maximizing the log likelihood on returns as we did in Table 1.

We obtain parameters by minimizing the dollar mean squared error

$$\$MSE = \frac{1}{N^T} \sum_{t,i} \left(C_{i,t}^D - C_{i,t}^M \right)^2 \tag{6.1}$$

where $C_{i,t}^D$ is the market price of option i at time t, $C_{i,t}^M$ is the model price, and $N^T = \sum_{t=1}^{T} N_t$.

---

[26]While the GARCH(1,1)-Jump model generates a steeper smirk in some cases, it underperforms the other models with respect to matching the level of implied volatilities. These two aspects of the model's performance explain the variations in performance across moneyness and maturity in Table 4B.

T is the total number of days included in the sample and N_t the number of options included in the sample at date t. The variance dynamic is used to update the variance from one Wednesday to the next using daily returns and the option valuation formula in (4.4) is used to compute the model prices on each Wednesday. The volatility updating rule is applied to the 500 days predating the first Wednesday used in the estimation exercise, and it is initialized at the model's unconditional variance.

Needless to say, this nonlinear least squares (NLS) estimation technique is much more computationally intensive than the simple MLE on returns in Table 1. For each function evaluation performed by the numerical optimizer, the model price must be calculated for every option in the sample. The optimizer performs many function evaluations for each parameter update and consequently it is crucial to be able to compute option prices quickly and reliably. The pricing formula in (4.4) makes this estimation technique feasible. As we unfortunately do not have a closed-form pricing formula for the GARCH(1,1)-Jump model, we do not consider that model in this section.

Table 5 presents parameter estimates obtained using the 1990-1992 options data and in-sample RMSEs for the 1990-1992 data, as well as out-of-sample RMSEs using the 1993 data. Note that the shortest maturity is seven days because options with very short maturities were filtered out. Table 6 presents parameter estimates obtained using options data for 1992-1994, as well as 1992-1994 in-sample and 1995 out-of-sample option RMSEs. Table 7 presents RMSE results by moneyness and maturity for the two in-sample and two out-of-sample periods.[27]

In Table 5 we present results for the 1990-1992 period (in-sample) and the 1993 period (out-of-sample). The standard errors indicate that almost all parameters are estimated significantly different from zero.[28] There are some differences with the parameters estimated from returns in Table 1, but the parameters are mostly of the same order of magnitude. This is also true for critical determinants of the models' performance, such as average annual volatility, average volatility of variance and average return correlation. Note also that the persistence of the short-run components and the long-run components is not dramatically different from Table 1. The persistence of the GARCH(1,1) process is higher than in Table 1 though. In fact, it is interesting to note that the persistence of the GARCH(1,1) model and the component GARCH model is close to one. This of course motivates the use of the persistent component model, where the persistence is restricted to be one. Note also that the average correlation between return and volatility is now close to minus one in all three models.

Table 5 contains two sets of RMSEs. The RMSEs in the leftmost columns (NLS) are obtained using the parameter values in the Table. In the rightmost column, we report RMSEs based on parameter values obtained from MLE in Table 1. First consider the RMSEs obtained using NLS. In the in-sample 1990-1992 period, the RMSE of the component model is 89.7% of that of the benchmark GARCH(1,1) model. For the out-of-sample 1993 period, the ratio of the RMSEs is 76.5%. For the persistent component model, the ratios are 94.8% and 93.3% respectively. Using the MLE estimates, the relative RMSEs are similar for the component model: 84.9% in 1990-1992 and 71.0% in 93. Using the MLE estimates, the persistent component model performs relatively worse in 1990-1992 with 113.0%, but better in 1993 with 56.5% of the RMSE in the

---

[27]Notice from the risk-neutral dynamics (4.2) and (4.3) that the parameter λ is not separately identified using option prices. We therefore simply set $\lambda = -0.5$ and we do not report λ in Tables 5-6.

[28]The standard errors are again computed using the outer product of the gradient at the optimum.

GARCH$(1, 1)$. When comparing across MLE and NLS estimates the RMSEs from NLS are typically much smaller than those from MLE. The information in option prices is clearly very valuable for estimating the models. The only example where the RMSE from MLE is close to that of the NLS counterpart is for the persistent component model in the 1993 out-of-sample period.

Table 6 presents the results for the 1992-1994 period (in-sample) and the 1995 period (out-of-sample). The results largely confirm those obtained in Table 5. The most important difference is that the in-sample and out-of-sample performance of the component model is even better relative to the benchmark, as compared with the results in Table 5. When using NLS estimates, the component model's RMSE is 77.3% of that of the GARCH$(1, 1)$ model for the 1992-1994 in-sample period, and for the 1995 out-of-sample period the ratio is 79.2%. For the persistent component model the ratios are 95.6% and 95.7%. When using MLE estimates the component model is 76.2% of the GARCH$(1, 1)$ in 1992-1994 and 69.2% in 1995. The persistent model performs very well relative to the GARCH$(1, 1)$ MLE generating a 70.1% relative RMSE for 1992-1994 and 45.0% in 1995.

Comparing RMSEs across NLS and MLE parameters, we again find that the option prices add important information and drive the NLS RMSEs down from their MLE levels. Again, the only case where the RMSE from MLE comes close to that from NLS is for the out-of-sample persistent model. Interesting differences with Table 5 are that the persistence of the short-run component is higher, and that the persistence of the GARCH$(1, 1)$ process in Table 6 is lower than in Table 5 but in line with the MLE estimate in Table 1.

Table 7 provides a more detailed analysis of the valuation errors by presenting RMSE results by moneyness and maturity, using the parameter estimates from Tables 5 and 6. The table contains the RMSE in dollars for the GARCH$(1, 1)$ model. To facilitate the interpretation of the table, the RMSEs for the component and persistent component models are normalized by the corresponding RMSE for the GARCH$(1, 1)$ model. It is clear that an overall RMSE which is not too different across the three models in Tables 5 and 6 can mask large differences in the models' performance for a given moneyness or maturity. Inspection of the out-of-sample results in Panels B and D is especially instructive. We conclude that the improved out-of-sample performance of the component models is due to the improved valuation of long-maturity options. This is an interesting affirmation of the intuition obtained previously using return-based estimates in Figures 1-2 and 8-9. The richer volatility dynamic in the component model enables more accurate modeling of variations in long-maturity option prices.

Overall, we conclude that based on the parameter values obtained using NLS, the performance of the component GARCH model is very impressive. Its RMSE is between 76.5% and 89.7% of the RMSE of the benchmark GARCH$(1, 1)$ model. The performance of the persistent component model is less impressive, both in-sample and out-of-sample. However the persistent component model performs relatively well in the out-of-sample experiments when MLE parameters are used. This suggests that the persistent component model may be valuable for option valuation in cases where no option price information is available for parameter estimation.

7 Conclusion and Directions for Future Work

This paper presents a new option valuation model based on the work by Engle and Lee (1999) and Heston and Nandi (2000). The empirical performance of the new variance component model is significantly better than that of the benchmark GARCH(1, 1) model, in-sample as well as out-of-sample, and regardless of the information used in estimation. This is an important finding because the literature has demonstrated that it is difficult to find empirical models that improve on the GARCH(1, 1) model or the Heston (1993) model. We also compare the component model to a GARCH(1, 1)-Jump model, which combines conditional heteroskedasticity with Poisson-normal jumps. The GARCH(1, 1)-Jump model achieves a better statistical fit to daily returns than the component model, but the component model performs better when using the parameter estimates to fit options. Overall, the GARCH(1, 1)-Jump model and the variance component model remedy some of the same biases left unaddressed by the GARCH(1, 1) model.

An important aspect of the model's improved performance is that its richer parameterization allows for improved joint modeling of long-maturity and short-maturity options. The model captures the stylized fact that shocks to current conditional volatility impact on the conditional variance forecast up to a year in the future, which results in a very different implied volatility term structure for at-the-money options. The component model also results in a different path for spot volatility compared to the GARCH(1, 1) model.

Because the estimated persistence of the model is close to one, we also investigate a special case of our model in which shocks to the variance never die out. When estimating model parameters by maximum likelihood using a historical time series of returns, the persistent component model is somewhat inferior to the component model when judged by the likelihood criterion. When the MLE parameters are used to price options, the persistent component model performs similarly to the component model in terms of overall fit. This is partly due to its ability to model long-maturity options, especially in high volatility periods.

When model parameters are estimated from option prices, the component model significantly outperforms the other models both in and out-of-sample. We also find that for a given model the parameters obtained from historical return data always lead to higher RMSEs than the parameters directly estimated from option data.

Given the success of the proposed component models, a number of further extensions to this work are warranted. First, the empirical performance of the model should of course be validated using other datasets. In particular, it would be interesting to test the model using LEAPS data, because the model may excel at modeling long-maturity LEAPS options. In this regard a direct comparison between component and fractionally integrated volatility models may be interesting. It could also be useful to combine the stylized features of the model with other modeling components that improve option valuation. One interesting experiment could be to replace the Poisson jump innovations considered in this paper with another non-Gaussian distribution. Combining the model in this paper with the inverse Gaussian shock model in Christoffersen, Heston and Jacobs (2006) may be a viable approach. Finally, in this paper we have proposed a component model that gives a closed form solution using results from Heston and Nandi (2000) who rely on an affine GARCH model. We believe that this is a logical first step, but the affine structure of the model may be restrictive in ways that are not immediately apparent. It may therefore prove worthwhile to investigate non-affine variance component models. A final remark is that the performance of the GARCH(1,1)-Jump model, which we use here as a benchmark for the

component models, is encouraging. It may prove worthwhile to further analyze its performance, for instance by estimating the model using option data.

8 Appendix A

This Appendix derives the moment generating function for the component GARCH process. The component GARCH process is given by

$$
\begin{aligned}
\ln(S_{t+1}) &= \ln(S_t) + r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1} \\
h_{t+1} &= q_{t+1} + \tilde{\beta} (h_t - q_t) + \alpha \left((z_t - \gamma_1 \sqrt{h_t})^2 - (1 + \gamma_1^2 h_t) \right) \\
q_{t+1} &= \omega + \rho q_t + \varphi \left((z_t - \gamma_2 \sqrt{h_t})^2 - (1 + \gamma_2^2 h_t) \right).
\end{aligned}
$$

Let $x_t = \ln(S_t)$. For convenience we will denote the time t conditional generating function of S_T (or equivalently the conditional moment generating function (MGF) of x_T) by f_t instead of the more cumbersome $f(t; T, \phi)$. By definition

$$
f_t = E_t[\exp(\phi x_T)]. \tag{8.1}
$$

We shall guess that the moment generating function has the log-linear form. We use the more parsimonious notation A_t to indicate $A(t; T, \phi)$ and similarly for $B_{1,t}$ and $B_{2,t}$ and write

$$
f_t = \exp\left(\phi x_t + A_t + B_{1,t}(h_{t+1} - q_{t+1}) + B_{2,t} q_{t+1} \right). \tag{8.2}
$$

We have the terminal condition $A_T = B_{1,T} = B_{2,T} = 0$. Applying the law of iterated expectations to f_t in (8.1) we get

$$
f_t = E_t[f_{t+1}] = E_t \exp\left(\phi x_{t+1} + A_{t+1} + B_{1,t+1}(h_{t+2} - q_{t+2}) + B_{2,t+1} q_{t+2} \right).
$$

