Option Valuation with Conditional Heteroskedasticity and Non-Normality*

Peter Christoffersen
McGill University and CREATES

Redouane Elkamhi
McGill University

Bruno Feunou
Université de Montréal

Kris Jacobs
McGill University and Tilburg

March 18, 2008

Abstract

We provide results for the valuation of European style contingent claims for a large class of specifications of the underlying asset returns. Our valuation results obtain in a discrete time, infinite state-space setup using an equivalent martingale measure and the no-arbitrage principle. Our approach allows for general forms of heteroskedasticity in returns, and valuation results for homoskedastic processes can be obtained as a special case. It also allows for conditional non-normal return innovations, which is critically important because heteroskedasticity alone does not suffice to capture the option smirk. For the rich class of infinitely divisible distributions we show that the risk-neutral return dynamics take the same distributional form as the physical return dynamics. Interestingly, non-normality drives a wedge between the physical and the risk-neutral variance which is in accordance with previous empirical observations. Our framework nests the valuation results obtained by Duan (1995) and Heston and Nandi (2000) by allowing for a time-varying price of risk and non-normal innovations. We provide extensions of our results to discrete time stochastic volatility models, and we analyze the relation between our results and those obtained for continuous time models. An empirical example demonstrates the usefulness of conditional non-normality for the modeling of the index option smirk.

JEL Classification: G12

Keywords: GARCH; risk-neutral valuation; no-arbitrage; non-normal innovations

---

*Christoffersen and Jacobs are also affiliated with CIRANO and CIREQ and want to thank FQRSC, IFM^2 and SSHRC for financial support. Elkamhi and Feunou were supported by grants from IFM^2. For helpful comments we thank Tim Bollerslev, Mo Chaudhury, Rene Garcia, Tom McCurdy, Nour Meddahi, Karim Mimouni, Eric Renault, and Marcel Rindisbacher. Any remaining inadequacies are ours alone. Correspondence to: Peter Christoffersen, Desautels Faculty of Management, McGill University, 1001 Sherbrooke Street West, Montreal, Quebec, Canada, H3A 1G5; Tel: (514) 398-2869; Fax: (514) 398-3876; E-mail: peter.christoffersen@mcgill.ca.

1 Introduction

A contingent claim is a security whose payoff depends upon the value of another underlying security. A valuation relationship is an expression that relates the value of the contingent claim to the value of the underlying security and other variables. The most popular approach for valuing contingent claims is the use of a Risk Neutral Valuation Relationship (RNVR).

Most of the literature on contingent claims and most of the applications of the RNVR have been cast in continuous time. While the continuous-time approach offers many advantages, the valuation of contingent claims in discrete time is also of substantial interest. For example, when hedging option positions, rebalancing decisions must be made in discrete time. In the case of American and exotic options, early exercise decisions must be made in discrete time as well. Moreover, as only discrete observations are available for empirical study, discrete time models are often more econometrically tractable.

As a result, most of the stylized facts characterizing underlying securities have been studied in discrete time models. One very important feature of returns is conditional heteroskedasticity, which can be addressed in the GARCH framework of Engle (1982) and Bollerslev (1986).[1] Presumably, because of this evidence, most of the recent empirical work on discrete time option valuation has also focused on GARCH processes.[2] The GARCH model amounts to an infinite state space setup, with the innovations for underlying asset returns described by continuous distributions. In this case the market is incomplete, and it is in general not possible to construct a portfolio containing combinations of the contingent claim and the underlying asset that make the resulting portfolio riskless.[3]

To obtain a RNVR, the GARCH option valuation literature builds on the approach of Rubinstein (1976) and Brennan (1979), who demonstrate how to obtain RNVRs for lognormal and normal returns in the case of constant mean return and volatility, by specifying a representative agent economy and characterizing sufficient conditions on preferences.[4] For a given dynamic of the underlying security, specific assumptions have to be made on preferences in order to obtain a risk neutralization result. The first order condition resulting from this economy yields an Euler equation that can be used to price any asset. For lognormal stock returns and a conditionally heteroskedastic (GARCH) volatility dynamic, the standard result is the one in Duan (1995). Duan's result relies on the existence of a representative agent with constant relative risk aversion or constant absolute risk aversion.[5]

[1] See for example French, Schwert and Stambaugh (1987) and Schwert (1989) for early studies on stock returns. The literature is far too voluminous to cite all relevant papers here. See Bollerslev, Chou and Kroner (1992) and Diebold and Lopez (1995) for reviews on GARCH modeling.

[2] See Bollerslev and Mikkelsen (1996), Satchell and Timmermann (1996), Garcia and Renault (1998), Heston and Nandi (2000), Christoffersen and Jacobs (2004), and Christoffersen, Heston and Jacobs (2006) for applications to option valuation.

[3] In a discrete time finite state space setting, Harrison and Pliska (1981) provide the mathematical framework to obtain the existence of the risk neutral probability measure, to demonstrate uniqueness in the case of complete markets, and to get a RNVR for any contingent claim. See also Harrison and Kreps (1979), Cox, Ross and Rubinstein (1979) and Cox and Ross (1976) for discrete-time finite state-space approaches.

[4] Brennan (1979) characterizes the bivariate distribution of returns on aggregate wealth and the underlying asset under which a risk-neutral valuation relationship obtains in the homoskedastic case. Camara (2003) uses this approach to obtain valuation results for transformed normal dynamics of returns and state variables. See also Schroder (2004).

[5] Amin and Ng (1993) also study the heteroskedastic case. Although they formulate the problem in terms of the economy's stochastic discount factor, they start by making an assumption on the bivariate distribution

However, because it is difficult to characterize the general equilibrium setup underlying a RNVR, very few valuation results are currently available for heteroskedastic processes with non-normal innovations.[6] In this paper, we argue that it is possible to investigate option valuation for a large class of conditionally non-normal heteroskedastic processes, provided that the conditional moment generating function exists. It is also possible to accommodate a large class of time-varying risk premia. Our framework differs from the approach in Brennan (1979) and Duan (1995), and is more intimately related to the approach adopted in continuous-time option valuation: we only use the no-arbitrage assumption and some technical conditions on the investment strategies to show the existence of an RNVR. We demonstrate the existence of an EMM and characterize it, without first making an explicit assumption on the utility function of a representative agent. We then show that the price of the contingent claim defined as the expected value of the discounted payoff at maturity is a no-arbitrage price and characterize the risk-neutral dynamic.

Why are we able to provide more general valuation results than the existing literature? In our opinion, the analysis in Brennan (1979) and Duan (1995) addresses two important questions simultaneously: First, a mostly technical question that characterizes the risk-neutral dynamic and the valuation of options; second, a more economic one that characterizes the equilibrium underlying the valuation procedure. The existing discrete-time literature for the most part has viewed these two questions as inextricably linked, and has therefore largely limited itself to (log)normal return processes as well as a few special non-normal cases. Our paper differs in a subtle but important way from most of these studies. We argue that it is possible and desirable to treat these questions one at a time. We do not attempt to characterize the bivariate distribution of preferences and returns that gives rise to the risk-neutral valuation relationship. Instead, we make the following two assumptions. First, we assume a class of Radon-Nikodym derivatives and search for the EMM within this class. Second, we specify the volatility risk premium. This allows us to provide some general results on the valuation of options under conditionally non-normal asset returns without resorting to equilibrium techniques. We also show how the normal model and the available conditional non-normal models are special cases of our setup.

The same approach to separate these two questions occurs in the literature on option valuation using continuous-time stochastic volatility models, such as for instance in Heston's (1993a) model. These models also yield different equivalent martingale measures for different specifications of the volatility risk premium. For a given specification of the volatility risk premium, we can find an EMM and characterize the risk-neutral dynamic using Girsanov's theorem. To derive this result, and to value options, there is no need to explicitly characterize the utility function underlying the volatility risk premium. The latter task is very interesting in its own right, but differs from characterizing the risk-neutral dynamic and the option value for a given physical return dynamic.[7] The latter is a purely mathematical exercise. The former provides the economic background behind a particular choice of volatility premium, and therefore helps us understand whether a particular choice of functional form for the risk premium, which is often made for convenience, is also reasonable from an economic

of the stochastic discount factor and the underlying return process.

[6] Duan, Ritchken and Sun (2005) analyze a heteroskedastic model with Poisson-normal innovations and Duan (1999) analyzes a conditionally fat-tailed heteroskedastic model. Christoffersen, Heston and Jacobs (2006) use a heteroskedastic return dynamic with inverse Gaussian innovations.

[7] See for instance Heston (1993a) and Bates (1996, 2000) for a discussion.

perspective.

Our paper provides a set of tools that can be used to value options for a large class of discrete-time return dynamics that are characterized by heteroskedasticity and non-normal innovations. Whether this valuation exercise makes sense from an economic perspective depends on the nature of the assumed risk premium, and the general equilibrium setup that gives rise to such risk premium. There are two questions: a mostly technical one that characterizes the risk-neutral dynamic and the valuation of options, and a more economic one that characterizes the equilibrium underlying this valuation procedure. In our opinion, the existing discrete-time literature for the most part has viewed these two questions as inextricably linked, and has therefore largely limited itself to (log)normal return processes. We argue that it is possible and desirable to treat these questions one at a time, and we provide new results on the question of option valuation with conditionally non-normal returns.

There is a growing literature that values options for discrete-time return dynamics with non-normal innovations. A number of other papers obtain risk-neutral valuation relationships either under the maintained assumption of a specific kind of non-normal innovations, or under the maintained assumption of a particular form of heteroskedasticity, or both.[8] Combining non-normality and heteroskedasticity attempts to correct the biases associated with the conditionally normal GARCH model. These biases are similar to those displayed by the Heston (1993) model, which the continuous-time literature has sought to remedy by adding (potentially correlated) jumps in returns and volatility. This paper is therefore also related to empirical studies of jump models.[9]

The paper proceeds as follows. In Section 2 we define a class of heteroskedastic stock return processes, and we characterize the condition for an EMM for this class of processes. We then show sufficient conditions for an EMM to exist and we derive the risk neutral distribution of returns. In Section 3 we derive the no-arbitrage option price corresponding to the EMM. Section 4 discusses several special cases of return dynamics that can be analyzed using our approach. Section 5 uses one of the special cases for an empirical illustration. Section 6 provides continuous time limits of some of the specific models. Section 7 introduces an extension to discrete time stochastic volatility models and discusses issues regarding option price uniqueness in discrete and continuous time. Section 8 concludes.

2 Conditionally heteroskedastic models

We define the probability space $(\Omega, \mathbb{F}, P)$ to describe the physical distribution of the states of nature. The financial market consists of a zero-coupon risk-free bond index and a stock. The dynamics of the bond are described by the process $\{B_t\}_{t=1}^{T}$ normalized to $B_0 = 1$ and the dynamics of the stock price by $\{S_t\}_{t=1}^{T}$. The information structure is given by the filtration $\mathbb{F} = \{F_t | \, t = 1, ..., T\}$ generated by the stock and the bond process.

[8] Madan and Seneta (1990) use the symmetric and i.i.d. variance gamma distribution. Heston (1993b) presents results for the gamma distribution and Heston (2004) analyzes a number of infinitely divisible distributions. Camara (2003) uses a transformed normal innovation and Duan (1999) uses a heteroskedastic model with a transformed normal innovation. Christoffersen, Heston and Jacobs (2006) analyze a heteroskedastic return process with inverse Gaussian innovations.

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

2.1 The stock price process

The underlying stock price process is assumed to follow the conditional distribution D under the physical measure P. We write

$$R_t \equiv \ln\left(\frac{S_t}{S_{t-1}}\right) = \mu_t - \gamma_t + \varepsilon_t \qquad\qquad \varepsilon_t|F_{t-1} \sim D(0,\sigma_t^2) \tag{2.1}$$

where S_t is the stock price at time t, and σ_t^2 is the conditional variance of the log return in period t. The mean correction factor, γ_t, is defined from

$$\exp(\gamma_t) \equiv E_{t-1}\left[\exp(\varepsilon_t)\right]$$

and it serves to ensure that the conditional expected gross rate of return, $E_{t-1}[S_t/S_{t-1}]$, is equal to $\exp(\mu_t)$. More explicitly,

$$\begin{aligned} E_{t-1}\left[S_t/S_{t-1}\right] &= E_{t-1}\left[\exp(\mu_t - \gamma_t + \varepsilon_t)\right] = \exp(\mu_t) \\ &\iff \exp(\gamma_t) = E_{t-1}\left[\exp(\varepsilon_t)\right] \end{aligned}$$

Note that our specification (2.1) does not restrict the risk premium in any way nor does it assume conditional normality.

We follow most of the existing discrete-time empirical finance literature by focusing on conditional means μ_t and conditional variances σ_t^2 that are F_{t-1} measurable. We do not constrain the interest rate r_t to be constant. It is instead assumed to be an element of F_{t-1} as well. Our framework is able to accommodate the class of ARCH and GARCH processes proposed by Engle (1982) and Bollerslev (1986) and used for option valuation by Amin and Ng (1993), Duan (1995, 1999), and Heston and Nandi (2000). Our results also hold for different types of GARCH specifications, such as the EGARCH model of Nelson (1991) or the specification of Glosten, Jagannathan and Runkle (1993).

In the following, we show that we can find an EMM by defining a probability measure that makes the discounted security process a martingale. We derive more general results on option valuation for heteroskedastic processes compared to the available literature, because we focus on the narrow question of option valuation while ignoring the economic question regarding the preferences of the representative agent that support this valuation argument in equilibrium.

We use a no-arbitrage argument that is similar to the one used in the continuous-time literature. We first prove the existence of an EMM. Subsequently we demonstrate the existence of a RNVR by demonstrating that the price of the contingent claim, defined as the expected value of the discounted payoff at maturity, is a no-arbitrage price under this EMM.[10] The proof uses an argument similar to the one used in the continuous-time literature, but is arguably more straightforward as it avoids the technical issues involved in the analysis of local and super martingales.

[10] Duan (1995) refers to RNVR as Local RNVR in the case of GARCH. The reason for the distinction is that the conditional volatility is identical under the two measures only one period ahead. In the remainder of the paper we will drop this distinction for ease of exposition. We emphasize that the result that the conditional volatility differs between the two measures for more than one period ahead is to be expected as volatility is random in this case. This feature is very similar to the continuous time case, which has random volatility for any horizon.

2.2 Specifying an equivalent martingale measure

The objective in this section is to find a measure equivalent to the physical measure P that makes the price of the stock discounted by the riskless asset a martingale. An EMM is defined as long as the Radon-Nikodym derivative is defined. We start by specifying a candidate Radon-Nikodym derivative of a probability measure. We then show that this Radon-Nikodym derivative defines an EMM that makes the discounted stock price process a martingale. This result in turn allows us to obtain the distribution of the stock return under this EMM.

For a given sequence of a random variable, ν_t, we define the following candidate Radon-Nikodym derivative

$$\left.\frac{dQ}{dP}\right| F_t = \exp\left(-\sum_{i=1}^{t}\left(\nu_i\varepsilon_i + \Psi_i\left(\nu_i\right)\right)\right) \tag{2.2}$$

where $\Psi_t(u)$ is defined as the natural logarithm of the moment generating function

$$E_{t-1}\left[\exp(-u\varepsilon_t)\right] \equiv \exp\left(\Psi_t(u)\right)$$

Note that we can think of the mean correction factor in (2.1) as $\gamma_t = \Psi_t(-1)$. Note also that in the normal case we have $\Psi_t(u) = \frac{1}{2}\sigma_t^2 u^2$ and $\gamma_t = \Psi_t(-1) = \frac{1}{2}\sigma_t^2$.

We can now show the following lemma

Lemma 1 $\left.\frac{dQ}{dP}\right| F_t$ *is a Radon-Nikodym derivative*

Proof. We need to show that $\left.\frac{dQ}{dP}\right| F_t > 0$ which is immediate. We also need to show that $E_0^P\left[\left.\frac{dQ}{dP}\right| F_t\right] = 1$. We have

$$E_0^P\left[\left.\frac{dQ}{dP}\right| F_t\right] = E_0^P\left[\exp\left(-\sum_{i=1}^{t}\left(\nu_i\varepsilon_i + \Psi_i\left(\nu_i\right)\right)\right)\right].$$

Using the law of iterative expectations we can write

$$\begin{aligned}
E_0^P\left[\left.\frac{dQ}{dP}\right| F_t\right] &= E_0^P\left[E_1^P ... E_{t-1}^P \exp\left(-\sum_{i=1}^{t}\left(\nu_i\varepsilon_i + \Psi_i\left(\nu_i\right)\right)\right)\right] \\
&= E_0^P\left[E_1^P ... E_{t-2}^P \exp\left(-\sum_{i=1}^{t-1}\nu_i\varepsilon_i - \sum_{i=1}^{t}\Psi_i\left(\nu_i\right)\right) E_{t-1}^P \exp\left(-\nu_t\varepsilon_t\right)\right] \\
&= E_0^P\left[E_1^P ... E_{t-2}^P \exp\left(-\sum_{i=1}^{t-1}\nu_i\varepsilon_{i+1} - \sum_{i=1}^{t}\Psi_i\left(\nu_i\right)\right) \exp\left(\Psi_t\left(\nu_t\right)\right)\right] \\
&= E_0^P\left[E_1^P ... E_{t-2}^P \exp\left(-\sum_{i=1}^{t-1}\nu_i\varepsilon_i - \sum_{i=1}^{t-1}\Psi_i\left(\nu_i\right)\right)\right]
\end{aligned}$$

Iteratively, using this result we get

$$\begin{aligned}
E_0^P\left[\left.\frac{dQ}{dP}\right| F_t\right] &= E_0^P\left[\exp\left(-\nu_1\varepsilon_1 - \Psi_1\left(\nu_1\right)\right)\right] \\
&= \exp\left(-\Psi_1\left(\nu_1\right)\right)\exp\left(\Psi_1\left(\nu_1\right)\right) \\
&= 1
\end{aligned}$$

and the lemma obtains. ■

We are now ready to show that we can specify an EMM using this Radon-Nikodym derivative.

Proposition 1 *The probability measure Q defined by the Radon-Nikodym derivative in (2.2) is an EMM if and only if*

$$\Psi_t(\nu_t - 1) - \Psi_t(\nu_t) - \gamma_t + \alpha_t \sigma_t^2 = 0$$

where $\alpha_t = \dfrac{\mu_t - r_t}{\sigma_t^2}$.

Proof. We need $E^Q\left[\left.\frac{S_t}{B_t}\right| F_{t-1}\right] = \frac{S_{t-1}}{B_{t-1}}$ or equivalently $E^Q\left[\left.\frac{S_t}{S_{t-1}} / \frac{B_t}{B_{t-1}}\right| F_{t-1}\right] = 1$. We have

$$\begin{aligned}
E^Q\left[\left.\frac{S_t}{S_{t-1}} / \frac{B_t}{B_{t-1}}\right| F_{t-1}\right] &= E^P\left[\left.\left(\frac{\frac{dQ}{dP}\big| F_t}{\frac{dQ}{dP}\big| F_{t-1}}\right)\frac{S_t}{S_{t-1}} / \frac{B_t}{B_{t-1}}\right| F_{t-1}\right] \\
&= E^P\left[\left.\left(\frac{\frac{dQ}{dP}\big| F_t}{\frac{dQ}{dP}\big| F_{t-1}}\right)\frac{S_t}{S_{t-1}}\exp(-r_t)\right| F_{t-1}\right] \\
&= E^P\left[\exp\left(-\nu_t\varepsilon_t - \Psi_t(\nu_t)\right)\exp(\mu_t - \gamma_t + \varepsilon_t)\exp(-r_t)\right| F_{t-1}] \\
&= \exp\left(-\Psi_t(\nu_t) + \mu_t - r_t - \gamma_t\right) E^P\left[\exp\left((1-\nu_t)\varepsilon_t\right)|F_{t-1}\right] \\
&= \exp\left(-\Psi_t(\nu_t) + \mu_t - r_t - \gamma_t + \Psi_t(\nu_t - 1)\right) \\
&= \exp(\Psi_t(\nu_t - 1) - \Psi_t(\nu_t) - \Psi_t(-1) + \alpha_t\sigma_t^2)
\end{aligned}$$

Thus Q is a probability measure that makes the stock discounted by a riskless asset a martingale if and only if

$$\Psi_t(\nu_t - 1) - \Psi_t(\nu_t) - \gamma_t + \alpha_t \sigma_t^2 = 0 \tag{2.3}$$

This result implies that we can construct an EMM by choosing the sequence, ν_t, to make (2.3) hold. ■

2.3 Solving for the EMM

In this section we develop various results on the existence of a solution to (2.3).

Note first that in the conditional normal special case we get the solution to be the well-known price of risk $\nu_t = \alpha_t = (\mu_t - r_t)/\sigma_t^2$. Note also that if we additionally specify the conditional mean of the excess return to be affine in σ_t^2, so that $\mu_t = r_t + \lambda\sigma_t^2$, ν_t is a constant, namely the price of risk λ.

When allowing for conditional non-normal returns, a multitude of possible return distributions and thus $\Psi_t(.)$ functions are possible, and we need to put some structure on $\Psi_t(.)$ in order to analyze the existence of a solution to (2.3). In Section 4 below we consider some important non-normal special cases where an explicit solution for ν_t can be found. In the more general case, we provide the following result.

Proposition 2 *If Ψ is strictly convex, twice differentiable, and tends to infinity at the boundaries of its domain (u_1, u_2) where $u_1 + 1 < u_2$, then there exists a unique solution to equation (2.3).*

Proof. See the Appendix. ■

Proposition 2 provides a set of sufficient conditions for a unique solution to exist but the conditions are not necessary. For example, a similar result can be obtained assuming that Ψ is strictly concave. For the class of infinitely divisible distributions we will consider below that strict convexity holds (Feller, 1968) and so we show the above proof for this case.[11]

In the absence of these types of sufficient conditions, we can find an explicit approximate solution to the EMM equation in (2.3) from the second-order approximations

$$\begin{aligned}\Psi_t(\nu_t - 1) &\approx \Psi_t(0) + \Psi_t'(0)(v_t - 1) + \tfrac{1}{2}\Psi_t''(0)(v_t - 1)^2 \\ \Psi_t(\nu_t) &\approx \Psi_t(0) + \Psi_t'(0)\, v_t + \tfrac{1}{2}\Psi_t''(0)\, v_t^2\end{aligned}$$

By definition of the mean-zero shock ε_t we have that $\Psi_t'(0) = 0$, and $\Psi_t''(0) = \sigma_t^2$ so that the approximation along with the EMM condition in (2.3) gives us

$$v_t \approx \frac{\mu_t - r_t}{\sigma_t^2} + \frac{1}{2} - \frac{\gamma_t}{\sigma_t^2} \tag{2.4}$$

Notice that this approximation is exact in the normal case where $\gamma_t = \frac{1}{2}\sigma_t^2$ and $\nu_t = (\mu_t - r_t)/\sigma_t^2$. This approximate solution can be used in place of the exact solution, or it can be used as a starting value in a numerical search for the exact v_t.

Note finally that equation (2.3) suggests that if one is willing to put more structure on the return process in (2.1) then the problem of finding a solution for ν_t can be circumvented altogether. Denote the risk premium, $\mu_t - r_t$, by π_t. Then if the return dynamic is specified such that

$$R_t = r_t + \pi_t - \gamma_t + \varepsilon_t \qquad \varepsilon_t | F_{t-1} \sim D(0, \sigma_t^2)$$

where

$$\pi_t = \Psi_t(\nu_t) - \Psi_t(\nu_t - 1) + \gamma_t \tag{2.5}$$

then the EMM condition in (2.3) is trivially satisfied for any value of v_t. Thus v_t can be set to a constant parameter, v, to be estimated as part of the return dynamic. This approach is viable but suffers from the drawback that the return mean dynamic is now chosen for convenience rather than for empirical relevance. Note that in the normal case this approach yields

$$\begin{aligned}R_t &= r_t + \Psi_t(\nu) - \Psi_t(\nu - 1) + \varepsilon_t & \varepsilon_t | F_{t-1} \sim N(0, \sigma_t^2) \\ &= r_t + v\sigma_t^2 + \varepsilon_t & \varepsilon_t | F_{t-1} \sim N(0, \sigma_t^2)\end{aligned}$$

which corresponds to an affine risk-premium.

[11] Gourieroux and Monfort (2007) found a similar result in the context of the specification of a stochastic discount factor. They do not relate their result to the class of inifinitely divisible distributions.

2.4 Characterizing the risk-neutral distribution

When pricing options using Monte Carlo simulation, knowing the risk neutral distribution is valuable. In this section, we derive an important result that shows that for the class of models we investigate, the risk neutral distribution is from the same family as the original physical distribution.

We first need the following lemma, where we recall that $\Psi_t(u)$, denotes the one-day log conditional moment generating function of the innovation ε_t

Lemma 2

$$E_{t-1}^{Q}\left[\exp\left(-u\varepsilon_t\right)\right] = \exp\left(\Psi_t\left(\nu_t + u\right) - \Psi_t\left(\nu_t\right)\right)$$

Proof.

$$\begin{aligned} E_{t-1}^{Q}\left[\exp\left(-u\varepsilon_t\right)\right] &= E^P\left[\left(\frac{\frac{dQ}{dP}|F_t}{\frac{dQ}{dP}|F_{t-1}}\right)\exp(-u\varepsilon_t)|F_{t-1}\right] \\ &= E^P\left[\exp\left(-\nu_t\varepsilon_t - \Psi_t\left(\nu_t\right)\right)\exp(-u\varepsilon_t)|F_{t-1}\right] \\ &= \exp\left(\Psi_t\left(\nu_t + u\right) - \Psi_t\left(\nu_t\right)\right) \end{aligned}$$

■

From this lemma, if we define $\Psi_t^Q(u)$ to be the log conditional moment generating function of ε_t under the risk neutral probability measure, then we have

$$\Psi_t^Q(u) = \Psi_t\left(\nu_t + u\right) - \Psi_t\left(\nu_t\right) \tag{2.6}$$

From this we can derive

$$E_{t-1}^{Q}\left[\varepsilon_t\right] = \left.\frac{\partial \Psi_t^Q(-u)}{\partial u}\right|_{u=0} = -\Psi_t'\left(\nu_t\right)$$

which represents the risk premium. Define the risk neutral innovation

$$\varepsilon_t^* \equiv \varepsilon_t - E_{t-1}^{Q}\left[\varepsilon_t\right] = \varepsilon_t + \Psi_t'\left(\nu_t\right) \tag{2.7}$$

The risk-neutral log-conditional moment generating function of ε_t^*, labeled $\Psi_{t-1}^{Q*}(u)$, is then

$$\Psi_t^{Q*}(u) = -u\Psi_t'\left(\nu_t\right) + \Psi_t^Q(u) \tag{2.8}$$

We are now ready to show the following

Proposition 3 *If the physical conditional distribution of ε_t is an infinitely divisible distribution with finite second moment, then the risk-neutral conditional distribution of ε_t^* is also an infinitely divisible distribution with finite second moment.*

Proof. See the Appendix. ■

In the special case of the normal distribution we get simply

$$\varepsilon_t^* = \varepsilon_t + \Psi_t'\left(\nu_t\right) = \varepsilon_t + \mu_t - r_t$$

and

$$\Psi_t^{Q*}(u) = \frac{1}{2}\sigma_t^2 u^2$$

so that the risk-neutral innovations are normal and correspond to the physical shifted by the equity risk premium. In the more general case the relationship will not necessarily be so simple.