Substituting the dynamics of x_{t+1} gives

$$
\begin{aligned}
f_t &= E_t \exp \left(\begin{array}{c} \phi(x_t + r) + \phi \lambda h_{t+1} + \phi \sqrt{h_{t+1}} z_{t+1} + A_{t+1} + B_{1,t+1}(h_{t+2} - q_{t+2}) + \\ B_{2,t+1} q_{t+2} \end{array} \right) \\
&= E_t \exp \left(\begin{array}{c} \phi(x_t + r) + \phi \lambda h_{t+1} + \phi \sqrt{h_{t+1}} z_{t+1} + A_{t+1} + \\ B_{1,t+1} \left(\tilde{\beta} (h_{t+1} - q_{t+1}) + \alpha \left((z_{t+1} - \gamma_1 \sqrt{h_{t+1}})^2 - (1 + \gamma_1^2 h_{t+1}) \right) \right) + \\ B_{2,t+1} \left(\omega + \rho q_{t+1} + \varphi \left((z_{t+1} - \gamma_2 \sqrt{h_{t+1}})^2 - (1 + \gamma_2^2 h_{t+1}) \right) \right) \end{array} \right) \\
&= E_t \exp \left(\begin{array}{c} \phi(x_t + r) + \phi \lambda h_{t+1} + A_{t+1} + B_{1,t+1} \tilde{\beta} (h_{t+1} - q_{t+1}) + \\ B_{2,t+1}(\omega + \rho q_{t+1}) - (\alpha B_{1,t+1} + \varphi B_{2,t+1}) + \\ (\alpha B_{1,t+1} + \varphi B_{2,t+1}) \left(z_{t+1} - \frac{\alpha \gamma_1 B_{1,t+1} + \varphi \gamma_2 B_{2,t+1} - 0.5\phi}{(\alpha B_{1,t+1} + \varphi B_{2,t+1})} \sqrt{h_{t+1}} \right)^2 - \\ \frac{(\alpha \gamma_1 B_{1,t+1} + \varphi \gamma_2 B_{2,t+1} - 0.5\phi)^2}{(\alpha B_{1,t+1} + \varphi B_{2,t+1})} h_{t+1} \end{array} \right).
\end{aligned}
$$

Using the following result for a standard normal variable, z,

$$E\left[\exp(x(z+y)^2)\right] = \exp(-\frac{1}{2}\ln(1-2x) + xy^2/(1-2x))$$

we get

$$f_t = E_t \exp \begin{pmatrix} \phi(x_t+r) + A_{t+1} - (\alpha B_{1,t+1} + \varphi B_{2,t+1}) \\ -1/2\ln(1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}) + B_{2,t+1}\omega + \\ B_{1,t+1}\tilde{\beta}(h_{t+1} - q_{t+1}) + B_{2,t+1}\rho q_{t+1} + \\ \left(\lambda\phi + 2\frac{(\alpha\gamma_1 B_{1,t+1} + \varphi\gamma_2 B_{2,t+1} - 0.5\phi)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}\right) h_{t+1} \end{pmatrix}. \tag{8.3}$$

Matching terms in (8.3) and (8.2) gives

$$A_t = A_{t+1} + r\phi - (\alpha B_{1,t+1} + \varphi B_{2,t+1}) - 1/2\ln(1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}) + B_{2,t+1}\omega$$

$$B_{1,t} = B_{1,t+1}\tilde{\beta} + \lambda\phi + 2\frac{(\alpha\gamma_1 B_{1,t+1} + \varphi\gamma_2 B_{2,t+1} - 0.5\phi)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}$$

$$B_{2,t} = B_{2,t+1}\rho + \lambda\phi + 2\frac{(\alpha\gamma_1 B_{1,t+1} + \varphi\gamma_2 B_{2,t+1} - 0.5\phi)^2}{1 - 2\alpha B_{1,t+1} - 2\varphi B_{2,t+1}}.$$

9 Appendix B

The physical component GARCH dynamic is given by

$$\ln(S_{t+1}) = \ln(S_t) + r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1}$$

$$h_{t+1} = q_{t+1} + \tilde{\beta}(h_t - q_t) + \alpha\left(\left(z_t - \gamma_1\sqrt{h_t}\right)^2 - \left(1 + \gamma_1^2 h_t\right)\right) \tag{9.1}$$

$$q_{t+1} = \omega + \rho q_t + \varphi\left(\left(z_t - \gamma_2\sqrt{h_t}\right)^2 - \left(1 + \gamma_2^2 h_t\right)\right) \tag{9.2}$$

Under the risk neutral measure, we need $E^*[S_{t+1}/S_t] = \exp(r)$, which requires that

$$\ln(S_{t+1}) = \ln(S_t) + r - 0.5h_{t+1} + \sqrt{h_{t+1}} z_{t+1}^*.$$

This implies in turn that

$$z_{t+1}^* = z_{t+1} + (\lambda + 0.5)\sqrt{h_{t+1}}. \tag{9.3}$$

We also need to ensure that the conditional variances are the same under the two measures

$$Var_t(R_{t+1}) = Var_t^*(R_{t+1}).$$

We therefore need to have the same variance innovations under the two measures. Thus we need

$$\left(z_t - \gamma_i\sqrt{h_t}\right)^2 = \left(z_t^* - \gamma_i^*\sqrt{h_t}\right)^2 \qquad i = 1, 2$$

which can be achieved by defining new risk neutral parameters $\gamma_i^* = \gamma_i + \lambda + 0.5$, $i = 1, 2$.

Consider the following candidate for the risk-neutral component GARCH dynamic

$$h_{t+1} \;=\; q_{t+1} + \tilde{\beta}^* \left(h_t - q_t\right) + \alpha \left(\left(z_t^* - \gamma_1^* \sqrt{h_t}\right)^2 - \left(1 + \gamma_1^{*2} h_t\right) \right) \tag{9.4}$$

$$q_{t+1} \;=\; \omega + \rho^* q_t + \varphi \left(\left(z_t^* - \gamma_2^* \sqrt{h_t}\right)^2 - \left(1 + \gamma_2^{*2} h_t\right) \right) \tag{9.5}$$

where $z_t^* \sim N(0,1)$ and the risk neutral parameters are defined as follows

$$\tilde{\beta}^* \;=\; \tilde{\beta} + \alpha \left(\gamma_1^{*2} - \gamma_1^2\right) + \varphi \left(\gamma_2^{*2} - \gamma_2^2\right) \tag{9.6}$$
$$\rho^* \;=\; \rho + \alpha \left(\gamma_1^{*2} - \gamma_1^2\right) + \varphi \left(\gamma_2^{*2} - \gamma_2^2\right).$$

For this candidate risk-neutral dynamic to be valid, we have to verify that it is consistent with (9.1) and (9.2). Using (9.3), (9.6) and (9.5) in (9.4) we get

$$h_{t+1} = \omega + \rho q_t + \varphi \left(\left(z_t - \gamma_2 \sqrt{h_t}\right)^2 - \left(1 + \gamma_2^2 h_t\right) \right) + \tilde{\beta} \left(h_t - q_t\right) + \alpha \left(\left(z_t - \gamma_1 \sqrt{h_t}\right)^2 - \left(1 + \gamma_1^2 h_t\right) \right)$$

which is identical to what we get using the physical dynamic in (9.1) and (9.2).

References

Adrian, T., Rosenberg, J., 2005. Stock returns and volatility: pricing the long-run and short-run components of market risk. Unpublished working paper, the Federal Reserve Bank of New York.

Ait-Sahalia, Y., Lo, A., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. Journal of Finance 53, 499-547.

Alizadeh, S., Brandt, M., Diebold, F., 2002. Range-based estimation of stochastic volatility models. Journal of Finance 57, 1047-1091.

Andersen, T., Bollerslev, T., Diebold, F., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 529-626.

Amin, K., Ng, V., 1993. ARCH processes and option valuation. Unpublished working paper, University of Michigan.

Andersen, T., Benzoni, L., Lund, J., 2002. An empirical investigation of continuous-time equity return models. Journal of Finance 57, 1239-1284.

Baillie, R., Bollerslev, T., Mikkelsen, H., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 74, 3-30.

Bakshi, C., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. Journal of Finance 52, 2003-2049.

Bates, D., 1996. Jumps and stochastic volatility: exchange rate processes implicit in deutsche mark options. Review of Financial Studies 9, 69-107.

Bates, D., 2000. Post-87 crash fears in S&P500 futures options. Journal of Econometrics 94, 181-238.

Bates, D., 2006. Maximum likelihood estimation of latent affine processes. Review of Financial Studies 19, 909-965.

Benzoni, L., 1998. Pricing options under stochastic volatility: an econometric analysis. Unpublished working paper, University of Minnesota.

Beveridge, S., Nelson, C., 1981. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle. Journal of Monetary Economics 7, 151-174.

Black, F., 1976. Studies of stock price volatility changes. In: Proceedings of the 1976 meetings of the business and economic statistics section, American Statistical Association, 177-181.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637-659.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307-327.

Bollerslev, T., Mikkelsen, H., 1996. Modeling and pricing long memory in stock market volatility. Journal of Econometrics 73, 151-184.

Bollerslev, T., Mikkelsen, H., 1999. Long-term equity anticipation securities and stock market volatility dynamics. Journal of Econometrics 92, 75-99.

Bollerslev, T., Zhou, H., 2002. Estimating stochastic volatility using conditional moments of integrated volatility. Journal of Econometrics 109, 33-65.

Brandt, M., Jones, C., 2006. Forecasting volatility with range-based EGARCH models. Journal of Business and Economic Statistics 24, 470-486.

Brennan, M., 1979. The pricing of contingent claims in discrete-time Models. Journal of Finance 34, 53-68.

Broadie, M., Chernov, M., Johannes, M., 2007. Model specification and risk premiums: evidence from futures options. Journal of Finance 62, 1453-1490.

Camara, A., 2003. A generalization of the Brennan-Rubinstein approach for the pricing of derivatives. Journal of Finance 58, 805-819.

Carr, P., Wu, L., 2004. Time-changed Levy processes and option pricing. Journal of Financial Economics 17, 113–141.

Chacko, G., Viceira, L., 2003. Spectral GMM estimation of continuous-time processes. Journal of Econometrics 116, 259-292.

Chernov, M., Ghysels, E., 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of option valuation. Journal of Financial Economics 56, 407-458.

Chernov, M., Gallant, R., Ghysels, E., Tauchen, G., 2003. Alternative models for stock price dynamics. Journal of Econometrics 116, 225-257.

Christoffersen, P., Heston, S., Jacobs, K., 2006. Option valuation with conditional skewness. Journal of Econometrics 131, 253-284.

Christoffersen, P., Jacobs, K., 2004. Which GARCH model for option valuation? Management Science 50, 1204-1221.

Christoffersen, P., Jacobs, K., Mimouni, K., 2005. An empirical comparison of affine and non-affine models for equity index options. Unpublished working paper, McGill University.

Comte, F., Coutin, L., Renault, E., 2001. Affine fractional stochastic volatility models. Unpublished working paper, University of Montreal.

Dai, Q., Singleton, K., 2000. Specification analysis of affine term structure models. Journal of Finance 55, 1943-1978.

Dai, Q., Le, A., and K. Singleton, 2006. Discrete-Time Dynamic Term Structure Models with Generalized Market Prices of Risk. Unpublished working paper, Stanford University.

Das, S., Sundaram, R., 1999. Of smiles and smirks: A term structure perspective. Journal of Financial and Quantitative Analysis 34, 211-239.

Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83-06.

Duan, J.-C., 1995. The GARCH option pricing model. Mathematical Finance 5, 13-32.

Duan, J.-C., Ritchken, P., Sun, Z., 2005. Jump starting GARCH: pricing and hedging options with jumps in returns and volatilities. Unpublished working paper, Rotman School, University of Toronto.

Duan, J.-C., Ritchken, P., Sun, Z., 2006. Approximating GARCH-jump models, jump-diffusion processes, and option pricing. Mathematical Finance 16, 21-52.

Duffee, G., 1999. Estimating the price of default risk. Review of Financial Studies 12, 197-226.