Because of the one-to-one mapping between moment generating functions and distribution functions, this result can be used to derive specific parametric risk-neutral distributions corresponding to the parametric physical distributions assumed by the researcher.

2.5 Characterizing the risk-neutral conditional variance

The conditional risk-neutral variance, σ_t^{*2} is of particular interest in the dynamic heteroskedastic models we consider. It can be obtained by taking the second order derivative of the risk neutral log moment generating function $\Psi_t^{Q*}(u)$ and evaluating it at $u = 0$. We can write

$$\sigma_t^{*2} = \left.\frac{\partial^2 \Psi_t^{Q*}(u)}{\partial u^2}\right|_{u=0}$$

Using equation (2.8) and (2.6) we get

$$\sigma_t^{*2} = \left.\frac{\partial^2 \Psi_t^{Q}(u)}{\partial u^2}\right|_{u=0} = \left.\frac{\partial^2 \Psi_t(v_t + u)}{\partial u^2}\right|_{u=0} = \Psi_t''(\nu_t)$$

Recall that by definition, the conditional variance under the physical measure is

$$\sigma_t^2 = \Psi_t''(0)$$

Thus in general we have the following relationship between conditional variance under the two measures

$$\sigma_t^{*2} = \sigma_t^2 \frac{\Psi_t''(\nu_t)}{\Psi_t''(0)}$$

In the case of conditionally normal returns, we have $\Psi_t(u) = \frac{1}{2}\sigma_t^2 u^2$ and $\nu_t = (\mu_t - r_t)/\sigma_t^2$, so that $\Psi_t''(\nu_t) = \Psi_t''(0)$ and thus $\sigma_t^{*2} = \sigma_t^2$ but this will not be true in general for non-normal distributions. Non-normality will drive a wedge between the physical and risk-neutral conditional variances. This phenomenon is often observed empirically, as physical volatility measures from historical returns are systematically lower than risk-neutral volatilities implied from options. See for example Carr and Wu (2007). We will give explicit examples of non-normal distributions that generate this interesting and important empirical feature below.

3 The valuation of European style contingent claims

We have demonstrated in a general return model with time-varying conditional mean and volatility and non-normal shocks, conditions under which there exists an EMM Q that makes the stock discounted by the riskless asset a martingale.

We now turn our attention to the pricing of European style contingent claims. Existing papers on the pricing of contingent claims in a discrete-time infinite state space setup, such as the literature on GARCH option pricing in Duan (1995), Amin and Ng (1993) and Heston and Nandi (2000) value such contingent claims by making an assumption on the bivariate distribution of the stock return and the endowment, or an equivalent assumption. While this approach, which most often amounts to the characterization of the equilibrium that supports the pricing, is an elegant way to deal with the incompleteness that characterizes these markets, we argue that it is not strictly necessary to characterize the equilibrium. Instead, we adopt an approach which is more prevalent in the continuous-time literature, and proceed to pricing derivatives using a no-arbitrage argument alone.

To understand our approach, the analogy with option valuation for the stochastic volatility model of Heston (1993a) is particularly helpful. In this incomplete markets setting, an infinity of no-arbitrage contingent claims prices exist, one for every different specification of the price of risk. When one fixes the price of volatility risk, however, there is a unique no-arbitrage price. For the purpose of option valuation, one can simply pick a price of volatility risk, and the resulting valuation exercise is purely mechanical.

The question whether a particular price of risk is reasonable is of substantial interest in its own right, and an analysis of the representative agent utility function that support a particular price of risk is very valuable. However, this question can be analyzed separately from the option valuation problem. For the heteroskedastic discrete-time models we consider, a similar remark applies. We can value options provided we specify the price of risk. The link between our approach and the utility-based approach in Brennan (1979), Rubinstein (1976) and Duan (1995) is that assumptions on the utility function are implicit in the specification of the risk premium in the return dynamic in our case.[12] The representative agent preferences underlying this assumption are of interest, but it is not necessary to analyze them in order to value options. Of course, we note that the main difference with the continuous-time stochastic volatility models is that GARCH models are one-shock models, and that therefore there is only one price of risk.

We have already found an EMM Q. We therefore want to demonstrate that the price at time t is defined as

$$C_t = E^Q \left[\frac{C_T(S_T)}{B_T} B_t \middle| F_t \right].$$

The proof proceeds in a number of steps and requires defining a number of concepts that are well-known in the literature. Fortunately, even though our methodology closely follows the continuous-time case, we economize on the number of technical conditions in the continuous-time setup, such as admissibility, and avoid the concepts of local martingale and super martingale. The reason is that the integration over an infinite number of trading times in the continuous-time case is replaced by a finite sum over the trading days in discrete time.

Definitions

1. We denote by η_t, δ_t and ψ_t the units of the stock, the contingent claim and the bond held at date t. We refer to the F_t predictable processes η_t, δ_t and ψ_t as investment

[12] See Bick (1990) and He and Leland (1993) for a discussion of assumptions on the utility function implicit in the specification of the return dynamic for the market portfolio. We proceed along the lines of Jacob and Shiryaev (1998), and Shiryaev (1999).

strategies.

2. The value process
$$V_t = \eta_t S_t + \delta_t C_t + \psi_t B_t$$
describes the total dollar amount available for investments at date t.

3. The gain process
$$G_t = \sum_{i=0}^{t-1} \eta_i (S_{i+1} - S_i) + \sum_{i=0}^{t-1} \delta_i (C_{i+1} - C_i) + \sum_{i=0}^{t-1} \psi_i (B_{i+1} - B_i).$$
captures the total financial gains between dates 0 and t.

4. We call the process $\{\eta_t, \delta_t, \psi_t\}_{t=0}^{T-1}$ a self financing strategy if and only if $V_t = G_t$ $\forall t = 1, ..., T$.

5. The definition of an arbitrage opportunity is standard: we have an arbitrage opportunity if a self financing strategy exists with either $V_0 < 0$, $V_T \geq 0$ a.s. or $V_0 \leq 0$, $V_T > 0$ a.s.

6. We denote the discounted stock price at time t as $S_t^B = \frac{S_t}{B_t}$ and the discounted contingent claim as $C_t^B = \frac{C_t}{B_t}$. Similarly, the discounted value process is denoted $V_t^B = \frac{V_t}{B_t}$ and the discounted gain process $G_t^B = \frac{G_t}{B_t}$.

Note that for a self financing strategy, we have $V_t^B = G_t^B$ because $V_t = G_t$ and $B_t > 0$. Furthermore, we can show the following.

Lemma 3 *For a self financing strategy we have*
$$G_t^B = \sum_{i=0}^{t-1} \eta_i (S_{i+1}^B - S_i^B) + \sum_{i=0}^{t-1} \delta_i (C_{i+1}^B - C_i^B) \qquad \forall t = 1, ..., T$$

Proof. The proof involves straightforward but somewhat cumbersome algebraic manipulations of the above definitions. See the Appendix for the details. ■

We know that under the EMM we defined, the stock discounted by the risk free asset is a martingale. We now need to show that the contingent claims prices obtained by computing the expected value of the final payoff discounted by the risk free asset also constitute a martingale under this EMM.

Lemma 4 *The stochastic process defined by the discounted values of the candidate contingent claims prices is an F_t martingale under the EMM.*

Proof. We defined our candidate process for the contingent claims price under the EMM as $C_t = E^Q \left[\frac{C_T(S_T)}{B_T} B_t \middle| F_t \right]$. The process for the discounted values of the contingent claims prices is then defined as
$$C_t^B \equiv \frac{C_t}{B_t} = E^Q \left[\frac{C_T(S_T)}{B_T} \middle| F_t \right]$$

We use the fact that the conditional expectation itself is a Q martingale. This in turn follows from the law of iterated expectations and the European style payoff function. Taking conditional expectations with respect to F_s on both sides of the above equation yields

$$E^Q\left[\frac{C_t}{B_t}\middle| F_s\right] = E^Q\left[E^Q\left[\frac{C_T(S_T)}{B_T}\middle| F_t\right]\middle| F_s\right] \qquad \forall t > s$$

Now using the law of iterated expectations we get

$$E^Q\left[\frac{C_t}{B_t}\middle| F_s\right] = E^Q\left[\frac{C_T(S_T)}{B_T}\middle| F_s\right] = \frac{C_s}{B_s} = C_s^B \qquad \forall t > s$$

which gives the desired result. ■

Lemma 5 *Under the EMM defined by (2.2), the discounted gain process is a martingale.*

Proof. Under the EMM Q, the process $\left\{S_t^B\right\}_{t=1}^T$ is a Q martingale. Using a standard property of martingales the process defined as $SS_t^B = \sum_{i=0}^{t-1} \eta_i(S_{i+1}^B - S_i^B)$ then is a Q martingale, since the investment strategy η_t is included in the information set.[13] Furthermore, from Lemma 3 we get that $\left\{C_t^B\right\}_{t=1}^T$ is also a Q martingale. Then using the fact that δ_t is an F_t predetermined process and using the same martingale property as above we get that the process $CC_t^B = \sum_{i=0}^{t-1} \delta_i(C_{i+1}^B - C_i^B)$ is a Q martingale. Then since from Lemma 2 the discounted gain process $\left\{G_t^B\right\}_{t=1}^T$ is the sum of two Q martingales, SS_t^B and CC_t^B, it is itself a Q martingale. ■

At this stage, we have all the ingredients to show the following main result.

Proposition 4 *If we have an EMM that makes the discounted price of the stock a martingale, then defining the price of any contingent claim as the expected value of its payoff, taken under this EMM and discounted at the riskless interest rate, constitutes a no-arbitrage price.*

Proof. From Lemma 4 G_t^B is a Q martingale. Because we are considering self financing strategies we get that V_t^B is a martingale. We prove the absence of arbitrage by contradiction. If we assume the existence of an arbitrage opportunity, then there exists a self financing strategy with type 1 arbitrage ($V_0 < 0, V_T \geq 0$ a.s.) or type 2 arbitrage ($V_0 \leq 0, V_T > 0$ a.s.). Both cases lead to a clear contradiction. Consider type 1 arbitrage: we start from the existence of a self financing strategy with $V_0 < 0$ that ends up with a positive final value. $V_0 < 0$ implies that $V_0^B < 0$ since the numeraire is always positive by definition. Also since $V_T \geq 0$ we have $V_T^B \geq 0$. Taking expectations and using the fact that V_t^B is a Q martingale yields $V_0^B = E_0^Q[V_T^B] \geq 0$. This is a contradiction because we assumed that we start with a negative value $V_0 < 0$. A similar argument works for type 2 arbitrage. Thus, the C_t from the EMM Q must be a no-arbitrage price. ■

In summary, we have demonstrated that in a discrete-time infinite state space setting, if we have an EMM that makes the underlying asset price a martingale, then the expected

[13] Note that because we are working in discrete time there is no need to investigate the integrability of SS_t^B.

value of the payoff of the contingent claim taken under this EMM, discounted at the riskless asset, is a no-arbitrage price. In Section 2.2, we derived such an EMM. Altogether, we have therefore demonstrated that for any contingent claim paying a final payoff $C_T(S_T)$ the current price C_t can be computed as

$$C_t = E^Q \left[\frac{C_T(S_T)}{B_T} B_t \middle| F_t \right].$$

4 Important special cases

In this section we demonstrate how a number of existing models are nested in our setup. We also give an example of a model that has not yet been discussed in the literature but can be handled by our setup.

4.1 Flexible risk premium specifications

One of the advantages of our approach is that we can allow for time-varying risk premia. Here we discuss some potentially interesting ways to specify the risk premium in the return process for the underlying asset. In order to demonstrate the link with the available literature and for computational simplicity, we assume conditional normal returns, although this assumption is by no means necessary.

The conditional normal models in the Duan (1995) and Heston and Nandi (2000) models are special cases of our set-up. In our notation, Duan (1995) assumes

$$r_t = r, \text{ and } \mu_t = r + \lambda\sigma_t$$

which in our framework corresponds to a Radon-Nikodym derivative of

$$\left.\frac{dQ}{dP}\right| F_t = \exp\left(-\sum_{i=1}^{t} \left(\frac{\varepsilon_i}{\sigma_i}\lambda - \frac{1}{2}\lambda^2 \right) \right)$$

and risk neutral innovations of the form

$$\varepsilon_t^* = \varepsilon_t + \lambda\sigma_t.$$

Heston and Nandi (2000) instead assume

$$r_t = r, \text{ and } \mu_t = r + \lambda\sigma_t^2 + \frac{1}{2}\sigma_t^2$$

which in our framework corresponds to a Radon-Nikodym derivative of

$$\left.\frac{dQ}{dP}\right| F_t = \exp\left(-\sum_{i=1}^{t} \left(\left(\lambda + \frac{1}{2}\right)\varepsilon_i - \frac{1}{2}\left(\lambda + \frac{1}{2}\right)^2 \sigma_i^2 \right) \right)$$

and risk neutral innovations of the form

$$\varepsilon_t^* = \varepsilon_t + \lambda\sigma_t^2 + \frac{1}{2}\sigma_t^2.$$

However, many empirically relevant cases are not covered by existing theoretical results. For example, in the original ARCH-M paper, Engle, Lilien and Robins (1987) find the strongest empirical support for a risk premium specification of the form

$$\mu_t = r_t + \lambda \ln(\sigma_t) + \tfrac{1}{2}\sigma_t^2$$

which cannot be used for option valuation using the available theory. In our framework it simply corresponds to a Radon-Nikodym derivative of

$$\left.\frac{dQ}{dP}\right| F_t = \exp\left(-\sum_{i=1}^{t}\left(\frac{\lambda \ln(\sigma_i) + \frac{1}{2}\sigma_i^2}{\sigma_i^2}\varepsilon_i - \frac{1}{2}\left(\frac{\lambda \ln(\sigma_i) + \frac{1}{2}\sigma_i^2}{\sigma_i^2}\right)^2 \sigma_i^2\right)\right)$$

and risk neutral innovations

$$\varepsilon_t^* = \varepsilon_t + \lambda \ln(\sigma_t) + \tfrac{1}{2}\sigma_t^2$$

Our approach allows for option valuation under such specifications whereas the existing literature does not.

4.2 Conditionally inverse Gaussian returns

Christoffersen, Heston and Jacobs (2006) analyze a GARCH model with an inverse Gaussian innovation, $y_t \sim IG(\sigma_t^2/\eta^2)$. We can write their return dynamic as

$$R_t = r + \left(\zeta + \eta^{-1}\right)\sigma_t^2 + \varepsilon_t \tag{4.1}$$

where ε_t is a zero-mean innovation defined by

$$\varepsilon_t = \eta y_t - \eta^{-1}\sigma_t^2$$

and where the conditional return variance, σ_t^2, is of the GARCH form.

From the MGF of an inverse Gaussian variable, we can derive the conditional log MGF of ε_t as

$$\Psi_t(u) = \left(u + \frac{1 - \sqrt{1 + 2u\eta}}{\eta}\right)\frac{\sigma_t^2}{\eta}$$

The EMM condition

$$\Psi_t(\nu_t - 1) - \Psi_t(\nu_t) - \Psi_t(-1) + \alpha_t \sigma_t^2 = 0$$

is now solved by the constant

$$\nu_t = \nu = \frac{1}{2\eta}\left[\frac{\left(2 + \zeta\eta^3\right)^2}{4\zeta^2\eta^2} - 1\right], \forall t$$

which in turn implies that the EMM is given by

$$\begin{aligned}
\left.\frac{dQ}{dP}\right| F_t &= \exp\left(-\sum_{i=1}^{t}\left(\nu\varepsilon_i + \left(\nu + \frac{1 - \sqrt{1 + 2\nu\eta}}{\eta}\right)\frac{\sigma_i^2}{\eta}\right)\right) \\
&= \exp\left(-\nu t \overline{\varepsilon}_t - \delta t \overline{\sigma_t^2}\right)
\end{aligned}$$

where $\bar{\varepsilon}_t = \frac{1}{t}\sum_{i=1}^{t}\varepsilon_i$, $\overline{\sigma_t^2} = \frac{1}{t}\sum_{i=1}^{t}\sigma_i^2$, and $\delta = \frac{\nu}{\eta} + \frac{1-\sqrt{1+2\nu\eta}}{\eta^2}$.

These expressions can be used to obtain the risk-neutral distribution from Christoffersen, Heston and Jacobs (2006) using the results in Section 2. Recall that in general the risk neutral log MGF is

$$\Psi_t^{Q*}(u) = -u\Psi_t'(\nu) + \Psi_t(\nu+u) - \Psi_t(\nu)$$

In the GARCH-IG case we can write

$$\Psi_t^{Q*}(u) = \left(u + \frac{1-\sqrt{1+2u\eta^*}}{\eta^*}\right)\frac{\sigma_t^{*2}}{\eta^*}$$

where

$$\begin{aligned}
\eta^* &= \frac{\eta}{1+2\nu\eta} \text{ and} \\
\sigma_t^{*2} &= \frac{\sigma_t^2}{(1+2\nu\eta)^{3/2}}
\end{aligned}$$

So that the risk-neutral variance in general will be different from the physical variance.

The risk neutral return model can be written as

$$R_t \equiv \ln\left(\frac{S_t}{S_{t-1}}\right) = r - \Psi_t^{Q*}(-1) + \varepsilon_t^* = r + \left(\zeta^* + \eta^{*-1}\right)\sigma_t^{*2} + \varepsilon_t^*$$

where

$$\begin{aligned}
\zeta^* &= \frac{1-2\eta^* - \sqrt{1-2\eta^*}}{\eta^{*2}} \text{ and} \\
\varepsilon_t^* &= \eta^* y_t^* - \eta^{*-1}\sigma_t^{*2}
\end{aligned}$$

The risk neutral process thus takes the same form as the physical process which is exactly what our Proposition 3 in Section 2.4 shows.

4.3 Conditionally Poisson-normal jumps

Another interesting model that can be nested in our framework is the heteroskedastic model with Poisson-normal innovations in Duan, Ritchken and Sun (2005).[14] For expositional simplicity, we consider the simplest version of the model. More complex models, for instance with time-varying Poisson intensities, can also be accommodated. We can write the underlying asset return as

$$R_t = \kappa_t + \varepsilon_t$$

The zero-mean innovation ε_t equals

$$\varepsilon_t = \sigma_t\left(J_t - \lambda\bar{\mu}\right)$$

where J_t is a Poisson jump process with N_t jumps each with distribution $N\left(\bar{\mu}, \bar{\gamma}^2\right)$ and jump intensity λ. The conditional return variance equals $\left(1+\lambda\left(\bar{\mu}^2+\bar{\gamma}^2\right)\right)\sigma_t^2$, where σ_t^2 is of the

[14] Maheu and McCurdy (2004) consider a different discrete-time jump model.

GARCH form. The log return mean κ_t is a function of σ_t^2 as well as the jump and risk premium parameters.

We can derive the conditional log MGF of ε_t as

$$\begin{aligned}\Psi_t(u) &= \ln(E_{t-1}[\exp(-u\sigma_t(J_t-\lambda\overline{\mu}))]) \\ &= u\lambda\overline{\mu}\sigma_t + \frac{1}{2}u^2\sigma_t^2 + \lambda\left[\exp\left(-\overline{\mu}u\sigma_t + \frac{1}{2}\overline{\gamma}^2u^2\sigma_t^2\right)-1\right]\end{aligned}$$

The approach taken in Duan et al (2005) corresponds to fixing $\nu_t = \nu$ and setting

$$\kappa_t = r + \Psi_t(\nu) - \Psi_t(\nu-1)$$

which in turn implies that the EMM is given by

$$\left.\frac{dQ}{dP}\right| F_t = \exp\left(-\nu t\overline{\varepsilon_t} - \nu\lambda\overline{\mu}t\overline{\sigma_t} - \frac{1}{2}t\nu^2\overline{\sigma_t^2} + \lambda t - \lambda\sum_{i=1}^{t}\exp\left(-\overline{\mu}\nu_i\sigma_i + \frac{1}{2}\overline{\gamma}^2\nu_i^2\sigma_i^2\right)\right)$$

where $\overline{\varepsilon_t}$ and $\overline{\sigma_t^2}$ are the historical averages as above.

We can again show that the risk-neutral distribution of the risk neutral innovation is from the same family as the physical

$$\begin{aligned}\Psi_t^{Q*}(u) &= \ln E_{t-1}^Q[\exp(-u\varepsilon_t^*)] \\ &= u\lambda_t^*\overline{\mu}_t^*\sigma_t + \frac{1}{2}u^2\sigma_t^2 + \lambda_t^*\left[\exp\left(-\overline{\mu}_t^*u\sigma_t + \frac{1}{2}\overline{\gamma}^2u^2\sigma_t^2\right)-1\right]\end{aligned}$$

where

$$\lambda_t^* = \lambda\exp\left(-\overline{\mu}\nu\sigma_t + \frac{1}{2}\overline{\gamma}^2\nu^2\sigma_t^2\right) \text{ and } \overline{\mu}_t^* = \overline{\mu} - \overline{\gamma}^2\sigma_t\nu$$

Note that in this model the mapping between the risk-neutral and physical returns is

$$\varepsilon_t^* = \varepsilon_t + \Psi_t'(\nu) = \varepsilon_t + \sigma_t(\lambda\bar{\mu} - \lambda_t^*\overline{\mu}_t^*)$$

and the mapping between the physical and risk-neutral conditional variance is

$$\sigma_t^{*2} = \sigma_t^2 + \lambda_t^*\sigma_t^2\left(\overline{\gamma}^2 + \overline{\mu}_t^{*2}\right)$$

4.4 Conditionally skewed variance gamma returns

We now introduce a new model where the conditional skewness, s, and excess kurtosis, k, are given directly by two parameters in the model.[15] Consider the return of the underlying asset specified as follows

$$\begin{aligned}R_t &= \mu_t - \gamma_t + \varepsilon_t \\ &= \mu_t - \gamma_t + \sigma_t z_t, \qquad\qquad z_t \overset{i.i.d}{\sim} SVG(0,1,s,k)\end{aligned}$$

[15] In Christoffersen, Heston and Jacobs (2006), conditional skewness and kurtosis are driven by functions of the same parameter.

The distribution of the shocks, $SVG(0,1,s,k)$, is a standardized skewed variance gamma distribution which is constructed as a mixture of two gamma variables.[16] The conditional variance, σ_t^2, can take on any GARCH specification. We will provide an empirical illustration in the next section using a leading GARCH dynamic.

Let z_1 and z_2 be independent draws from two gamma distributions

$$z_{i,t} \overset{i.i.d}{\sim} \Gamma\left(4/\tau_i^2\right), \qquad i=1,2$$

parameterized as

$$\tau_1 = \sqrt{2}\left(s - \sqrt{\frac{2}{3}k - s^2}\right) \text{ and } \tau_2 = \sqrt{2}\left(s + \sqrt{\frac{2}{3}k - s^2}\right)$$

If we construct the SVG random variable from the two gamma variables as

$$z_t = \frac{1}{2\sqrt{2}}\left(\tau_1 z_{1,t} + \tau_2 z_{2,t}\right) - \sqrt{2}\left(\frac{1}{\tau_1} + \frac{1}{\tau_2}\right)$$

then z_t will have a mean of zero, a variance of one, a skewness of s, and an excess kurtosis of k, thus allowing for conditional skewness and kurtosis in the GARCH model, as intended.[17]

The log moment generating function of ε_t can be derived from the gamma distribution MGF as

$$\Psi_t(u) = \sqrt{2}\left(\tau_1^{-1} + \tau_2^{-1}\right)u\sigma_t - 4\tau_1^{-2}\ln\left(1 + \frac{1}{2\sqrt{2}}\tau_1 u\sigma_t\right) - 4\tau_2^{-2}\ln\left(1 + \frac{1}{2\sqrt{2}}\tau_2 u\sigma_t\right)$$

so that the mean correction variable, γ_t, for the return can be found as $\gamma_t = \Psi_t(-1)$.

Using the formula for the risk neutral conditional log MGF

$$\Psi_t^{Q*}(u) = -u\Psi_t'(\nu_t) + \Psi_t(\nu_t + u) - \Psi_t(\nu_t)$$

we can show that the risk neutral model is

$$R_t = r_f - \gamma_t^* + \varepsilon_t^* \tag{4.2}$$

where

$$\Psi_t^{Q*}(u) = \sqrt{2}\left(\tau_1^{-1}\sigma_{1,t}^* + \tau_2^{-1}\sigma_{2,t}^*\right)u - 4\tau_1^{-2}\ln\left(1 + \frac{1}{2\sqrt{2}}\tau_1\sigma_{1,t}^* u\right) - 4\tau_2^{-2}\ln\left(1 + \frac{1}{2\sqrt{2}}\tau_2\sigma_{2,t}^* u\right)$$

with

$$\sigma_{i,t}^* = \frac{\sigma_t}{\sqrt{2} + \frac{1}{2}\tau_i\sigma_t\nu_t}, \qquad \text{for } i=1,2. \tag{4.3}$$

We see that $\Psi_t^{Q*}(u)$ is exactly of the same form as $\Psi_t(u)$, and therefore that $\gamma_t^* = \Psi_t^{Q*}(-1)$. This model will be investigated empirically below.

[16] See Madan and Seneta (1990) for an early application of the symmetric and i.i.d. variance gamma distribution in finance.

[17] The special cases where τ_1 or τ_2 are zero can be handled easily by drawing from the normal distribution for the relevant mixing variable $z_{1,t}$ or $z_{2,t}$. When both τ_1 and τ_2 are zero then the normal distribution obtains for z_t.

5 Empirical illustration

In this section we demonstrate how the greater flexibility and generality allowed for by our approach can lead to more realistic option valuation models. To do so, we analyze the GARCH-SVG model in Section 4.4, which allows for conditional skewness and kurtosis, and which has not yet been analyzed in the literature. We compare its empirical implications with the more standard conditional normal model. We compute option prices from both models using parameters estimated from return data, and subsequently construct option implied volatility smiles. We also compare the two heteroskedastic models to two benchmark models with independent returns.

5.1 Parameter estimates from index returns and stylized facts

We start by illustrating some key stylized facts of daily equity index returns using the S&P500 as a running example.

Figure 1 shows a normal quantile-quantile plot (QQ plot) of daily S&P500 returns, using data from January 2, 1980 through December 30, 2005 for a total of 6,564 observations. The returns are standardized by the sample mean and standard deviation. The data quantiles on the vertical axis are plotted against the normal distribution quantiles on the horizontal axis. The plot reveals the well-known stark deviations from normality in daily asset returns: actual returns include much more extreme observations than the normal distribution allows for in a sample of this size. The largest negative return is the famous 20 standard deviation crash in October 1987, but the normal distribution has trouble fitting a large number of extremes in both tails of the return distribution. The actual returns range from -20 to +9 standard deviations but the normal distribution only ranges from -4 to +4 standard deviations in a sample of this size.

Figure 2 shows the sample autocorrelation function of the squared daily returns for the sample. The significantly positive correlations at short lags suggest the need for a dynamic volatility model allowing for clustering in volatility.

Figures 1 and 2 clearly suggest the need for a GARCH model which can capture potentially both the volatility clustering in Figure 2 and the non-normality in Figure 1.

As a benchmark, we use the conditional normal NGARCH model of Engle and Ng (1993)

$$R_t = \mu_t - \gamma_t + \sigma_t z_t, \qquad z_t \overset{i.i.d}{\sim} N(0,1) \tag{5.1}$$

where

$$\begin{aligned} \mu_t &= r_t + \lambda\sigma_t \\ \gamma_t &= \tfrac{1}{2}\sigma_t^2 \\ \sigma_t^2 &= \beta_0 + \beta_1\sigma_{t-1}^2 + \beta_2\sigma_{t-1}^2\left(z_{t-1} - \beta_3\right)^2 \end{aligned}$$

Notice that the β_3 parameter in the GARCH variance specification allows for an asymmetric variance response to positive versus negative shocks, z_{t-1}. This captures the so-called leverage effect, which is another important empirical regularity in daily equity index returns.