Duffie, D., Singleton, K., 1999. Modeling term structures of defaultable bonds. Review of Financial Studies 12, 687-720.

Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. Econometrica 68, 1343-1376.

Dumas, B., Fleming, J., Whaley, R., 1998. Implied volatility functions: empirical tests. Journal of Finance 53, 2059-2106.

Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. Econometrica 50, 987-1008.

Engle, R., Mustafa, C., 1992. Implied ARCH models from options prices. Journal of Econometrics 52, 289-311.

Engle, R., Lee, G., 1999. A permanent and transitory component model of stock return volatility. In: Engle, R., White, H. (Eds.), Cointegration, causality, and forecasting: a festschrift in honor of Clive W.J. Granger. Oxford University Press, New York, pp. 475-497.

Engle, R., Ng, V., 1993. Measuring and testing the impact of news on volatility. Journal of Finance 48, 1749-1778.

Engle, R., Patton, A., 2001. What good is a volatility model? Quantitative Finance 1, 237-245.

Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance 59, 1367-1403.

Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. Journal of Finance 58, 1269-1300.

Fama, E., French, K., 1988. Permanent and temporary components of stock prices. Journal of Political Economy 96, 246-273.

French, K., Schwert, G.W., Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3-30.

Hentschel, L., 1995. All in the family: nesting symmetric and asymmetric GARCH models. Journal of Financial Economics 39, 71-104.

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., 2000. A closed-form GARCH option pricing model. Review of Financial Studies 13, 585-626.

Hsieh, K., Ritchken, P., 2005. An empirical comparison of GARCH option pricing models. Review of Derivatives Research 8, 129-150.

Huang, J.-Z., Wu, L., 2004. Specification analysis of option pricing models based on time-changed Levy processes. Journal of Finance 59, 1405–1439.

Hull, J., White, A., 1987. The pricing of options with stochastic volatilities. Journal of Finance 42, 281-300.

Jones, C., 2003. The dynamics of stochastic volatility: evidence from underlying and options markets. Journal of Econometrics 116, 181-224.

Maheu, J., 2005, Can GARCH models capture long-range dependence? Studies in Nonlinear Dynamics and Econometrics 9, Article 1.

Maheu, J., McCurdy, T., 2004. News arrival, jump dynamics and volatility components for individual stock returns. Journal of Finance 59, 755–793.

Melino, A., Turnbull, S., 1990. Pricing foreign currency options with stochastic volatility. Journal of Econometrics 45, 239-265.

Nandi, S., 1998. How important is the correlation between returns and volatility in a stochastic volatility model? Empirical evidence from pricing and hedging in the S&P 500 index options market. Journal of Banking and Finance 22, 589-610.

Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study. Journal of Financial Economics 63, 3–50.

Pearson, N., Sun, T., 1994. Exploiting the conditional density in estimating the term structure: an application to the Cox, Ingersoll, and Ross model. Journal of Finance 49, 1279-1304.

Poterba, J., Summers, L., 1988. Mean reversion in stock returns: evidence and implications. Journal of Financial Economics 22, 27-60.

Ritchken, P., Trevor, R., 1999. Pricing options under generalized GARCH and stochastic volatility processes. Journal of Finance 54, 377-402.

Rubinstein, M., 1976. The valuation of uncertain income streams and the pricing of options. Bell Journal of Economics 7, 407-425.

Schroder, M., 2004. Risk-neutral parameter shifts and derivative pricing in discrete time. Journal of Finance 59, 2375-2401.

Scott, L., 1987. Option pricing when the variance changes randomly: theory, estimators and applications. Journal of Financial and Quantitative Analysis 22, 419-438.

Summers, L., 1986. Does the stock market rationally reflect fundamental values? Journal of Finance 41, 591-600.

Taylor, S., Xu, X., 1994. The term structure of volatility implied by foreign exchange options. Journal of Financial and Quantitative Analysis 29, 57-74.

Wiggins, J., 1987. Option values under stochastic volatility: theory and empirical evidence. Journal of Financial Economics 19, 351-372.

Figure 1. Term Structure of Variance with High and Low Initial Variance.
Component Model and GARCH(1,1). Normalized by Unconditional Variance.
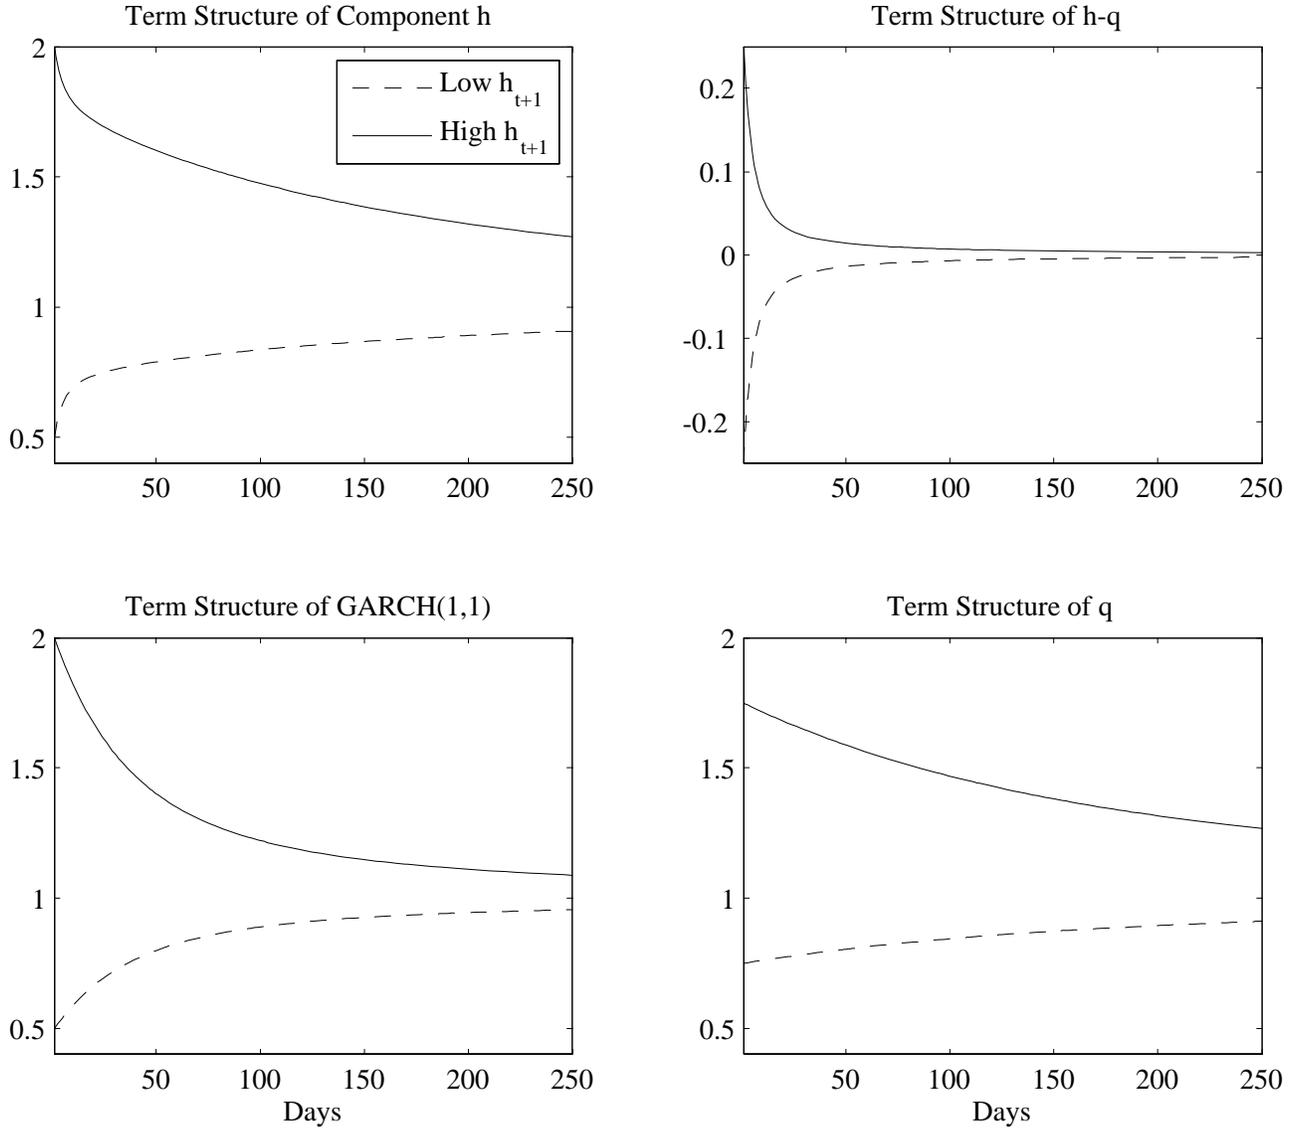
Notes to Figure: In the left-side panels we plot the variance term structure implied by the component GARCH and GARCH$(1, 1)$ models for 1 through 250 days. In the right-side panels we plot the term structure of the individual components for the component model. The parameter values are obtained from MLE estimation on returns in Table 1. For the low initial variance (dashed lines), the initial value of q_{t+1} is set to $0.75\sigma^2$, the initial value of h_{t+1} is set to $0.5\sigma^2$, and the initial value for h_{t+1} in the GARCH$(1, 1)$ is set to $0.5\sigma^2$. For the high initial variance (solid lines), the initial value of q_{t+1} is set to $1.75\sigma^2$, the initial value of h_{t+1} is set to $2\sigma^2$, and the initial value for h_{t+1} in the GARCH$(1, 1)$ is set to $2\sigma^2$. All values are normalized by the unconditional variance σ^2.