Table 1 reports the maximum likelihood estimates of the GARCH parameters. We also report parameter estimates for a version of the model where the GARCH dynamics have been shut down, that is, where $\beta_1 = \beta_2 = \beta_3 = 0$. Notice the large increase in the Log-likelihood function from including the GARCH dynamics.

Figure 3 shows the autocorrelation function for the observed squared GARCH shocks, z_t^2. If the GARCH model has adequately captured the volatility clustering then the shocks should be independent and in particular the squared shocks should be uncorrelated. Figure 3 suggests that the GARCH model does a good job of capturing the volatility dynamics in the daily index returns.

Figure 4 assesses the conditional normality assumption by plotting a QQ plot of z_t against the normal distribution. It is clear from Figure 4 that much of the non-normality in the raw returns has been removed by the GARCH model. This is particularly true for the right tail, where the non-normality was least pronounced to begin with. Unfortunately, the left tail of the shock distribution still exhibits strong evidence of non-normality with negative shocks as large as -10 standard deviations compared with the normal distribution's -4.

From Figures 3 and 4, we conclude that while the normal GARCH model appears to provide adequate dynamics for capturing volatility clustering, the conditional normality assumption is violated in the data and must be modified in the model.

For the implementation of the GARCH-SVG model, μ_t and σ_t^2 are the same as in the conditional normal model in (5.1). We can calibrate the s and k parameters in the GARCH-SVG model from Section 4.4 by simply equating them to the sample moments from the z_t sequence from the QMLE estimation of the GARCH model. These sample moments are reported in Table 1.

Figure 5 shows the QQ plot of the GARCH shocks against the SVG distribution. Compared with the normal QQ plot in Figure 4, we see that the SVG captures the left tail of the shock distribution much better than the normal does. Impressively, the SVG model only has trouble fitting the two most extreme negative shocks, whereas the normal distribution misses a whole string of large negative shocks.

5.2 Option prices and implied volatilities

Armed with estimated return processes we are ready to assess the option pricing implications of the different models. From Section 3 we have the general option price relationship which for a European call option with strike price K is

$$C_t(T, K) = E_t^Q \left[Max(S_T - K, 0) \frac{B_t}{B_T} \right]$$

Using the estimated physical process from Table 1 we can now simulate future paths for S_T from the current S_t and compute the option price as the simulated sample analogue to this discounted expectation.

We present evidence on the option pricing properties of the various models in Figures 6 and 7. Figure 6 considers an i.i.d. normal and an i.i.d. SVG model where the GARCH dynamics have been shut down ($\beta_1 = \beta_2 = \beta_3 = 0$), and s and k have been set to the the sample skewness and kurtosis from the raw returns which are reported in Table 1. Figure 7 considers the normal GARCH-Normal and GARCH-SVG models. The parameter estimates used are again from Table 1.

We first compute option prices for various moneyness and maturities and we then compute implied Black and Scholes (1973) volatilities from the model option prices. Implied volatilities are plotted against moneyness on the horizontal axis. The three panels correspond to maturities of 1 day, 1 week, and 1 month respectively.

The i.i.d. SVG model in Figure 6 (solid lines) shows a strong implied volatility "smile" for the 1-day maturity driven by the large excess kurtosis of 27.33 from Table 1. Interestingly, as the maturity increases the smile becomes an asymmetric "smirk" driven by the skewness parameter of -1.21 in Table 1. The i.i.d. normal model in Figure 6 (dashed line) results in a flat implied volatility curve.

The GARCH-SVG model in Figure 7 shows a smirk at the 1 day maturity compared with the flat implied volatility for the GARCH-Normal model where the conditional 1 day distribution is normal. The GARCH-Normal model generates a non-trivial volatility smirk for horizons beyond 1 day where the conditional distribution is no longer normal. However, the GARCH-SVG model is capable of capturing much more non-normality than the GARCH-Normal model at all horizons. This is important because the empirical option valuation literature often finds that existing models are unable to fit short term option prices where the implied degree of non-normality is large.[18]

From this empirical illustration we conclude that it is possible to build relatively simple models capturing the conditional volatility and non-normality found in index returns data, and more importantly that such models provide the flexibility needed to price options.

6 Some continuous time limits

In order to anchor our work in the continuous time literature we now explore the links between some of the discrete time models we have analyzed above and standard continuous-time models. We study three important cases: a homoskedastic model with normal innovations, a homoskedastic model with non-normal (inverse Gaussian) innovations, and a heteroskedastic model with normal innovations.

6.1 Homoskedastic normal returns

Consider the homoskedastic i.i.d. normal model for discrete time interval Δ,

$$R_t = \ln(S_t) - \ln(S_{t-\Delta}) = \mu\Delta - \tfrac{1}{2}\sigma^2\Delta + \sigma\sqrt{\Delta}z_t \qquad z_t|F_{t-1} \sim N(0,1) \tag{6.1}$$

and for simplicity also consider a constant risk-free rate.

The EMM condition (2.3) is solved by choosing a constant $\nu = (\mu - r)/\sigma^2$, and the discrete-time risk-neutral dynamic is given by

$$\ln(S_t) - \ln(S_{t-\Delta}) = r\Delta - \tfrac{1}{2}\sigma^2\Delta + \sigma\sqrt{\Delta}z_t^* \qquad z_t^*|F_{t-1} \sim N(0,1) \tag{6.2}$$

The continuous-time limit of the risk-neutral process is given by

$$d(\ln(S_t)) = \left(r - \tfrac{1}{2}\sigma^2\right)dt + \sigma dz^*(t)$$

where $z^*(t)$ is a Wiener process under Q. This is the risk-neutral process in the Black-Scholes-Merton (BSM) model. In the diffusion limit the options are thus priced using the BSM formula.

[18] See Bates (2003) for an excellent discussion of this and other stylized facts in option markets.

Consider a European option with strike price K with $T-t=M\Delta$ days to maturity. The call price can be written as

$$C_{\Delta,t}=e^{-rM\Delta}S_tE_t^Q\left[e^{R_{t,M}}I[R_{t,M}>\ln(K/S_t)]\right]-e^{-rM\Delta}KP_t^Q[R_{t,M}>\ln(K/S_t)]$$

where $R_{t,M}=\ln\left(S_{t+M\Delta}\right)-\ln\left(S_t\right)$ and where $I\left[*\right]$ is the indicator function. Using the i.i.d. normal risk-neutral process in (6.2) we can rewrite the call price as

$$C_{\Delta,t}=e^{-rM\Delta}P_{1,t,\Delta}-e^{-rM\Delta}KP_{2,t,\Delta}$$

where

$$P_{1,t,\Delta}=e^{rM\Delta}\Phi\left(\frac{\ln(S_t/K)+\left(r+\frac{1}{2}\sigma^2\right)\Delta M}{\sigma\sqrt{\Delta M}}\right),\qquad P_{2,t,\Delta}=\Phi\left(\frac{\ln(S_t/K)+\left(r-\frac{1}{2}\sigma^2\right)\Delta M}{\sigma\sqrt{\Delta M}}\right)$$

where Φ is the c.d.f. of the standard normal distribution.

Note therefore that the i.i.d. normal discrete time process and its parameterization in (6.1), and the choice of Radon-Nikodym derivative (and thus EMM) gives the BSM price for any Δ so long as r and σ are provided in the proper daily log-return units.

6.2 Homoskedastic Inverse Gaussian returns

Consider now a homoskedastic version of the Inverse Gaussian (IG) model in (4.1) written for a discrete time interval Δ,

$$\begin{aligned}
R_t &= r\Delta+\left(\zeta(\Delta)+\eta(\Delta)^{-1}\right)\sigma^2(\Delta)+\varepsilon_t && (6.3)\\
\varepsilon_t &= \eta(\Delta)y_t-\eta(\Delta)^{-1}\sigma^2(\Delta) && (6.4)\\
y_t &\sim IG\left(\frac{\sigma^2(\Delta)}{\eta(\Delta)}\right) && (6.5)
\end{aligned}$$

As shown above for the heteroskedastic case, the risk neutral return distribution is in the same family as the historical model and can be written as follows

$$\begin{aligned}
R_t &= r\Delta+\left(\zeta^*(\Delta)+\eta^*(\Delta)^{-1}\right)\sigma^{*2}(\Delta)+\varepsilon_t^* && (6.6)\\
\varepsilon_t^* &= \eta^*(\Delta)y_t^*-\eta^*(\Delta)^{-1}\sigma^{*2}(\Delta) && (6.7)\\
y_t^* &\overset{Q}{\sim} IG\left(\frac{\sigma^{*2}(\Delta)}{\eta^*(\Delta)}\right) && (6.8)
\end{aligned}$$

where

$$\begin{aligned}
\eta^*(\Delta) &= \frac{\eta(\Delta)}{1+2\nu(\Delta)\eta(\Delta)}\\
\sigma^{*2}(\Delta) &= \frac{\sigma^2(\Delta)}{(1+2\nu(\Delta)\eta(\Delta))^{3/2}}\\
\zeta^*(\Delta) &= \frac{1-2\eta^*(\Delta)-\sqrt{1-2\eta^*(\Delta)}}{\eta^*(\Delta)^2}
\end{aligned}$$

and where $\nu(\Delta)$ solves (2.3) and is given by

$$\nu(\Delta) = \frac{1}{2\eta(\Delta)}\left[\frac{\left(2+\zeta(\Delta)^2\eta(\Delta)^3\right)^2}{4\zeta(\Delta)^2\eta(\Delta)^2} - 1\right]$$

Consider a European option with strike price K with $T-t = M\Delta$ days to maturity. The call price can be written as

$$C_{\Delta,t} = e^{-rM\Delta}P_{1,t,\Delta} - e^{-rM\Delta}KP_{2,t,\Delta} \tag{6.9}$$

where using Fourier inversion of the risk-neutral log MGF of $R_{t,M}$, namely $\Psi^{Q*}_{t,M}(u)$, we have

$$\begin{aligned} P_{1,t,\Delta} &= \frac{e^{rM\Delta}}{2} + \int_0^{+\infty} \mathrm{Re}\left[\frac{\exp\left(\Psi^{Q*}_{t,M}(-1-iu) - iu\ln\left(\frac{K}{S_t}\right)\right)}{i\pi u}\right] du \\ P_{2,t,\Delta} &= \frac{1}{2} + \int_0^{+\infty} \mathrm{Re}\left[\frac{\exp\left(-iu\ln\left(\frac{K}{S_t}\right) + \Psi^{Q*}_{t,M}(-iu)\right)}{i\pi u}\right] du \end{aligned}$$

where we have

$$\Psi^{Q*}_{t,M}(u) \equiv \ln\left(E^{Q}_{t}\left[\exp\left(uR_{t,M}\right)\right]\right) = -\left(r\Delta + \zeta(\Delta)\sigma^2(\Delta)\right)Mu + \frac{\left(1-\sqrt{1+2u\eta^*(\Delta)}\right)\sigma^{*2}(\Delta)M}{\eta^*(\Delta)^2}$$

Christoffersen, Heston and Jacobs (2006) show that in the heteroskedastic case, the stochastic volatility model in Heston (1993a) with perfectly correlated shocks can be obtained as a limit of the IG-GARCH model when discrete time IG-GARCH process converges to the continuous time stochastic volatility Heston (1993a) process (see below) with perfectly correlated shocks when Δ and $\eta(\Delta)$ go to zero. This limit obtains when using a particular parameterization for the IG-GARCH model and the parameterization $\zeta(\Delta) = \lambda - \eta(\Delta)^{-1}$ for the return mean. As the homoskedastic IG model is a special case of the IG-GARCH model it will converge to the homoskedastic Heston (1993) process which is simply the geometric Brownian motion underlying the Black-Scholes model. The continuous-time limit of the risk-neutral process is thus again given by

$$d(\ln(S_t)) = \left(r - \tfrac{1}{2}\sigma^2\right)dt + \sigma dz^*(t)$$

Figure 8 illustrates the convergence of the homoskedastic IG option price in (6.9) to the BSM price when Δ goes to zero. In the figure we let $\eta(\Delta) = \eta\Delta$, $\sigma^2(\Delta) = \sigma^2\Delta$, and set $\lambda = 1$, $\eta = -0.001$ per day, and $\sigma = .10.$ per year. Note that skewness is $3\eta/\sigma$ in the IG model.

Christoffersen, Heston and Jacobs (2004) also show that an alternative pure jump limit can be obtained in the Inverse Gaussian model.

6.3 Heteroskedastic normal returns

Consider the Heston and Nandi (2000) model

$$\begin{aligned} R_t &= r\Delta + \lambda\sigma_t^2 + \sigma_t z_t \\ \sigma_{t+\Delta}^2 &= \omega + \beta\sigma_t^2 + \alpha(z_t - \gamma\sigma_t)^2 \end{aligned} \tag{6.10}$$

Defining $v_{t+\Delta} = \sigma_{t+\Delta}^2/\Delta$, we have

$$v_{t+\Delta} = \omega_v + \beta v_t + \alpha_v\left(z_t - \gamma_v\sqrt{v_t}\right)^2 \tag{6.11}$$

with $\omega_v = \omega/\Delta$, $\alpha_v = \alpha/\Delta$ and $\gamma_v = \gamma\sqrt{\Delta}$.

Note that the conditional correlation

$$Corr_{t-\Delta}\left(v_{t+\Delta}, R_t\right) = -\frac{sign(\gamma_v)\sqrt{2\gamma_v^2 v_t}}{\sqrt{1+2\gamma_v^2 v_t}}$$

so that the correlation goes to plus or minus one when the interval shrinks to zero. Using the parametrization $\alpha(\Delta) = \frac{1}{4}\sigma^2\Delta^2$, $\beta(\Delta) = 0$, $\omega(\Delta) = (\kappa\theta - \frac{1}{4}\sigma^2)\Delta^2$, and $\gamma(\Delta) = \frac{2}{\sigma\Delta} - \frac{\kappa}{\sigma}$, and following Foster and Nelson (1994), Heston and Nandi derive the diffusion limit for the physical process

$$\begin{aligned} d\ln(S_t) &= (r + \lambda v)dt + \sqrt{v}dz \\ dv &= \kappa(\theta - v)dt + \sigma\sqrt{v}dz \end{aligned} \tag{6.12}$$

which corresponds to a special case of the stochastic volatility model in Heston (1993) with perfectly correlated shocks to stock price and volatility.

The Heston-Nandi discrete time option price is

$$C_{t,\Delta} = S_t P_{1,t,\Delta} - e^{-rM\Delta} K P_{2,t,\Delta}$$

where the formulas for $P_{1,t,\Delta}$ and $P_{2,t,\Delta}$, which rely on Fourier inversion, are provided in Heston and Nandi (2000).

Figure 9 shows the convergence of the Heston and Nandi (2000) discrete time GARCH option price to the continuous time SV option price in Heston (1993). We use $r = 0$, $K = 100$, $S = 100$, $M\Delta = 180$, $\sigma = 0.1$, $V = 0.01$, $\kappa = 2$, $\theta = .01$, and shock correlation $\rho = -1$.

Note that markets are complete in the limiting case with $\rho = -1$ because there is only one source of uncertainty. Below we analyze the more general case of a discrete-time two-shock stochastic volatility model and its continuous time limit where $-1 < \rho < 1$ and so where markets are incomplete even in continuous time.

Figure 9 shows that convergence is fast suggesting that the added incompleteness arising from discrete time is minimal. By comparison convergence is slower in Figure 8 because of the conditional skewness in the discrete time process which is not present in the limiting process considered here.

7 Extensions and Discussion

In this section we first develop a discrete-time two-shock stochastic volatility model and derive its continuous time limit. Subsequently we compare the risk neutralization for this model with risk neutralization in the continuous time SV model. We discuss the GARCH model as a special case of this approach. We also discuss the issue of market incompleteness and the resulting non-uniqueness of option prices, again by discussing similarities and differences between continuous and discrete time.

7.1 A discrete-time stochastic volatility model

Popular continuous-time stochastic volatility models such as Heston (1993a) contain two (correlated) innovations, whereas the GARCH processes in this paper contain a single innovation. Nelson (1991) and Duan (1997) derive a continuous-time two-innovation stochastic volatility model as the limit of a GARCH model, but as noted by Corradi (2000) for instance, a given discrete-time model can have several continuous-time limits and vice versa.[19] As shown above, Heston and Nandi (2000) derive a limit to their proposed GARCH process that contains two perfectly correlated shocks. This limit amounts to a one-shock process, and is therefore intuitively similar to a GARCH process.

With this in mind, we now analyze the limits of a class of discrete-time stochastic volatility processes, which contain two (potentially correlated) shocks.[20] We derive the continuous-time limits for these processes, and then analyze the GARCH limit as a special case.

Consider the return and volatility dynamics

$$\begin{aligned} R_t &= \ln(S_t/S_{t-1}) = \mu_t + \sigma_t z_{1,t} \\ \sigma^2_{t+1} &= f(\sigma^2_t, z_{2,t}) \end{aligned}$$

where

$$z_t \equiv (z_{1,t}, z_{2,t})' \sim N\left(\left((0,0)', \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right)\right)$$

The log MGF is given by

$$\Psi_t(u_1, u_2) = \log[E_{t-1}(\exp(-u_1 z_{1,t} - u_2 z_{2,t}))] = \frac{1}{2}\left[(u_1 + \rho u_2)^2 + \left(1 - \rho^2\right) u_2^2\right]$$

By analogy with the one-shock case, we define the following Radon-Nikodym derivative

$$\frac{dQ}{dP}|F_t = \exp\left(-\sum_{i=1}^{t}\left(\nu_{1,i} z_{1,i} + \nu_{2,i} z_{2,i} + \Psi_i(\nu_{1,i}, \nu_{2,i})\right)\right)$$

The probability measure Q defined by the Radon-Nikodym derivative is an EMM if and only if

$$\Psi_t(\nu_{1,t} - \sigma_t \nu_{2,t}) - \Psi_t(\nu_{1,t}, \nu_{2,t}) + \mu_t - r = \frac{1}{2}\sigma_t^2 - (\nu_{1,t} + \rho\nu_{2,t})\sigma_t + \mu_t - r = 0 \qquad (7.1)$$

[19] See also Nelson and Foster (1994), Foster and Nelson (1996), Nelson (1996) and Ritchken and Trevor (1999) for limit results.

[20] See Ghysels, Harvey and Renault (1995) for a review of discrete-time stochastic volatility models.

Note that this is one equation in two unknowns, namely $\nu_{1,t}$ and $\nu_{2,t}$. Thus the second shock provides a new source of non-uniqueness to be discussed further below.

The risk neutral log MGF is given by

$$\begin{aligned} E^Q_{t-1}[\exp(-u_1 z_{1,t} - u_2 z_{2,t})] &= E^P_{t-1}\left[\left(\frac{\frac{dQ}{dP}|F_t}{\frac{dQ}{dP}|F_{t-1}}\right)\exp(-u_1 z_{1,t} - u_2 z_{2,t})\right] \\ &= \exp(\Psi_t(u_1+\nu_{1,t}, u_2+\nu_{2,t}) - \Psi_t(\nu_{1,t}, \nu_{2,t})) \end{aligned}$$

where

$$z_t = (z_{1,t}, z_{2,t})' \overset{Q}{\sim} N\left(\left((-\nu_{1,t} - \rho\nu_{2,t}, -\nu_{1,t}\rho - \nu_{2,t})', \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right)\right) \tag{7.2}$$

We now illustrate this risk-neutralization for a specific parametric example

$$\begin{aligned} R_t &= r + \lambda\sigma_t^2 - \frac{1}{2}\sigma_t^2 + \sigma_t z_{1,t} \\ \sigma_{t+1}^2 &= \omega + \beta\sigma_t^2 + \alpha(z_{2,t} - \gamma\sigma_t)^2 \end{aligned} \tag{7.3}$$

The dynamic in (7.3) can be thought of as a stochastic volatility (two-shock) generalization of the GARCH dynamic in Heston and Nandi (2000).

According to (7.2) the risk-neutral model is given by

$$\begin{aligned} R_t &= r - \frac{1}{2}\sigma_t^2 + \sigma_t z_{1,t}^* \\ \sigma_{t+1}^2 &= \omega + \beta\sigma_t^2 + \alpha(z_{2,t}^* - \nu_{1,t}\rho - \nu_{2,t} - \gamma\sigma_t)^2 \end{aligned} \tag{7.4}$$

where

$$z_t^* = \begin{pmatrix} z_{1,t}^* = z_{1,t} + \nu_{1,t} + \rho\nu_{2,t} \\ z_{2,t}^* = z_{1,t} + \nu_{1,t}\rho + \nu_{2,t} \end{pmatrix} \overset{Q}{\sim} N\left(\left((0,0)', \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right)\right)$$

In the one-shock GARCH case above, we could simply solve (2.3) by choosing the scalar ν_t as a function of the GARCH parameters. Determining $\nu_{1,t}$ and $\nu_{2,t}$ in a model with two innovations is somewhat more complex, but the intuition underlying the procedure is critical to understanding the link with the continuous-time literature. From (7.1) and (7.3) we have $\nu_{1,t} + \nu_{2,t}\rho = \lambda\sigma_t$. We then note that if we want to preserve the affine structure in (7.4) we need $\nu_{2,t} = \nu_2\sigma_t$, which yields the risk neutral dynamic

$$\begin{aligned} R_t &= r - \frac{1}{2}\sigma_t^2 + \sigma_t z_{1,t}^* \\ \sigma_{t+1}^2 &= \omega + \beta\sigma_t^2 + \alpha(z_{2,t}^* - \gamma^*\sigma_t)^2 \end{aligned} \tag{7.5}$$

with $\gamma^* = \gamma + \nu_2(1-\rho^2) + \lambda\rho$.

Note that while λ, which is the price of equity risk, can be estimated from returns, ν_2, which arises from the new separate volatility shock is not identified from the return on the underlying asset only. It must be estimated from option prices as it only appears under the risk neutral measure. This is of course also the case in the continuous time SV models.

Using an approach similar to that taken in Heston and Nandi (2000) option valuation in this discrete-time SV model can be done via Fourier inversion of the conditional characteristic function.

7.2 A diffusion limit of the discrete-time stochastic volatility model

We first write the discrete-time stochastic volatility model as

$$
\begin{aligned}
R_t &= r\Delta + \lambda\sigma_t^2 - \frac{1}{2}\sigma_t^2 + \sigma_t z_{1,t} && (7.6)\\
\sigma_{t+\Delta}^2 &= \omega + \beta\sigma_t^2 + \alpha(z_{2,t} - \gamma\sigma_t)^2 && (7.7)
\end{aligned}
$$

Reparametrizing $v_{t+\Delta} = \sigma_{t+\Delta}^2/\Delta$, we have

$$
v_{t+\Delta} = \omega_v + \beta v_t + \alpha_v \left(z_{2,t} - \gamma_v\sqrt{v_t}\right)^2 \tag{7.8}
$$

with $\omega_v = \omega/\Delta$, $\alpha_v = \alpha/\Delta$ and $\gamma_v = \gamma\sqrt{\Delta}$.

By analogy with Heston and Nandi (2000) we use the parametrization $\alpha(\Delta) = \frac{1}{4}\sigma^2\Delta^2$, $\beta(\Delta) = 0$, $\omega(\Delta) = (\kappa\theta - \frac{1}{4}\sigma^2)\Delta^2$, and $\gamma(\Delta) = \frac{2}{\sigma\Delta} - \frac{\kappa}{\sigma}$. As $\Delta \to 0$ the dynamic in (7.6) and (7.8) converges to

$$
\begin{aligned}
d\ln(S_t) &= (r + \lambda v - \frac{1}{2}v)dt + \sqrt{v}dz_1 && (7.9)\\
dv &= \kappa(\theta - v)dt + \sigma\sqrt{v}dz_2
\end{aligned}
$$

where z_1 and z_2 are two Wiener process such that $dz_1dz_2 = \rho dt$. Note that the discrete time conditional correlation is given by

$$
corr_{t-\Delta}\left(v_{t+\Delta}, R_t\right) = -\frac{\rho sign(\gamma_v)\sqrt{2\gamma_v^2 v_t}}{\sqrt{1 + 2\gamma_v^2 v_t}}
$$

As $\Delta \to 0$, the variance asymmetry parameter $\gamma_v(\Delta)$ approaches positive or negative infinity, and therefore the correlation approaches ρ or $-\rho$ in the limit.

We can use a similar argument to demonstrate that as $\Delta \to 0$, the risk neutral discrete-time stochastic volatility model (7.5) converges to the following dynamic

$$
\begin{aligned}
d\ln(S_t) &= (r - \frac{1}{2}v)dt + \sqrt{v}dz_1^* && (7.10)\\
dv &= [\kappa(\theta - v) + \sigma(\nu_2(1 - \rho^2) + \lambda\rho)v]dt + \sigma\sqrt{v}dz_2^*
\end{aligned}
$$

where z_1^* and z_2^* are two Wiener processes such that $dz_1^*dz_2^* = \rho dt$.

7.3 The relationship with the continuous time affine SV model

Both (7.9) and (7.10) are square root stochastic volatility models of the type proposed by Heston (1993). We now link our discrete time stochastic volatility model and its risk-neutralization to the conventional risk-neutralization in the Heston (1993) model. Assume for simplicity that the parameterization of the conditional mean dynamic under the physical measure is given by (7.9). Heston (1993) proposes the following risk neutralization

$$
\begin{aligned}
d\ln(S_t) &= (r - \frac{1}{2}v)dt + \sqrt{v}dz_1^* && (7.11)\\
dv &= [\kappa(\theta - v) - \sigma\chi^* v]dt + \sigma\sqrt{v}dz_2^*
\end{aligned}
$$

where z_1^* and z_2^* are two Wiener process under the risk neutral probability Q and[21]

$$\begin{aligned} dz_1^* &= dz_1 + \left(\lambda - \frac{1}{2}\right)\sqrt{v}dt \\ dz_2^* &= dz_2 + \chi^*\sqrt{v}dt \end{aligned} \tag{7.12}$$

In the discrete-time stochastic volatility model, the parameter λ in (7.3) captures the price of equity risk, and ν_2 captures the price of volatility risk. In the Heston model, the price of equity risk λ plays the same role as in the discrete-time model, and we have also a price of volatility risk χ^* which ensures the affine structure of the risk-neutral process. Comparing (7.11) and (7.10), we find

$$\chi^* = \nu_2(1-\rho^2) + \lambda\rho. \tag{7.13}$$

Note that for $\rho = 0$, the continuous time price of volatility risk χ^* is not related to λ, but is simply equal to the discrete-time price of volatility risk ν_2. Moreover, this mapping between the price of volatility risk in discrete-time and continuous-time stochastic volatility models also provides insight into the relationship between the discrete-time GARCH model and the available continuous-time literature. While the GARCH model contains a single innovation, it can usefully be thought of as a special case of the two-shock discrete-time stochastic volatility model in (7.4), for $\rho = 1$ (or $\rho = -1$). In this case, from (7.13), $\chi^* = \lambda$. Because the GARCH model contains a single shock, the specification of the equity risk premium λ does double duty: it also implicitly defines the price of volatility risk, which is perfectly correlated with the price of equity risk by design. In other words, the GARCH return dynamic implicitly makes an assumption about the volatility risk premium. The parameter governing the equity risk premium also determines the volatility risk premium. Strictly speaking therefore, in the case of the GARCH model the only assumption we make in our approach is on the form of the Radon-Nikodym derivative. All other assumptions needed for risk-neutral valuation are implicit in the specification of the return dynamic, or, in other words, the assumptions on the equilibrium supporting the valuation problem are implicitly incorporated in the risk premium assumption for the return dynamic.