Figure 2. Variance Term Structure Impulse Response to Positive and Negative Return Shocks. Component Model Versus GARCH(1,1). Normalized by Unconditional Variance.
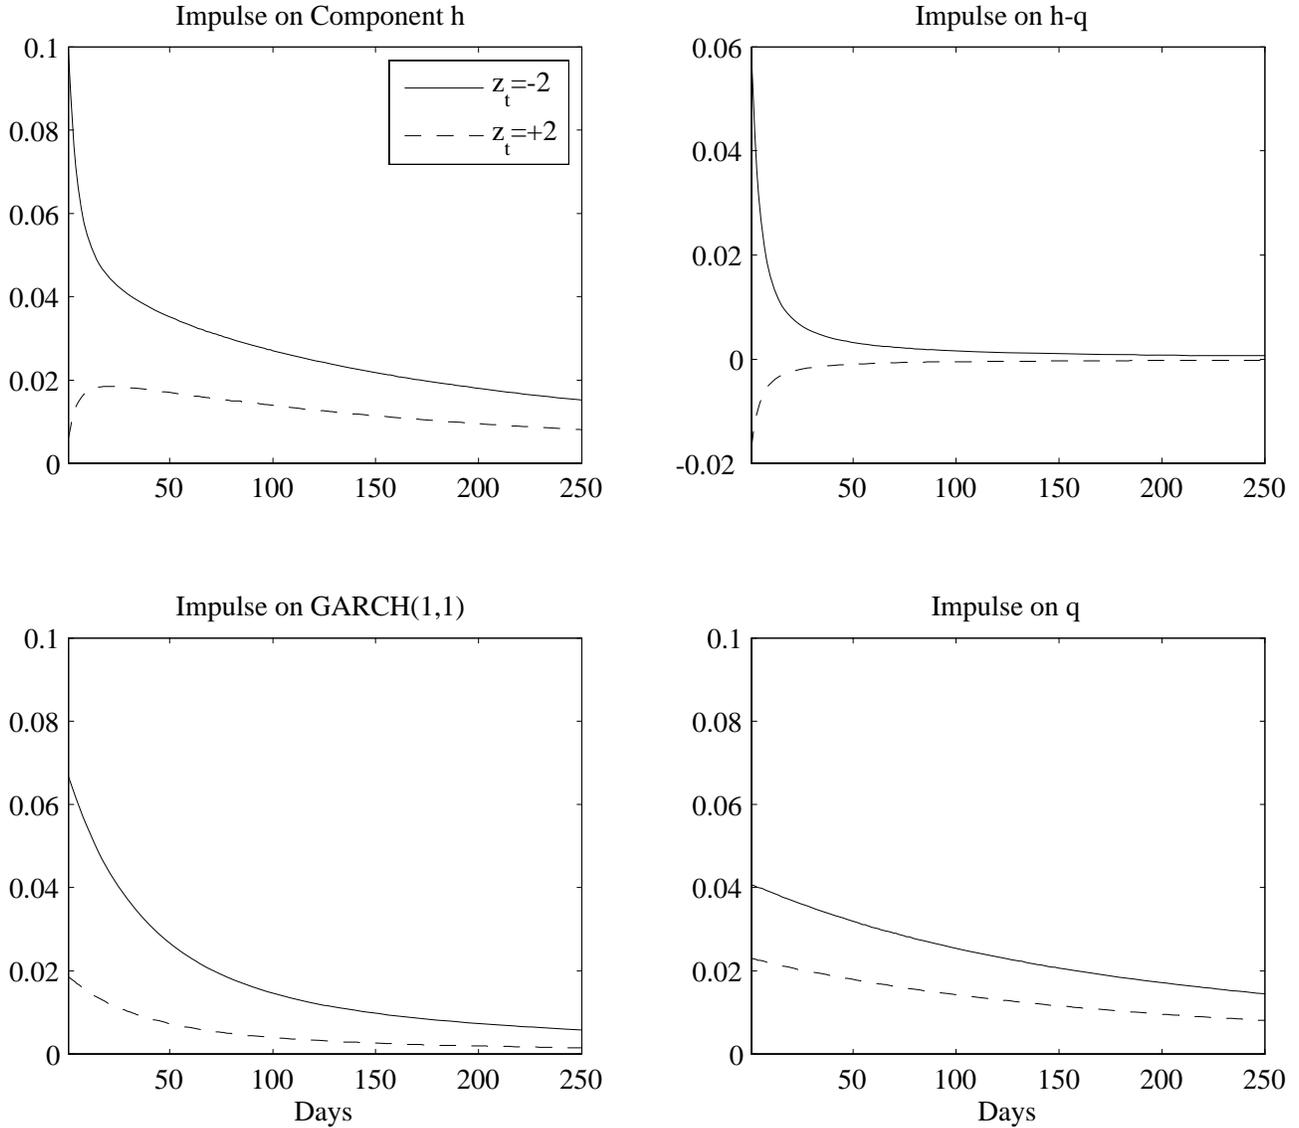
Notes to Figure: We plot the variance term structure responses to two shocks: $z_t = 2$ (dashed lines) and $z_t = -2$ (solid lines). In the left-hand panels we plot the responses to shocks to returns in the component and GARCH(1, 1) models. For the component model, the right-hand panels show the response of the individual components. The parameter values are obtained from the MLE estimation on returns in Table 1. The current variance is set equal to the unconditional value. All values are normalized by the unconditional variance.

Figure 3. Term Structure of Skewness and Kurtosis.
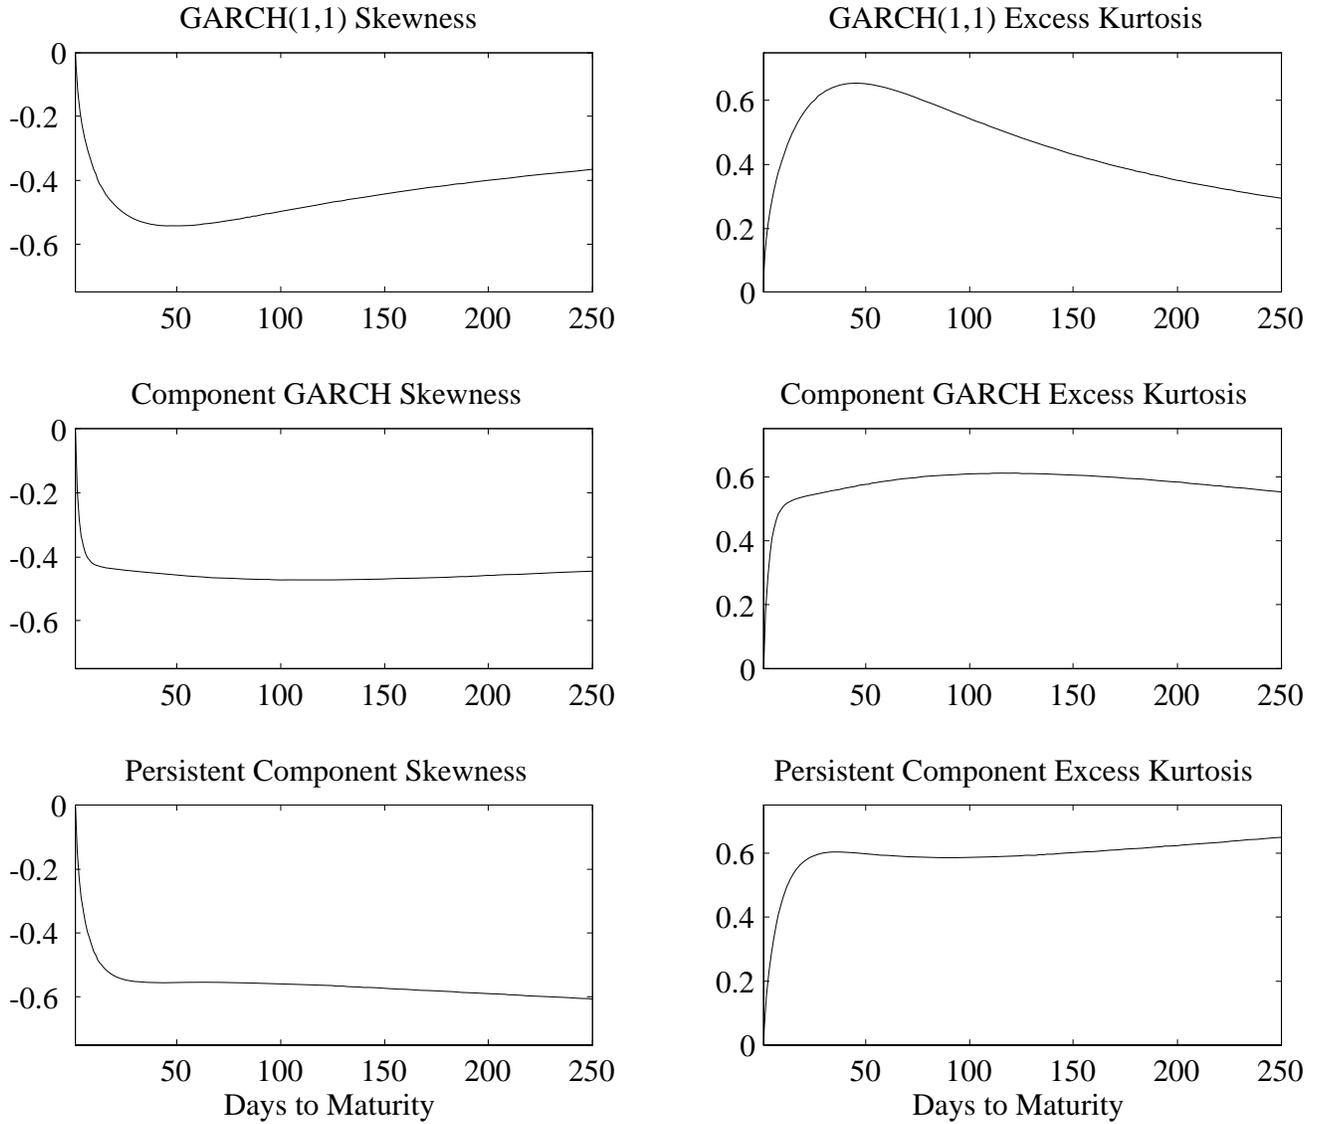
Notes to Figure: We use the numerical derivatives of $\ln(f(t, T; \phi))$ in (4.1) to compute the term structure of conditional skewness and kurtosis in the three GARCH models. The initial volatility is set to its long run value in the GARCH$(1, 1)$ and component GARCH models. In the persistent component model the initial volatility is set to the unconditional volatility from the component model. The parameter values are obtained from the MLE estimates on returns in Table 1.

Figure 4. Spot Variance of Component GARCH versus GARCH(1,1).
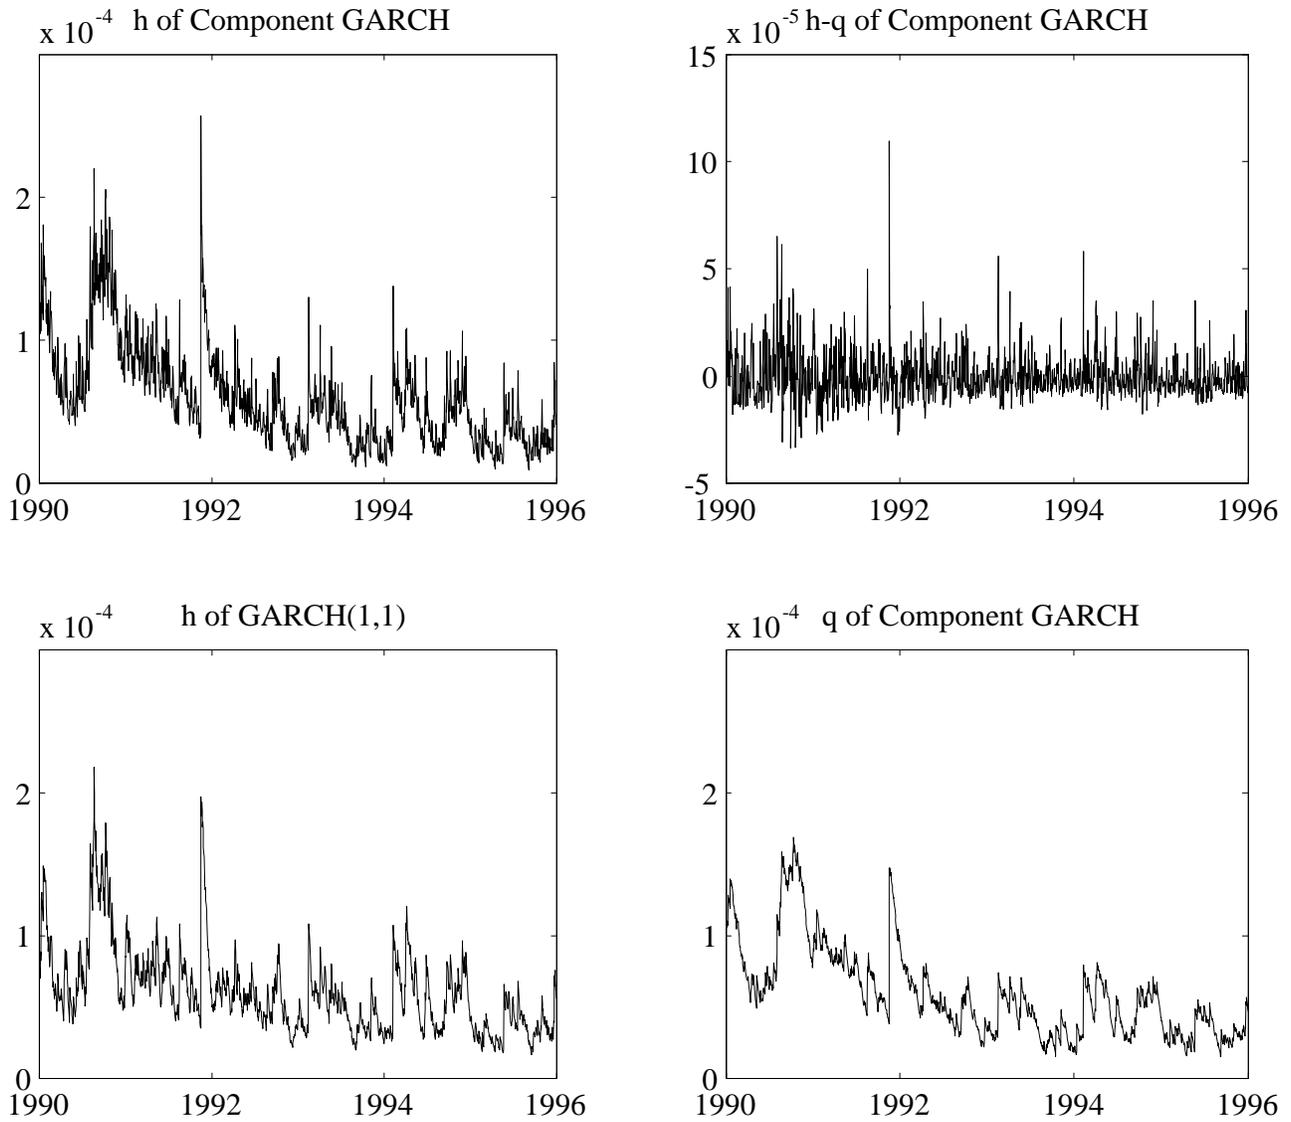
Notes to Figure: The left-hand panels plot the variance paths from the component and GARCH(1, 1) models. The right-hand panels plot the individual components. The parameter values are obtained from MLE estimation on returns in Table 1.

Figure 5. Conditional Standard Deviation of Variance and Conditional Correlation.
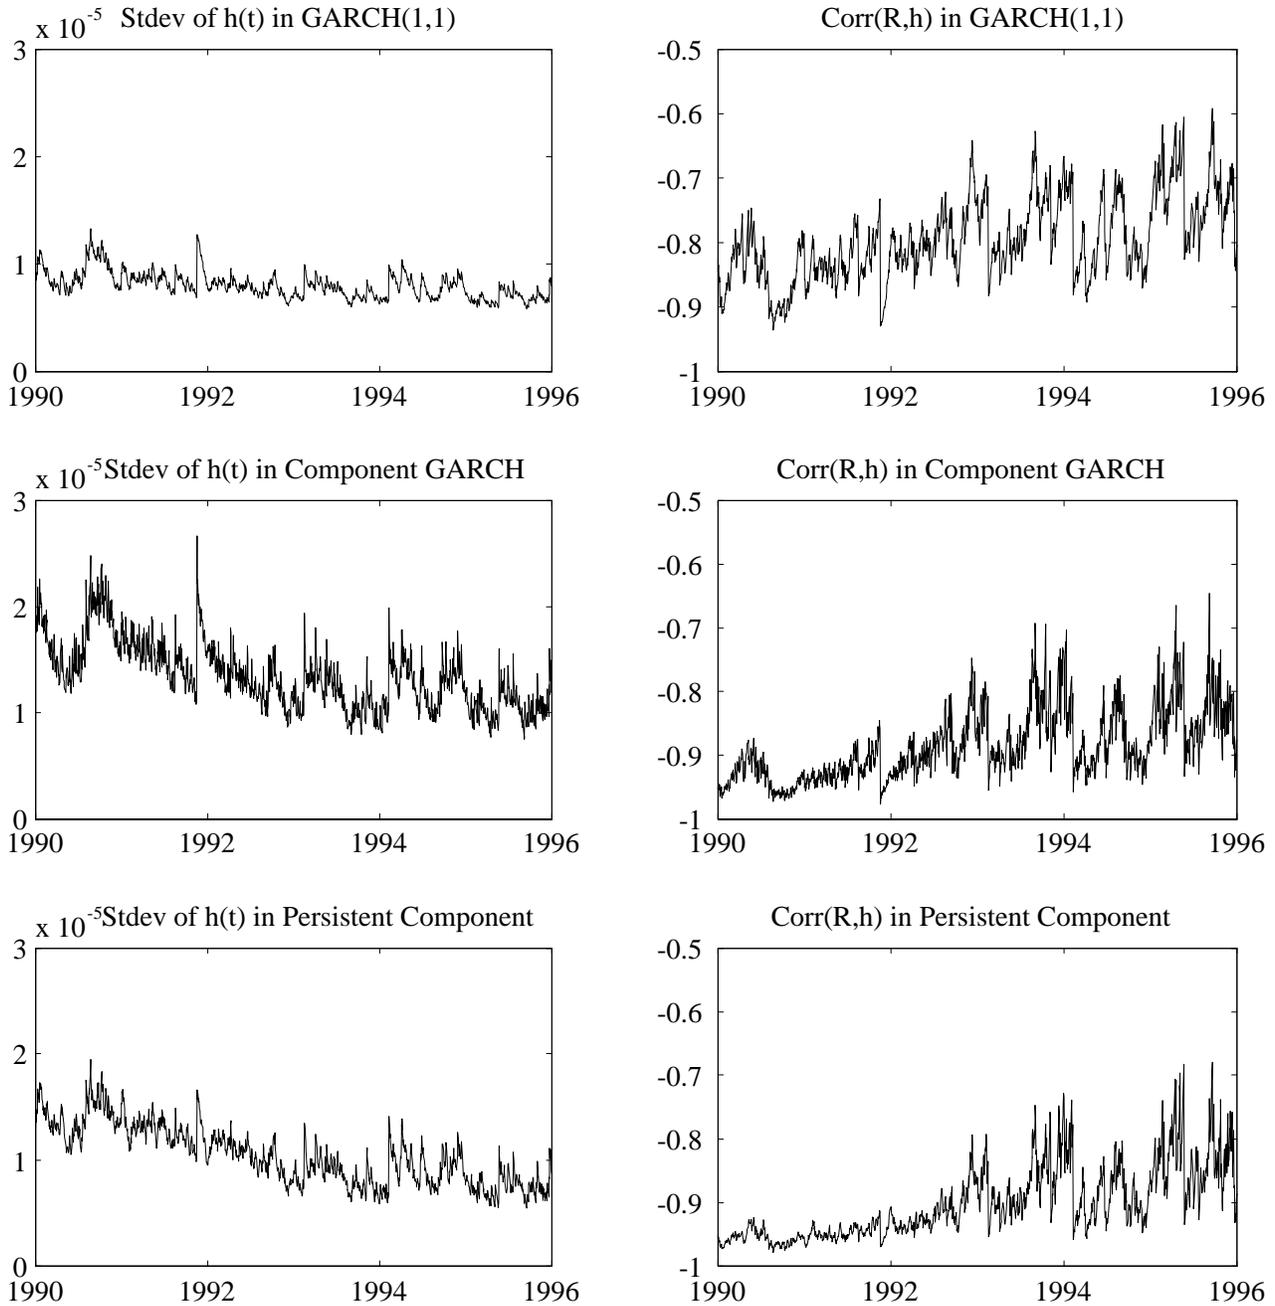
Notes to Figure: In the left panels we plot the conditional standard deviation of conditional variance as implied by the GARCH models and in the right panels we plot the conditional correlation between return and variance. The scales are identical across the rows of panels in order to facilitate comparison across models. The parameter values are obtained from the MLE estimates on returns in Table 1.

Figure 6. Correlation Between Return and Variance as a Function of Volatility Level.
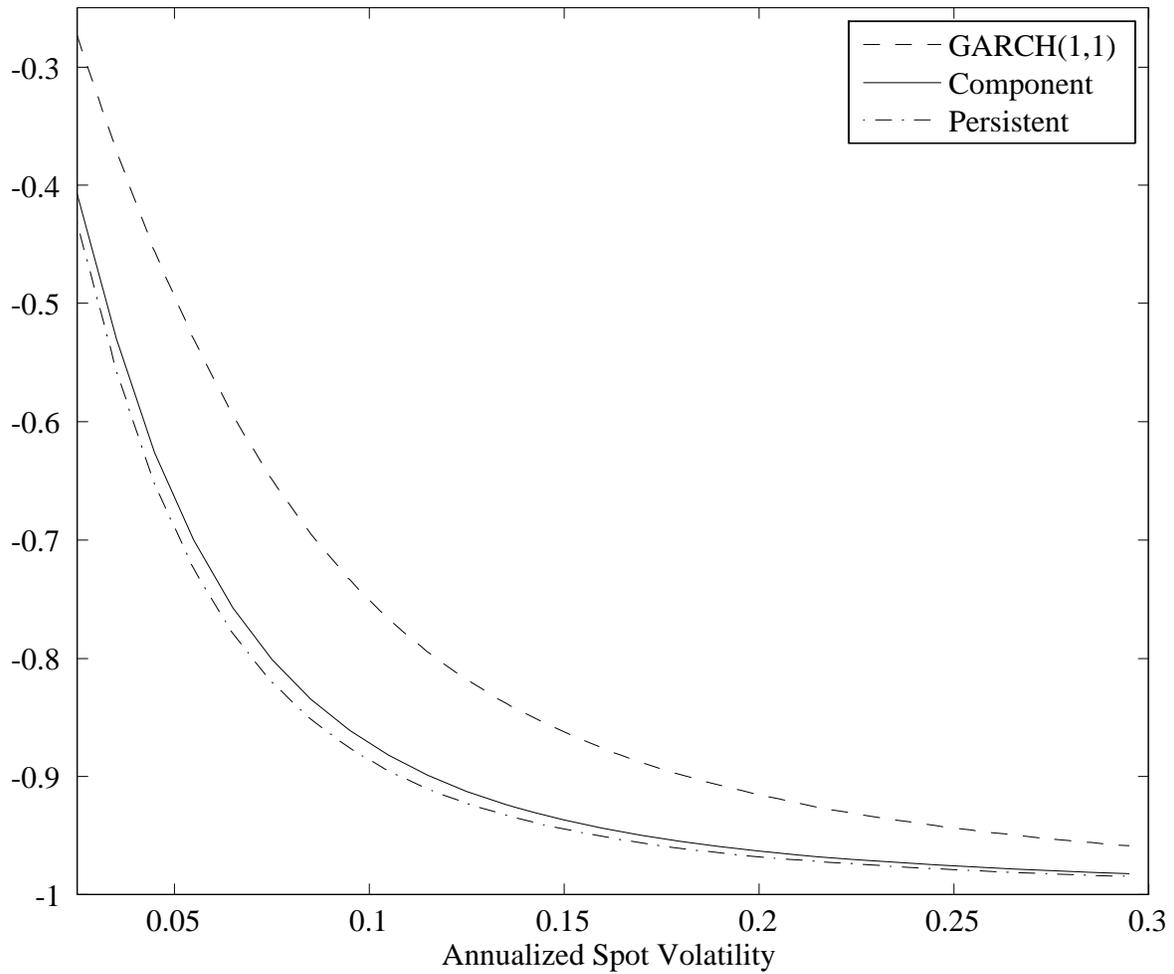
Notes to Figure: We plot the conditional correlation between the return on the underlying index and the daily variance. This conditional correlation is plotted against the level of volatility annualized. The dashed line corresponds to the GARCH$(1, 1)$, the solid line to the component model and the dash-dots to the persistent component model. The parameter estimates are from Table 1.