7.4 Option Price Uniqueness

In Section 6.3 and 7.2, we have analyzed continuous-time limits for four discrete-time returns dynamics: a homoskedastic model with normal innovations, a homoskedastic model with non-normal innovations, a heteroskedastic GARCH model with normal innovations, and a discrete time stochastic volatility model. We now further discuss these limit results by discussing the implications of three model features: normality, heteroskedasticity and the number of innovations (one versus two). We pay particular attention to the uniqueness of the option price in the discrete-time models we consider and their continuous-time limits.

First, to illustrate the importance of the conditional normality assumption, compare the results for the homoskedastic normal model in Section 6.1 with the results for the homoskedastic IG model in Section 6.2. The option value in the homoskedastic normal model is always equal to the BSM value, regardless of the value of μ. The existence of a unique option

[21] Notice that for ease of interpretation, in our notation the price of volatility risk χ^* has been rescaled by $1/\sigma$ compared to the notation in Heston (1993a).

value, rather than bounds on option values, may seem surprising because the structure of the underlying discrete-time economy is very different from that of the continuous-time limit economy. Markets are not complete and therefore the EMM is not necessarily unique. However, the intuition for this finding is simple: although there is a unique solution to the EMM condition (2.3), this finding is conditional on the assumption regarding the Radon-Nikodym derivative. Presumably other solutions may be available given alternative assumptions on the Radon-Nikodym derivative, but an investigation of this issue is beyond the scope of this paper.

In contrast to the case of conditionally normal innovations, the homoskedastic IG option value is generally different from the BSM value. The option value depends on the skewness parameter, as well as on the equity premium. In the limit, as skewness goes to zero, the option value converges to BSM.

Second, Section 7.3 suggests that it is instructive to distinguish between two types of heteroskedastic discrete-time models: one-shock models, such as GARCH models, and two-shock stochastic volatility models. We focus on normal innovations in order to facilitate the comparison with continuous-time models. We analyze a discrete-time stochastic volatility model which has the popular Heston (1993) model as a limit, and we are able to relate its risk neutralization to that used in the Heston model. In both cases option valuation takes place in an incomplete-markets setup, and in both cases this involves the use of a price of volatility risk which can be conveniently chosen to preserve the affine structure of the risk-neutral model dynamic. The intimate relationship between the risk-neutralization in the discrete-time and continuous-time models is due to three model characteristics: first, both models are affine; second, both models contain two innovations; and third, the discrete-time model has normal innovations. The Wiener processes in the continuous-time model correspond to a normality assumption. While continuous-time specifications can model non-normal return distributions, the instantaneous distribution is normal.[22]

GARCH models can be studied as special cases of discrete-time stochastic volatility models. In the GARCH model, one parameter determines the volatility risk premium as well as the equity risk premium. This is consistent with the interpretation of the GARCH model as a one-shock model with perfectly correlated equity and volatility innovations.[23]

In summary, in our opinion the analysis of discrete-time option valuation models has a lot of similarities with the analysis of continuous-time option valuation models. We provide a unique EMM solution subject to the choice of Radon-Nikodym derivative. Because the discrete-time, infinite state-space setup is characterized by incomplete markets, other Radon-Nikodym derivatives may of course exist. But this is analogous to the standard continuous time stochastic volatility setting, where markets are also incomplete, and where a Radon-Nikodym derivative or a volatility risk premium specification must be chosen in order to

[22] In a discrete-time model, non-normality can be built in by assuming that the multi-period distribution is non-normal, as in a continuous-time model, or by assuming that the one-period return is conditionally non-normal. This does of course not mean that the discrete-time setup is superior to the continuous-time setup. In the end the suitability of either approach will be determined by the trade-off between its flexibility and its analytical convenience.

[23] While it could be argued that this structure limits the usefulness of the GARCH model, one has to keep in mind that this structure is exactly what makes the GARCH model econometrically tractable. Indeed, the success of the GARCH model in modeling returns, and its growing popularity in modeling options, are precisely due to the fact that despite its simple structure it provides a very good fit.

select a unique option price.

Conceptually, the biggest difference between discrete and continuous time is the classic i.i.d. normal case, when markets are complete in continuous time but incomplete in discrete time. The discrete time option price is now unique only subject to the choice of a particular Radon-Nikodym derivative. While this striking continuous-time result is of great theoretical importance and interest, we also note that the evidence indicates that non i.i.d heteroskedastic models characterized by market incompleteness are more empirically relevant for option valuation, both in continuous and in discrete time. Bakshi, Cao and Chen (1997) make a strong case for this using continuous time stochastic volatility models, and Heston and Nandi (2000) and Christoffersen and Jacobs (2004) provide evidence using discrete time GARCH models.

Most existing papers on option pricing in discrete time using an infinite state space setup[24] assume normally distributed returns and, in the words of Rubinstein (1976), "complete" the markets by assuming a representative agent with certain preferences, such as for instance constant relative risk aversion. Our approach, much like the one used in the continuous-time stochastic volatility literature, is to let the researcher specify an empirically realistic return dynamic for the underlying asset, and subsequently provide an equivalent martingale measure that allows us to find option prices using a no-arbitrage assumption. Proposition 1 provides the form of the EMM and Proposition 4 provides the no-arbitrage option pricing result. Instead of making explicit assumptions regarding a representative agent with particular risk preferences, we assume a form of the Radon-Nikodym derivative, and a price of volatility risk, much like in the continuous time stochastic volatility literature. Whereas the assumption on the representative agent's utility function "completes" the market in the standard normal discrete time setting, the Radon-Nikodym derivative "completes" the market in our setup.

7.5 Option Price Bounds

There is a rich literature that derives bounds on option valuation using stochastic dominance techniques. While some of these papers focus on i.i.d. returns (see for example Perrakis (1986), Perrakis and Ryan (1984), Ritchken and Kuo (1988), and Constantinides, Jackwerth and Perrakis (2008)), recently Perrakis and Oancea (2007) have used stochastic dominance to derive option pricing bounds in the presence of GARCH processes for returns. Perrakis and Oancea (2007) show that in the GARCH case, stochastic dominance alone does not suffice to derive option bounds because of nonconvexities, and additional assumptions have to be made.[25] Under these types of assumptions, our approach leads to a unique price. Presumably bounds similar to the ones in Perrakis and Oancea (2007) may be obtained using other assumptions on the RN derivative, but this exploration is left for future work.

8 Conclusion

This paper provides valuation results for contingent claims in a discrete time infinite state space setup. Our valuation argument applies to a large class of conditionally normal and non-

[24] See for example Rubinstein (1976), Brennan (1979), and Duan (1995).

[25] Perrakis and Oancea (2007) discuss normal innovations. It is not clear how to relate our results to theirs for non-normal innovations.

normal stock returns with flexible time-varying mean and volatility, as well as a potentially time-varying price of risk. This setup generalizes the result in Duan (1995) in the sense that we do not restrict the returns to be conditionally normal, nor do we restrict the price of risk to be constant.

Our results apply to some of the most widely used discrete time processes in finance, such as GARCH processes. We also apply our approach to the analysis of discrete-time processes with multiple innovations, such as discrete-time stochastic volatility processes. For the class of processes we analyze, the risk neutral return dynamic is the same as the physical dynamic, but with a different parameterization which we are able to characterize completely. To provide intuition for our findings, we extensively discuss the relationship between our results and existing results for continuous-time stochastic volatility models, which can be derived as limits of our discrete-time dynamics.

To demonstrate the empirical relevance of our approach, we provide an empirical analysis of a heteroskedastic return dynamic with a standardized skewed variance gamma distribution, which is constructed as the mixture of two gamma variables. In the resulting dynamic, conditional skewness and kurtosis are directly governed by two distinct parameters. We estimate the model on return data using quasi maximum likelihood and compare its performance with the heteroskedastic conditional normal model which is standard in the literature. Diagnostics clearly indicate that the conditionally non-normal model outperforms the conditionally normal model, and an analysis of the option smirk demonstrates that this model provides substantially more flexibility to value options.

It is likely that our risk neutralizations can equivalently be derived using the specification of a candidate stochastic discount factor or using the Esscher transform, rather than through our approach which starts with the specification of a Radon-Nikodym derivative and derives the EMM.[26] However, in most applications that we are aware of, existing work actually starts out by assuming a bivariate distribution for the stochastic discount factor and the stock return (see for example Amin and Ng (1993)).[27] This assumption clearly goes beyond the existence of no-arbitrage and is closer in spirit to the general equilibrium setup of Duan (1995) and Brennan (1979). See Garcia, Ghysels and Renault (2006) for a discussion on how some of these assumed joint distributions effectively amount to degenerate distributions. Our approach is also related to the risk-neutral valuation argument used in Heston (1993b, 2004) and Christoffersen, Heston and Jacobs (2006), but in our opinion our approach is more transparent. Duan, Ritchken and Sun (2005) use a risk neutralization for a Poisson-normal heteroskedastic model that has some similarities with our approach. However, they do not apply their principle to the investigation of more general return dynamics.

We leave several important issues unaddressed. First, while we obtain a unique EMM given the choice of Radon-Nikodym derivative, we cannot exclude that even for a given specification of the risk premium, there exist other EMMs corresponding to different functional forms of the Radon-Nikodym derivative. Second, it would be interesting to more fully explore the relationship between our findings and those in the stochastic dominance literature that derives option bounds. Third, while we advocate separating the valuation issue and the general equilibrium setup that supports it, the general equilibrium foundations of our results

[26] See for example Hansen and Richard (1987) for a characterization of risk neutralization using the stochastic discount factor. See also Buhlman et al (1998) and Siu et al (2004) for applications of the Esscher transform in option valuation.

[27] See Gourieroux and Monfort (2007) for a notable exception.

are of course very important. It may prove possible to characterize the equilibrium setup that gives rise to the risk neutralization proposed for some of the processes considered in this paper, such as the empirically interesting dynamics considered in Section 5. However, this is by no means a trivial problem, and it is left for future work.

9 Appendix

9.1 Proof of Proposition 2.

Define $f(v) = \Psi(v) - \Psi(v-1)$. We want to show that $f(v) = E[R_t - r] \equiv \pi$ has a unique solution for any given value of π. Existence is obtained if $f(v)$ can take any real value. The assumption that Ψ tends to infinity at the boundaries of its domain implies that we have $\Psi(u_1) = \infty$ and $\Psi(u_2) = \infty$. Note that the domain of $f(v)$ is (u_1+1, u_2). Note also that Ψ is continuous because by assumption it is twice differentiable on its domain. Since Ψ is continuous $f(.)$ is also continuous. Now let us evaluate the function $f(.)$ on u_1+1 and u_2. $f(u_1+1) = \Psi(u_1+1) - \Psi(u_1) = -\Psi(u_1) = -\infty$ since $\Psi(u_1) = \infty$. $f(u_2) = \Psi(u_2) - \Psi(u_2 - 1) = \Psi(u_2) = \infty$ since $\Psi(u_2) = \infty$. Hence $f(.)$ is continuous and can attain $-\infty$ or ∞. This implies that for any value π between $(-\infty, \infty)$, there exists a value v in the domain of the continuous function $f(.)$ such that $f(v) = \pi$. Furthermore, we have that $f'(u) = \Psi'(u) - \Psi'(u-1)$. Convexity of Ψ implies that if $\Psi''(u) > 0$, then $\Psi'(.)$ is increasing. A direct consequence of this is that if $f'(u) = \Psi'(u) - \Psi'(u-1) > 0$ then $f(.)$ is increasing. We then have that $f(.)$ is increasing and continuous, which implies that $f(.)$ is a bijection, and uniqueness follows.

9.2 Proof of Lemma 2.

For a self financing strategy we have

$$\begin{aligned} G_{t+1} = V_{t+1} &= \eta_t S_{t+1} + \delta_t C_{t+1} + \psi_t B_{t+1} \\ &= \eta_{t+1} S_{t+1} + \delta_{t+1} C_{t+1} + \psi_{t+1} B_{t+1} \end{aligned}$$

We also have

$$G_t = \sum_{i=0}^{t-1} \eta_i (S_{i+1} - S_i) + \sum_{i=0}^{t-1} \delta_i (C_{i+1} - C_i) + \sum_{i=0}^{t-1} \psi_i (B_{i+1} - B_i).$$

It follows that

$$G_{t+1} - G_t = \eta_t (S_{t+1} - S_t) + \delta_t (C_{t+1} - C_t) + \psi_t (B_{t+1} - B_t)$$

We can trivially also write

$$G^B_{t+1} - G^B_t = G^B_{t+1} - G^B_t + \underbrace{\left(\frac{G_{t+1}}{B_t} - \frac{G_{t+1}}{B_t} \right)}_{=0}$$

This implies that

$$G_{t+1}^B - G_t^B = (\eta_t S_{t+1} + \delta_t C_{t+1} + \psi_t B_{t+1}) \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right)$$
$$+ \frac{1}{B_t} \left(\eta_t (S_{t+1} - S_t) + \delta_t (C_{t+1} - C_t) + \psi_t (B_{t+1} - B_t) \right).$$

$$\begin{aligned} = \; & \eta_t \left[S_{t+1} \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right) + \frac{1}{B_t} (S_{t+1} - S_t) \right] \\ & + \delta_t \left[C_{t+1} \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right) + \frac{1}{B_t} (C_{t+1} - C_t) \right] \\ & + \underbrace{\psi_t B_{t+1} \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right) + \frac{1}{B_t} \psi_t (B_{t+1} - B_t)}_{=0} \end{aligned}$$

Then

$$G_{t+1}^B - G_t^B = \eta_t \left[S_{t+1} \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right) + \frac{1}{B_t} (S_{t+1} - S_t) \right] + \delta_t \left[C_{t+1} \left(\frac{1}{B_{t+1}} - \frac{1}{B_t} \right) + \frac{1}{B_t} (C_{t+1} - C_t) \right]$$
$$= \eta_t (S_{t+1}^B - S_t^B) + \delta_t (C_{t+1}^B - C_t^B) + \left(\eta_t \frac{S_{t+1}}{B_t} - \eta_t \frac{S_{t+1}}{B_t} \right) + \left(\delta_t \frac{C_{t+1}}{B_t} - \delta_t \frac{C_{t+1}}{B_t} \right)$$

and therefore

$$G_{t+1}^B - G_t^B = \eta_t (S_{t+1}^B - S_t^B) + \delta_t (C_{t+1}^B - C_t^B). \qquad \forall t = 1, ..., T-1$$

Because $G_0 = G_0^B = 0$ the discounted gain can be written as the sum of past changes

$$G_t^B = \sum_{i=0}^{t-1} (G_{i+1}^B - G_i^B) \; \forall t = 1, ..., T.$$

Therefore the discounted gain can be written

$$G_t^B = \sum_{i=0}^{t-1} \eta_i (S_{i+1}^B - S_i^B) + \sum_{i=0}^{t-1} \delta_i (C_{i+1}^B - C_i^B)$$

and the proof is complete.