Figure 7. Weekly Implied Volatility Bias for At-the-Money Options.
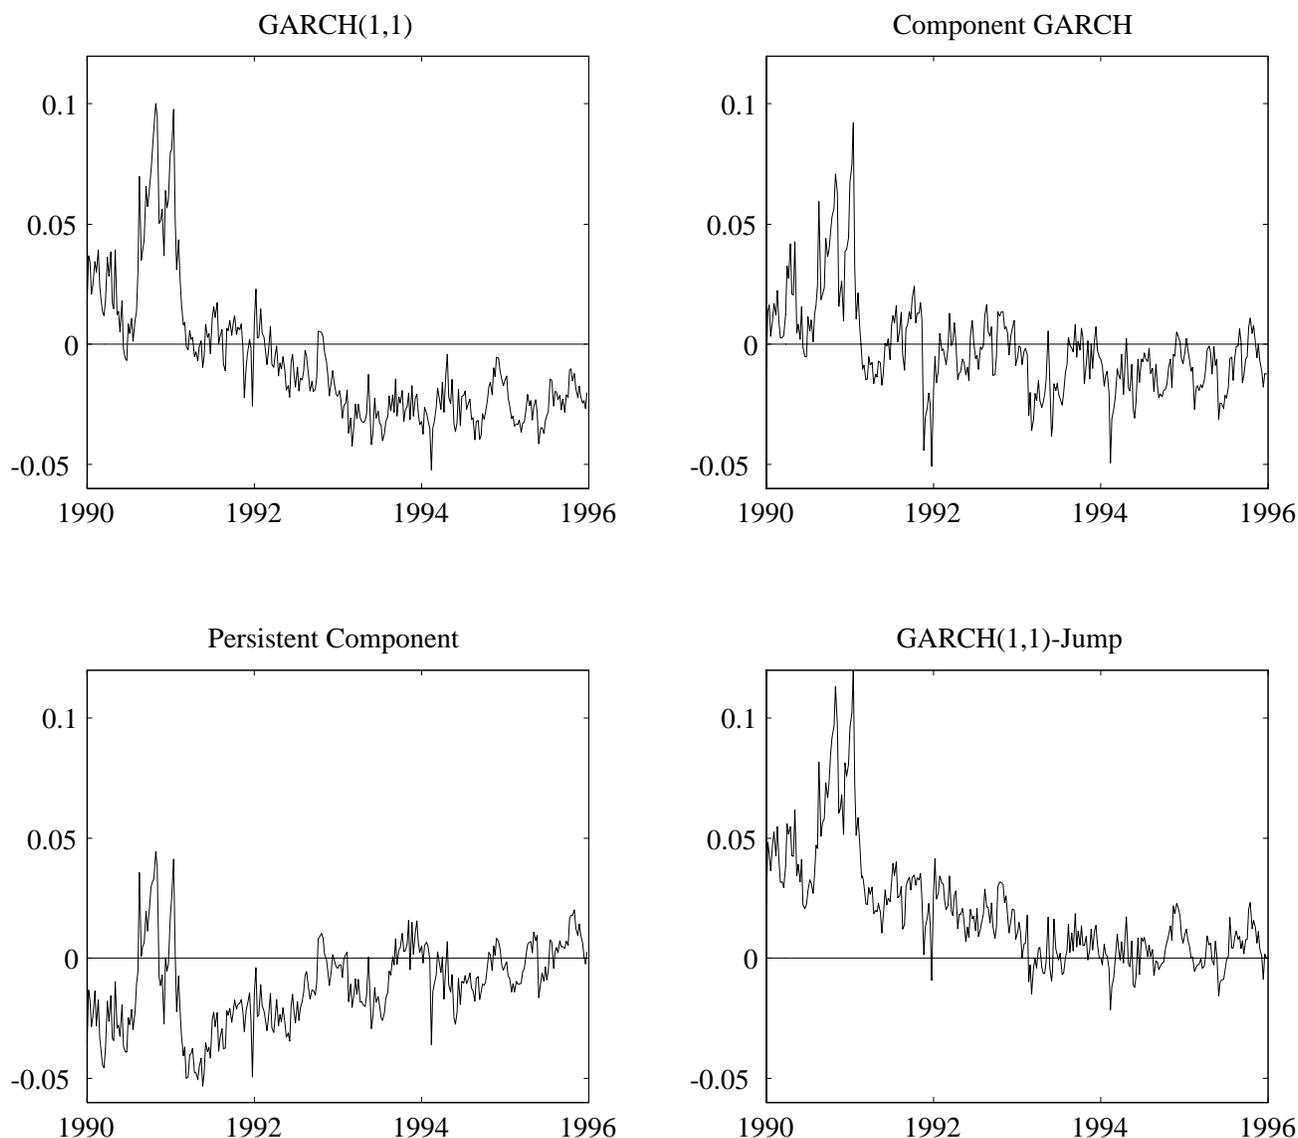
Notes to Figure: Each Wednesday we compute the Black-Scholes implied volatility for each at-the-money option contract. Options with moneyness (index value over strike price) between 0.975 and 1.025 are considered at-the-money. The implied volatility is computed both for the market price and for each model price. We plot the weekly average difference between the market and model implied volatility. The top-left panel shows the GARCH(1, 1) model, the top-right panel shows the component model, the bottom-left panel shows the persistent component model, and the bottom-right panel shows the GARCH(1, 1)-Jump model. The MLE estimates from Tables 1 and 4 are used.

Figure 8. Implied Volatility Term Structures for At-the-Money Options.
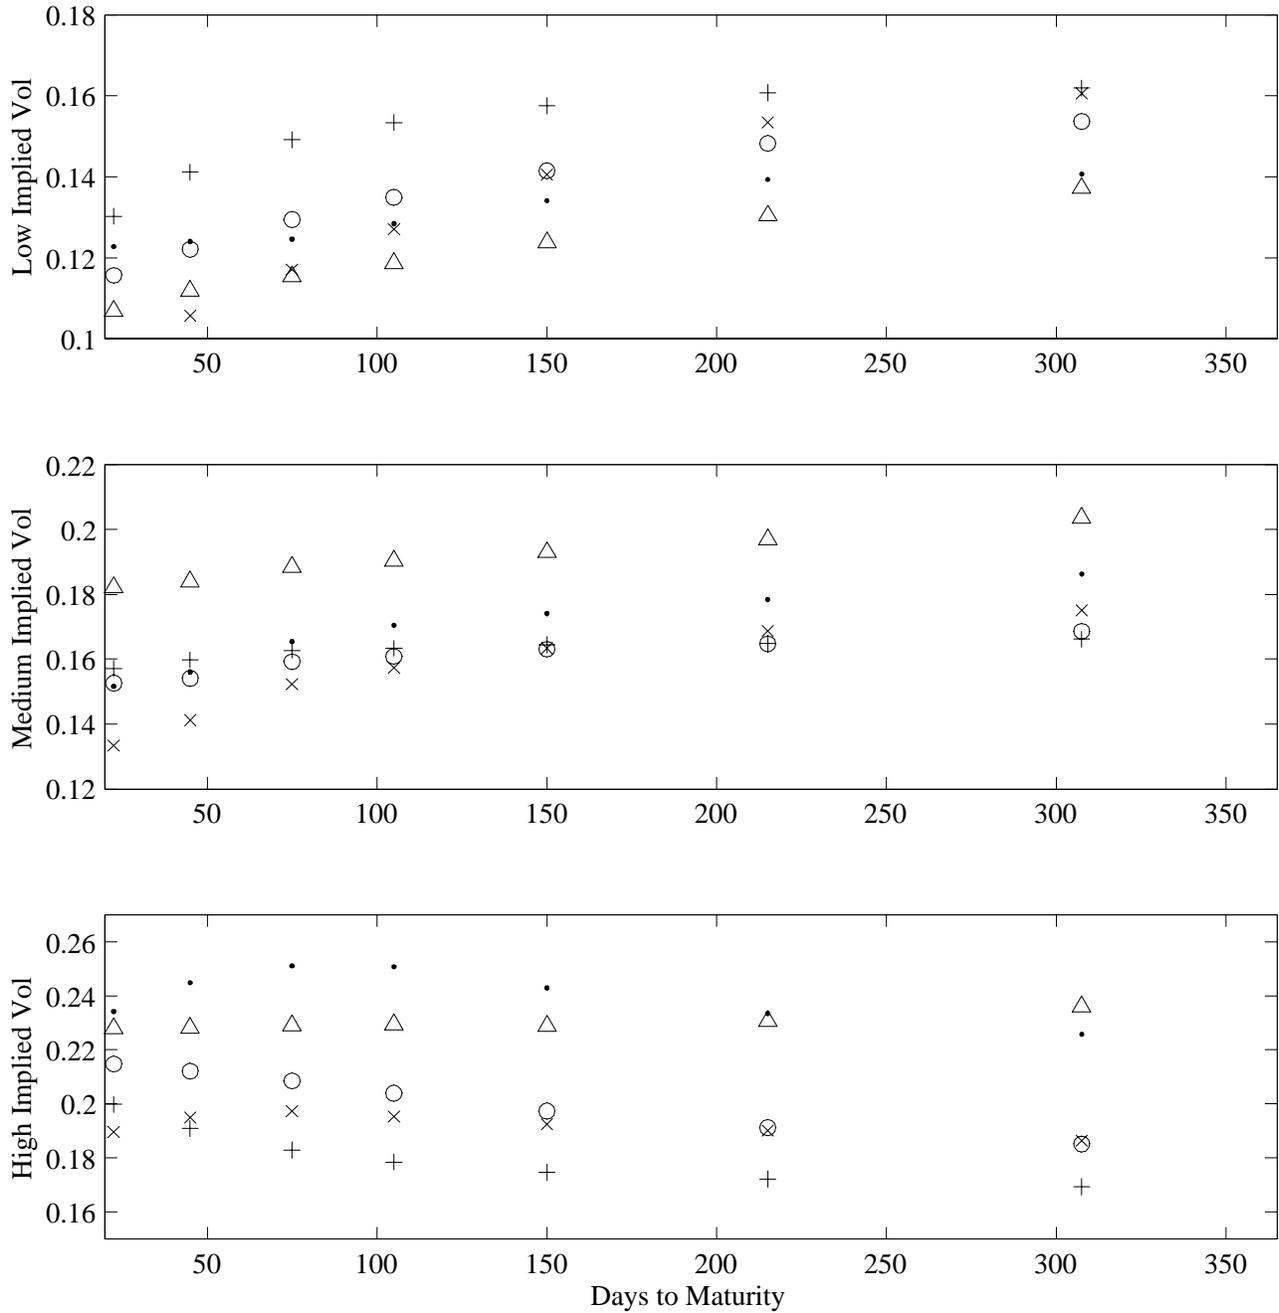
Notes to Figure: We compute average implied annualized Black-Scholes volatilities based on the data and the models for three different periods: a high volatility period (bottom panel), a medium volatility period (middle panel) and a low volatility period (top panel). We report on the three GARCH models and the GARCH-Jump model for options with different maturity. A dot represents the implied volatility of the data, a + represents the GARCH(1,1) model, a circle represents the component GARCH model, a triangle represents the persistent component model, and an x represents the GARCH(1,1)-Jump model. The maturity is on the horizontal axis. The MLE estimates from Tables 1 and 4 are used.

Figure 9. Implied Volatility Smirks.
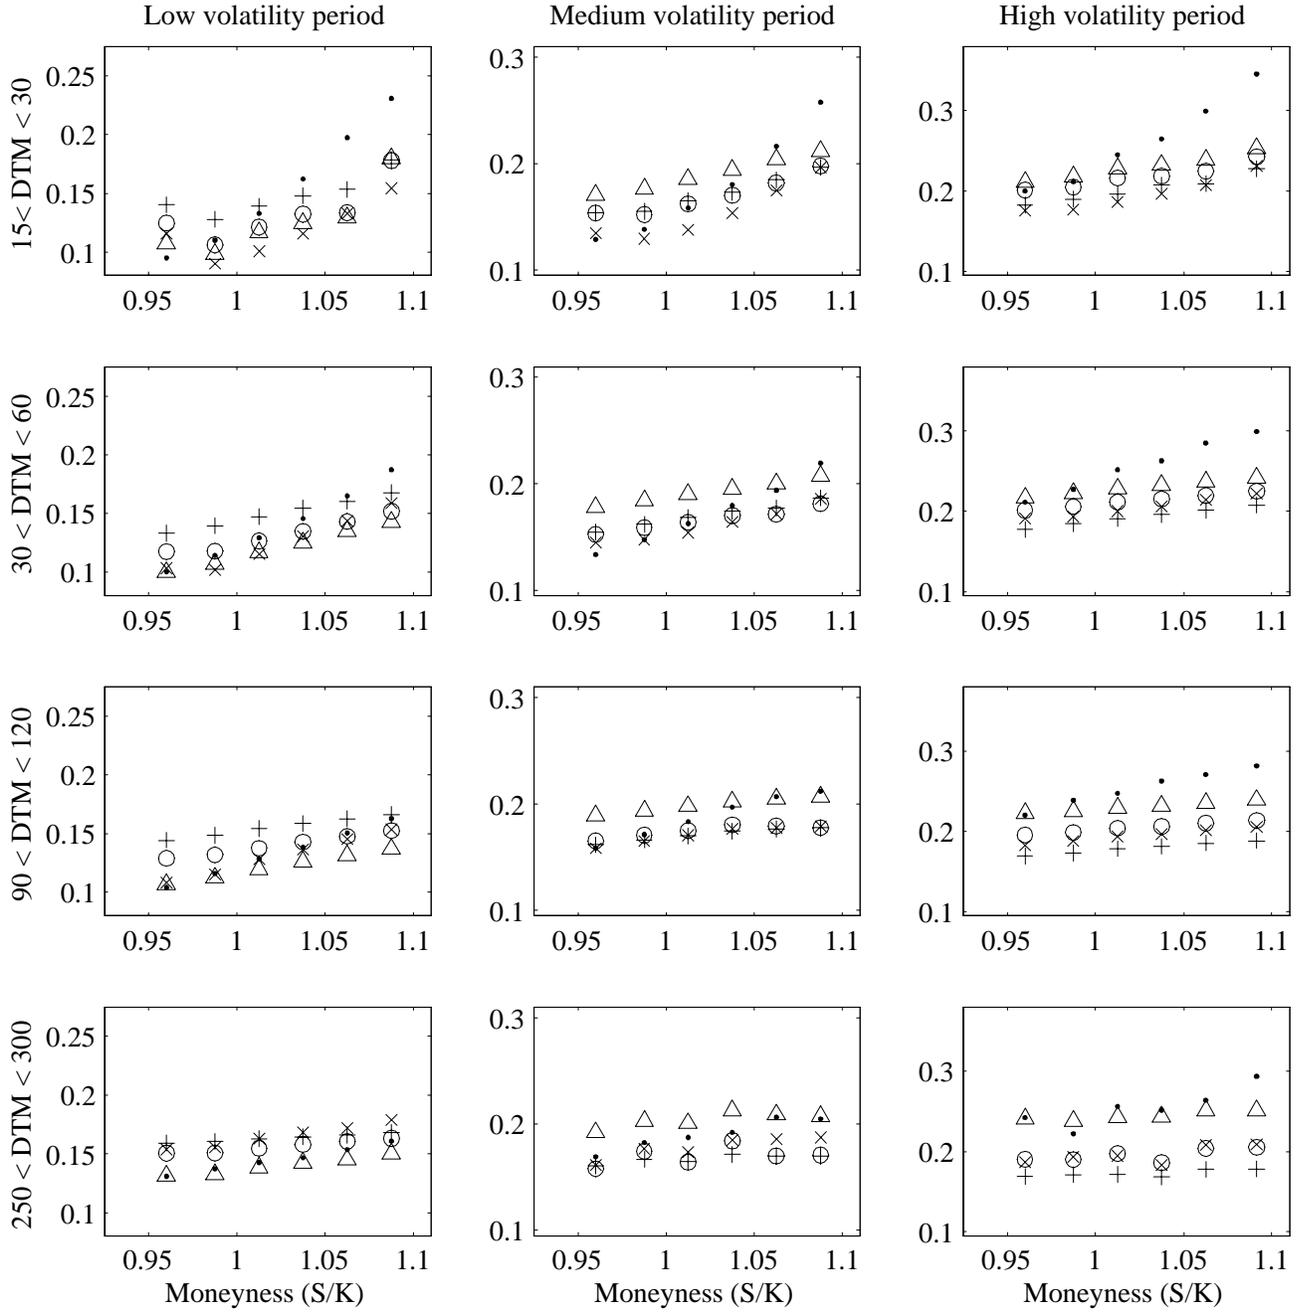
Notes to Figure: We compute average implied annualized Black-Scholes volatilities based on the data and the models for three different periods: a high volatility period, a medium volatility period and a low volatility period. We report on the three normal GARCH models and the GARCH(1,1)-Jump model for various moneyness and maturity. A dot represents the implied volatility of the data, a + represents the GARCH(1,1) model, a circle represents the component GARCH model, a triangle represents the persistent component model, and an x represents the GARCH(1,1)-Jump model. The moneyness is on the horizontal axis and each row of panels corresponds to a different maturity. The MLE estimates from Tables 1 and 4 are used.

Table 1. MLE Estimates and Properties
Estimation Sample: Daily Returns, 1962-2001

GARCH(1,1)			Component GARCH			Persistent Component		
Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error
λ	2.231E+00	1.123E+00	λ	2.092E+00	7.729E-01	λ	2.017E-07	4.316E-01
w	2.101E-17	1.120E-07	$\tilde{\beta}$	6.437E-01	2.759E-02	$\tilde{\beta}$	8.822E-01	9.931E-03
b	9.012E-01	4.678E-03	α	1.580E-06	2.430E-07	α	2.057E-06	1.539E-07
a	3.317E-06	1.380E-07	γ_1	4.151E+02	6.341E+01	γ_1	2.516E+02	2.237E+01
c	1.276E+02	8.347E+00	γ_2	6.324E+01	5.300E+00	γ_2	1.187E+02	1.126E+01
			ω	8.208E-07	7.620E-08	ω	1.187E-07	1.393E-08
			φ	2.480E-06	1.160E-07	φ	7.966E-07	4.599E-08
			ρ	9.896E-01	9.630E-04	ρ	1.000E+00	

Ln Likelihood	33,955		Ln Likelihood	34,102		Ln Likelihood	34,005	
Persistence	0.9553		Persistence	0.9963		Persistence	1.0000	
Average Annual Vol	0.1206		Average Annual Vol	0.1174		Average Annual Vol	0.1239	
Average Vol of Var	7.997E-06		Average Vol of Var	1.341E-05		Average Vol of Var	1.044E-05	
Average Correlation	-0.7940		Average Correlation	-0.8849		Average Correlation	-0.9014	
Option RMSE	2.236		Option RMSE	1.706		Option RMSE	1.705	
Normalized	1.000		Normalized	0.763		Normalized	0.763	

Notes to Table: We use daily total returns from July 1, 1962 to December 31, 2001 on the S&P500 index to estimate the three models using Maximum Likelihood. Robust standard errors are calculated from the outer product of the gradient at the optimum parameter values. Persistence refers to the persistence of the conditional variance as defined in the text. Average Annual Vol refers to the average annualized standard deviation during 1990-95. Average Vol of Var refers to the average standard deviation of the conditional variance during 1990-95. Average Correlation refers to the average correlation between the return and the conditional variance during 1990-95. Ln Likelihood refers to the logarithm of the likelihood at the optimal parameter values. Option RMSE refers to the fit of the models on option prices observed during 1990-95.

Table 2. S&P 500 Index Call Option Data (1990-1995)

Panel A. Number of Call Option Contracts

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	101	1,884	1,931	1,769	5,685
0.975<S/X<1.00	283	1,272	706	477	2,738
1.00<S/X<1.025	300	1,212	726	526	2,764
1.025<S/X<1.05	261	1,167	654	409	2,491
1.05<S/X<1.075	245	1,039	582	390	2,256
1.075<S/X	549	2,345	1,679	1,245	5,818
All	1,739	8,919	6,278	4,816	21,752

Panel B. Average Call Price

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.88	2.30	6.25	11.94	6.62
0.975<S/X<1.00	2.29	6.83	15.19	27.50	12.12
1.00<S/X<1.025	8.35	13.60	22.48	34.41	19.32
1.025<S/X<1.05	17.57	22.00	30.11	42.14	26.97
1.05<S/X<1.075	27.11	30.84	38.14	48.83	35.43
1.075<S/X	50.67	52.79	58.99	68.34	57.70
All	24.32	23.66	28.68	36.07	27.91

Panel C. Average Implied Volatility from Call Options

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.1625	0.1269	0.1350	0.1394	0.1342
0.975<S/X<1.00	0.1308	0.1296	0.1449	0.1562	0.1383
1.00<S/X<1.025	0.1527	0.1459	0.1558	0.1606	0.1520
1.025<S/X<1.05	0.1915	0.1647	0.1665	0.1656	0.1681
1.05<S/X<1.075	0.2433	0.1828	0.1775	0.1739	0.1865
1.075<S/X	0.3897	0.2356	0.1961	0.1868	0.2283
All	0.2434	0.1703	0.1622	0.1607	0.1717

Notes to Table: We use European call options on the S&P500 index. The prices are taken from quotes within 30 minutes from closing on each Wednesday during the January 1, 1990 to December 31, 1995 period. The moneyness and maturity filters used by Bakshi, Cao and Chen (1997) are applied here as well. The implied volatilities are calculated using the Black-Scholes formula.

Table 3. 1990-1995 RMSE and Ratio RMSE by Moneyness and Maturity
Parameters Estimated from Daily Returns, 1962-2001

Panel A. GARCH(1,1) RMSE

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.454	1.778	3.032	4.155	3.090
0.975<S/X<1.00	0.671	2.116	3.087	3.548	2.603
1.00<S/X<1.025	0.638	1.650	2.574	2.955	2.154
1.025<S/X<1.05	0.595	1.204	2.099	2.487	1.700
1.05<S/X<1.075	0.735	1.013	1.879	2.227	1.516
1.075<S/X	0.759	1.024	1.424	1.917	1.360
All	0.683	1.503	2.448	3.228	2.236

Panel B. Ratio of Component GARCH to GARCH(1,1) RMSE

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.782	0.684	0.712	0.782	0.749
0.975<S/X<1.00	0.788	0.631	0.657	0.739	0.678
1.00<S/X<1.025	0.870	0.655	0.669	0.733	0.691
1.025<S/X<1.05	0.968	0.832	0.744	0.755	0.773
1.05<S/X<1.075	1.043	1.000	0.849	0.800	0.870
1.075<S/X	1.000	1.037	0.974	0.907	0.962
All	0.949	0.750	0.735	0.784	0.763

Panel C. Ratio of Persistent Component to GARCH(1,1) RMSE

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	0.985	0.726	0.774	0.842	0.808
0.975<S/X<1.00	0.978	0.640	0.635	0.722	0.669
1.00<S/X<1.025	0.982	0.672	0.600	0.681	0.653
1.025<S/X<1.05	0.947	0.773	0.613	0.660	0.675
1.05<S/X<1.075	1.010	0.909	0.663	0.718	0.750
1.075<S/X	1.002	0.981	0.836	0.784	0.856
All	0.990	0.746	0.719	0.796	0.763

Notes to Table: We use the MLE estimates from Table 1 to compute the root mean squared option valuation error (RMSE) for various moneyness and maturity bins during 1990-1995. Panel A shows the RMSEs for the GARCH(1,1) model. Panels B and C show the ratio of the RMSEs of the component and persistent component models over the RMSE of the GARCH(1,1) model.

Table 4. GARCH(1,1)-Jump Model
Parameters Estimated from Daily Returns, 1962-2001

Panel A. GARCH(1,1)-Jump MLE Estimates

Parameter	Estimate	Std. Error
λ_z	2.793E+00	1.864E+00
λ_y	-3.419E-03	1.153E-02
w	-1.186E-06	9.457E-08
b	9.551E-01	3.979E-03
a	2.191E-06	1.306E-07
c	1.088E+02	8.983E+00
χ	9.845E-03	2.100E-03
μ	-1.174E-02	4.879E-03
τ	2.799E-02	2.048E-03

Ln Likelihood	34,260
Option RMSE	1.939

Panel B. Ratio of GARCH(1,1)-Jump to GARCH(1,1) RMSE

	DTM<20	20<DTM<80	80<DTM<180	DTM>180	All
S/X<0.975	1.265	0.538	0.654	0.631	0.629
0.975<S/X<1.00	1.166	0.655	0.758	0.809	0.749
1.00<S/X<1.025	1.578	1.027	0.989	0.908	0.978
1.025<S/X<1.05	1.405	1.467	1.231	1.016	1.224
1.05<S/X<1.075	0.989	1.556	1.377	1.201	1.346
1.075<S/X	1.177	1.277	1.348	1.223	1.274
All	1.247	0.942	0.902	0.800	0.867

Notes to Table: In Panel A, we use daily total returns from July 1, 1962 to December 31, 2001 on the S&P500 index to estimate the GARCH(1,1)-Jump model using Maximum Likelihood. Robust standard errors are calculated from the outer product of the gradient at the optimum parameter values. In Panel B, we compute the ratio of the option root mean squared error (RMSE) from the GARCH(1,1)-Jump model to the RMSE of the GARCH(1,1) in Panel A of Table 3.

Table 5. NLS Estimates and Properties
Sample: 1990-1992 (in-sample) 1993 (out-of-sample)

GARCH(1,1)			Component GARCH			Persistent Component		
Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error
w	3.891E-14	3.560E-12	$\tilde{\beta}$	7.050E-01	2.565E-01	$\tilde{\beta}$	7.201E-01	1.021E-01
b	6.801E-01	3.211E-03	α	1.770E-06	3.444E-07	α	1.597E-06	2.279E-06
a	2.666E-07	6.110E-09	γ_1	5.617E+02	1.494E+02	γ_1	7.481E+02	8.974E+01
c	1.090E+03	5.432E+01	γ_2	5.638E+02	1.555E+02	γ_2	4.767E+02	1.246E+02
			ω	2.424E-07	1.212E-07	ω	5.343E-08	1.345E-08
			φ	5.249E-07	3.525E-07	φ	5.123E-07	1.285E-07
			ρ	9.981E-01	3.519E-03	ρ	1.000E+00	

Persistence	0.9970		Persistence	0.9994		Persistence	1.0000	
Average Annual Vol	0.1347		Average Annual Vol	0.1405		Average Annual Vol	0.1431	
Average Vol of Var	4.283E-06		Average Vol of Var	1.962E-05		Average Vol of Var	2.197E-05	
Average Correlation	-0.9967		Average Correlation	-0.9876		Average Correlation	-0.9914	

	NLS	MLE		NLS	MLE		NLS	MLE
RMSE (90-92)	1.038	1.896	RMSE (90-92)	0.931	1.609	RMSE (90-92)	0.984	2.143
Normalized	1.000	1.000	Normalized	0.897	0.849	Normalized	0.948	1.130
RMSE (93)	1.284	2.229	RMSE (93)	0.983	1.584	RMSE (93)	1.198	1.260
Normalized	1.000	1.000	Normalized	0.765	0.710	Normalized	0.933	0.565

Notes to Table: We use Wednesday option prices from from January 1, 1990 to December 31, 1992 on the S&P500 index to estimate the three GARCH models using Nonlinear Least Squares on the valuation errors. Robust standard errors are calculated from the outer product of the gradient at the optimum parameter values. RMSE refers to the square root of the mean-squared valuation errors. RMSE(90-92) refers to the 1990-1992 in-sample period, and RMSE(93) refers to the 1993 out-of-sample period. NLS refers to the model estimated using option data and MLE refers to the model estimated using returns only. Normalized values are divided by the respective RMSE from GARCH(1,1).

Table 6. NLS Estimates and Properties

Sample: 1992-1994 (in-sample) 1995 (out-of-sample)

GARCH(1,1)			Component GARCH			Persistent Component		
Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error	Parameter	Estimate	Std. Error
w	7.521E-16	3.498E-09	$\tilde{\beta}$	9.297E-01	3.346E-02	$\tilde{\beta}$	9.587E-01	3.821E-05
b	4.694E-01	1.251E-01	α	1.808E-06	1.320E-07	α	1.943E-06	1.614E-06
a	1.936E-06	3.986E-07	γ_1	5.854E+02	2.362E+02	γ_1	2.589E+02	8.383E+01
c	5.061E+02	1.041E+02	γ_2	5.749E+02	4.025E+02	γ_2	2.254E+02	5.063E+02
			ω	2.204E-07	3.470E-08	ω	6.927E-08	1.262E-08
			φ	2.835E-07	1.586E-07	φ	6.971E-07	1.253E-08
			ρ	9.966E-01	1.277E-03	ρ	1.000E+00	

Persistence	0.9654		Persistence	0.9998		Persistence	1.0000	
Average Annual Vol	0.1074		Average Annual Vol	0.1129		Average Annual Vol	0.1082	
Average Vol of Var	1.423E-05		Average Vol of Var	1.838E-05		Average Vol of Var	1.085E-05	
Average Correlation	-0.9701		Average Correlation	-0.9781		Average Correlation	-0.9095	

	NLS	MLE		NLS	MLE		NLS	MLE
RMSE (92-94)	1.107	2.000	RMSE (92-94)	0.855	1.524	RMSE (92-94)	1.058	1.402
Normalized	1.000	1.000	Normalized	0.773	0.762	Normalized	0.956	0.701
RMSE (95)	1.227	2.775	RMSE (95)	0.972	1.920	RMSE (95)	1.174	1.249
Normalized	1.000	1.000	Normalized	0.792	0.692	Normalized	0.957	0.450

Notes to Table: See notes to Table 5. RMSE(92-94) is now in-sample and RMSE(95) is out-of-sample.

Table 7. RMSE and Ratio RMSE by Moneyness and Maturity

	Panel A. 1990-1992 in Sample			Panel B. 1993 Out of Sample		
	GARCH(1,1)	Component	Persistent	GARCH(1,1)	Component	Persistent
Moneyness	RMSE	RMSE Ratio	RMSE Ratio	RMSE	RMSE Ratio	RMSE Ratio
S/X<0.975	1.078	0.873	0.987	1.461	0.556	1.028
0.975<S/X<1.00	1.054	0.927	0.977	1.631	0.644	0.671
1.00<S/X<1.025	0.956	0.947	1.003	1.356	0.749	0.763
1.025<S/X<1.05	0.919	0.936	0.941	1.008	0.883	0.876
1.05<S/X<1.075	1.032	0.902	0.891	0.860	0.902	1.071
1.075<S/X	1.079	0.883	0.861	1.166	0.977	1.068
All	1.038	0.897	0.948	1.284	0.765	0.933
Maturity						
DTM<20	0.610	0.923	0.903	0.813	0.994	0.994
20<DTM<80	0.976	0.865	0.859	1.124	0.828	0.782
80<DTM<180	1.124	0.883	0.891	1.258	0.778	0.914
DTM>180	1.151	0.960	1.121	1.822	0.652	1.079
All	1.038	0.897	0.948	1.284	0.765	0.933

	Panel C. 1992-1994 in Sample			Panel D. 1995 Out of Sample		
	GARCH(1,1)	Component	Persistent	GARCH(1,1)	Component	Persistent
Moneyness	RMSE	RMSE Ratio	RMSE Ratio	RMSE	RMSE Ratio	RMSE Ratio
S/X<0.975	1.122	0.604	1.010	1.771	0.757	0.781
0.975<S/X<1.00	1.275	0.747	0.964	1.546	0.783	0.985
1.00<S/X<1.025	1.228	0.760	0.918	1.389	0.708	0.950
1.025<S/X<1.05	1.002	0.790	0.900	1.110	0.705	1.021
1.05<S/X<1.075	0.991	0.785	0.856	0.896	0.796	1.092
1.075<S/X	1.032	0.919	0.969	0.681	1.026	1.286
All	1.107	0.773	0.956	1.227	0.792	0.957
Maturity						
DTM<20	0.857	0.909	0.899	0.744	0.978	0.952
20<DTM<80	1.009	0.808	0.927	0.846	0.929	1.129
80<DTM<180	1.045	0.717	1.000	1.187	0.847	1.016
DTM>180	1.422	0.744	0.959	1.848	0.670	0.848
All	1.107	0.773	0.956	1.227	0.792	0.957

Notes to Table: We use the NLS estimates from Tables 5 and 6 to compute the root mean squared option valuation error (RMSE) for various moneyness and maturity bins. The RMSE is reported in levels for the GARCH(1,1) model and for the component and persistent component models we report the RMSE ratio with the GARCH(1,1) model.

Research Papers
2008

CREATES
Center for Research in Econometric Analysis of Time Series

2007-44	Charles S. Bos, Siem Jan Koopman and Marius Ooms: Long memory modelling of inflation with stochastic variance and structural breaks
2007-45	James Davidson and Nigar Hashimzade: Representation and Weak Convergence of Stochastic Integrals with Fractional Integrator Processes
2008-01	John A. Carlson, Christian M. Dahl, and Carol L. Osler: Short-run Exchange-Rate Dynamics: Theory and Evidence
2008-02	Peter Reinhard Hansen: Reduced-Rank Regression: A Useful Determinant Identity
2008-03	Søren Johansen, Katarina Juselius, Roman Frydman and Michael Goldberg: Testing hypotheses in an I(2) model with applications to the persistent long swings in the Dmk/$ rate
2008-04	Olaf Posch: Explaining output volatility: The case of taxation
2008-05	Annastiina Silvennoinen and Timo Teräsvirta: Modelling Multivariate Autoregressive Conditional Heteroskedasticity with the Double Smooth Transition Conditional Correlation GARCH Model
2008-06	Annastiina Silvennoinen and Timo Teräsvirta: Multivariate GARCH models. To appear in T. G. Andersen, R. A. Davis, J.-P. Kreiss and T. Mikosch, eds. Handbook of Financial Time Series. New York: Springer.
2008-07	Changli He, Annastiina Silvennoinen and Timo Teräsvirta: Parameterizing unconditional skewness in models for financial time series
2008-08	Cristina Amado and Timo Teräsvirta: Modelling Conditional and Unconditional Heteroskedasticity with Smoothly Time-Varying Structure
2008-09	Søren Johansen and Bent Nielsen: An analysis of the indicator saturation estimator as a robust regression estimator
2008-10	Peter Christoffersen, Kris Jacobs, Christian Dorion and Yintian Wang: Volatility Components, Affine Restrictions and Non-Normal Innovations
2008-11	Peter Christoffersen, Kris Jacobs, Chayawat Ornthanalai and Yintian Wang: Option Valuation with Long-run and Short-run Volatility Components