9.3 Proof of Proposition 3.

From Lukacs (1970), page 119, we have the Kolmogorov canonical representation of the log-moment generating function of an infinitely divisible distribution function. This result stipulates that a function Ψ is the log-moment generating function of an infinitely divisible distribution with finite second moment if, and only if, it can be written in the form

$$\Psi(u) = -uc + \int_{-\infty}^{+\infty} \left(e^{-ux} - 1 + ux \right) \frac{dK(x)}{x^2}$$

where c is a real constant while $K(u)$ is a nondecreasing and bounded function such that $K(-\infty) = 0$. Applying this theorem gives the following form for $\Psi_t(u)$,

$$\Psi_t(u) = -uc_{t-1} + \int_{-\infty}^{+\infty} \left(e^{-ux} - 1 + ux\right) \frac{dK_{t-1}(x)}{x^2} \tag{9.1}$$

where c_{t-1} is a random variable known at $t-1$, and $K_{t-1}(x)$ is a function known at $t-1$, which is nondecreasing and bounded so that $K_{t-1}(-\infty) = 0$. Using relation (2.8) and the characterisation (9.1) we can write $\Psi_t^{Q*}(u)$ as

$$\Psi_t^{Q*}(u) = \int_{-\infty}^{+\infty} \left(e^{-ux} - 1 + ux\right) \frac{dK_{t-1}^*(x)}{x^2}$$

where

$$K_{t-1}^*(x) = \int_{-\infty}^{x} e^{-\nu_{t-1}y} dK_{t-1}(y)$$

This implies that

$$K_{t-1}^*(-\infty) = 0$$

$K_{t-1}^*(x)$ is obviously non-decreasing since $K_{t-1}(x)$ is non-decreasing, $K_{t-1}^*(\infty) < \infty$, because $K_{t-1}(\infty) < \infty$, and $e^{-\nu_t y}$ is a decreasing function of y which converge to 0. Recall that ν_t is the price of risk, which is positive and known at time $t-1$.

In conclusion we have constructed a constant c_{t-1}^* $(= 0)$ and a non-decreasing bounded function $K_{t-1}^*(x)$, with $K_{t-1}^*(-\infty) = 0$, such that

$$\Psi_t^{Q*}(u) = -uc_{t-1}^* + \int_{-\infty}^{+\infty} \left(e^{-ux} - 1 + ux\right) \frac{dK_{t-1}^*(x)}{x^2}.$$

Hence, according to the Kolmogorov canonical representation, the conditional distribution of ε_t^* is infinitely divisible.

References

[1] Amin, K. and V. Ng (1993), ARCH Processes and Option Valuation, Manuscript, University of Michigan.

[2] Bakshi, C., Cao, C. and Z. Chen (1997), Empirical Performance of Alternative Option Pricing Models, Journal of Finance, 52, 2003-2049.

[3] Bates, D. (1996), Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, Review of Financial Studies, 9, 69-107.

[4] Bates, D. (2000), Post-'87 Crash Fears in the S&P 500 Futures Option Market, Journal of Econometrics, 94, 181-238.

[5] Bates, D. (2003), Empirical Option Pricing: a Retrospection, Journal of Econometrics, 116, 387-404.

[6] Bick, A. (1990), On Viable Diffusion Price Processes of the Market Portfolio, Journal of Finance, 45, 673-689.

[7] Black, F., and M. Scholes (1973), The Pricing of Options and Corporate Liabilities, Journal of Political Economy, 81, 637-659.

[8] Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, Journal of Econometrics, 31, 307-327.

[9] Bollerslev, T., Chou, R. and K. Kroner (1992), ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence, Journal of Econometrics, 52, 5-59.

[10] Bollerslev, T., M. Gibbons, and H. Zhou (2005), Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities, Manuscript, Duke University.

[11] Bollerslev, T. and H.O. Mikkelsen (1996), Long-Term Equity AnticiPation Securities and Stock Market Volatility Dynamics, Journal of Econometrics, 92, 75-99.

[12] Brennan, M. (1979), The Pricing of Contingent Claims in Discrete-Time Models, Journal of Finance, 34, 53-68.

[13] Broadie, M., Chernov, M. and M. Johannes (2006), Model Specification and Risk Premia: Evidence from S&P500 Futures Options Market, forthcoming in the Journal of Finance.

[14] Buhlmann, H., Delbaen, F., Embrechts, P., Shiryaev, A., 1998. On Esscher Transforms in Discrete Finance Models. ASTIN Bulletin, 28, 171–86.

[15] Camara, A. (2003), A Generalization of the Brennan-Rubenstein Approach for the Pricing of Derivatives, Journal of Finance, 58, 805-819.

[16] Carr, P. and L. Wu (2004), Time-Changed Levy Processes and Option Pricing, Journal of Financial Economics, 17, 113–141.

[17] Carr, P. and L. Wu (2007), Variance Risk Premia, Review of Financial Studies, forthcoming.

[18] Christoffersen, P., S. Heston, and K. Jacobs (2006), Option Valuation with Conditional Skewness, Journal of Econometrics, 131, 253-284.

[19] Christoffersen, P. and K. Jacobs (2004), Which GARCH Model for Option Valuation? Management Science, 50, 1204-1221.

[20] Corradi, V. (2000), Reconsidering the Continuous Time Limit of the GARCH(1,1) Process, Journal of Econometrics, 96, 145-153.

[21] Cox, J. and S. Ross (1976), The Valuation of Options for Alternative Stochastic Processes, Journal of Financial Economics, 3, 145-166.

[22] Cox, J., Ross, S. and M. Rubinstein (1979), Option Pricing: A Simplified Approach, Journal of Financial Economics, 7, 229-264.

[23] Diebold, F. and J. Lopez (1995), Modeling Volatility Dynamics, in Macroeconometrics: Developments, Tensions and Prospects, Kevin Hoover (ed.). Boston: Kluwer Academic Press.

[24] Duan, J.-C. (1995), The GARCH Option Pricing Model, Mathematical Finance, 5, 13-32.

[25] Duan, J.-C. (1997), Augmented GARCH(p,q) Process and Its Diffusion Limit, Journal of Econometrics, 79, 97-127.

[26] Duan, J.-C. (1999), Conditionally Fat-Tailed Distributions and the Volatility Smile in Options, Manuscript, University of Toronto.

[27] Duan, J.-C., Ritchken, P. and Z. Sun (2005), Jump Starting GARCH: Pricing and Hedging Options with Jumps in Returns and Volatilities, Working Paper, Rotman School, University of Toronto.

[28] Engle, R. (1982), Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation, Econometrica, 50, 987-1008.

[29] Engle, R., D. M. Lilien and R.P. Robins (1987), Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model, Econometrica, 55, 391-407.

[30] Engle, R. and V. Ng (1993), Measuring and Testing the Impact of News on Volatility, Journal of Finance, 48, 1749-1778.

[31] Eraker, B. (2004), Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, Journal of Finance, 59, 1367-1403.

[32] Eraker, B., M. Johannes and N. Polson (2003), The Impact of Jumps in Volatility and Returns, Journal of Finance, 58, 1269-1300.

[33] Feller, W. (1968), An Introduction to Probability Theory and its Applications. John Wiley & Sons.

[34] Foster, D. P. and D.B. Nelson (1996), Continuous Record Asymptotics for Rolling Sample Variance Estimators, Econometrica, 64, 139-174.

[35] French, K., G. W. Schwert, and R. Stambaugh (1987), Expected Stock Returns and Volatility, Journal of Financial Economics, 19, 3-30.

[36] Garcia, R., E. Ghysels, and E. Renault (2006), The Econometrics of Option Pricing, forthcoming in the Handbook of Financial Econometrics, Yacine Aït-Sahalia and Lars Peter Hansen, eds., Elsevier-North Holland, Amsterdam.

[37] Garcia, R. and E. Renault (1998), A Note on Hedging in ARCH-type Option Pricing Models, Mathematical Finance, 8, 153-161.

[38] Glosten, L. R., R. Jagannathan, and D. Runkle (1993), On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, Journal of Finance, 48, 1779-1801.

[39] Ghysels, E., A. Harvey, and E. Renault (1995), Stochastic Volatility, Handbook of Statistics 14, Statistical Methods in Finance, G.S. Maddala and C.R. Rao (eds.), North Holland. Amsterdam.

[40] Gourieroux, C. and A. Monfort (2007), Econometric Specifications of Stochastic Discount Factor Models, forthcoming in the Journal of Econometrics, 136, 509-530.

[41] Hansen, L. and S. Richard (1987), The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models, Econometrica, 55, 587-613.

[42] Harrison, M. and D. Kreps (1979), Martingales and Arbitrage in Multi-period Securities Markets, Journal of Economic Theory, 20, 381-408.

[43] Harrison, M., and S. Pliska (1981), Martingales and Stochastic Integrals in the Theory of Continuous Trading, Stochastic Processes and their Applications, 11, 215–260.

[44] He, H. and H. Leland (1993), On Equilibrium Asset Price Processes, Review of Financial Studies, 6, 593-617.

[45] Heston, S. (1993a), A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, Review of Financial Studies, 6, 327-343.

[46] Heston, S. (1993b), Invisible Parameters in Option Prices, Journal of Finance, 48, 933-947.

[47] Heston, S. (2004), Option Valuation with Infinitely Divisible Distributions, Quantitative Finance, 4, 515-524.

[48] Heston, S. and S. Nandi (2000), A Closed-Form GARCH Option Pricing Model, Review of Financial Studies, 13, 585-626.

[49] Huang, J.-Z. and L. Wu (2004), Specification Analysis of Option Pricing Models Based on Time-Changed Levy Processes, Journal of Finance, 59, 1405–1439.

[50] Jacob, J. and A. Shiryaev (1998), Local Martingales and the Fundamental Asset Pricing Theorems in the Discrete-time Case, Finance and Stochastics, 259-273.

[51] Karatzas, I., and S. Shreve (1998), Methods of Mathematical Finance, Springer-Verlag.

[52] Lukacs, E. (1970), Characteristic Functions, Griffin. London

[53] Madan, D. B., and E. Seneta (1990), The Variance Gamma Model from Share Market Returns, Journal of Business, 63, 511-524.

[54] Maheu, J., and T. McCurdy (2004), News Arrival, Jump Dynamics and Volatility Components for Individual Stock Returns, Journal of Finance, 59, 755–793.

[55] Nelson, D. B. (1991), Conditional Heteroskedasticity in Asset Returns: A New Approach, Econometrica, 59, 347-370.

[56] Nelson, D. B. (1996), Asymptotically Optimal Smoothing with Arch Models, Econometrica, 64, 561-573.

[57] Nelson, D. B. and D. P. Foster (1994), Asymptotic Filtering Theory for Univariate Arch Models, Econometrica, 62, 1-41.

[58] Pan, J. (2002), The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, Journal of Financial Economics, 63, 3–50.

[59] Perrakis, S. (1986), Option Bounds in Discrete Time: Extensions and the Pricing of the American Put, Journal of Business, 59, 119-141.

[60] Perrakis, S. and P. Ryan (1984), Option Bounds in Discrete Time, Journal of Finance, 39, 519-525.

[61] Ritchken, P. and S. Kuo (1988), Option Bounds with Finite Revision Opportunities, Journal of Finance, 43, 301-308.

[62] Ritchken, P. and R. Trevor (1999), Option Pricing Under GARCH and Stochastic Volatility, Journal of Finance, 54, 377-402.

[63] Rubinstein, M. (1976), The Valuation of Uncertain Income Streams and the Pricing of Options, Bell Journal of Economics, 7, 407-425.

[64] Satchell, S., and A. Timmermann (1996), Option Pricing with GARCH and Systematic Consumption Risk, Derivatives Use, Trading & Regulation, Part II, 1, 353-67.

[65] Schroder, M. (2004), Risk Neutral Parameter Shifts and Derivatives Pricing in Discrete Time, Journal of Finance, 59, 2375-2402.

[66] Schwert, G.W. (1989), Why Does Stock Market Volatility Change over Time? Journal of Finance, 44, 1115-1153.

[67] Shiryaev, A. (1999), Essentials of Stochastic Finance. Facts, Models, Theory, World Scientific.

[68] Siu, T.K., H. Tong, and H. Yang (2004), On Pricing Derivatives under GARCH Models: A Dynamic Gerber-Shiu Approach, North American Actuarial Journal, 8, 17-31.

Figure 1: Quantile-Quantile Plot of S&P500 Returns Against the Normal Distribution

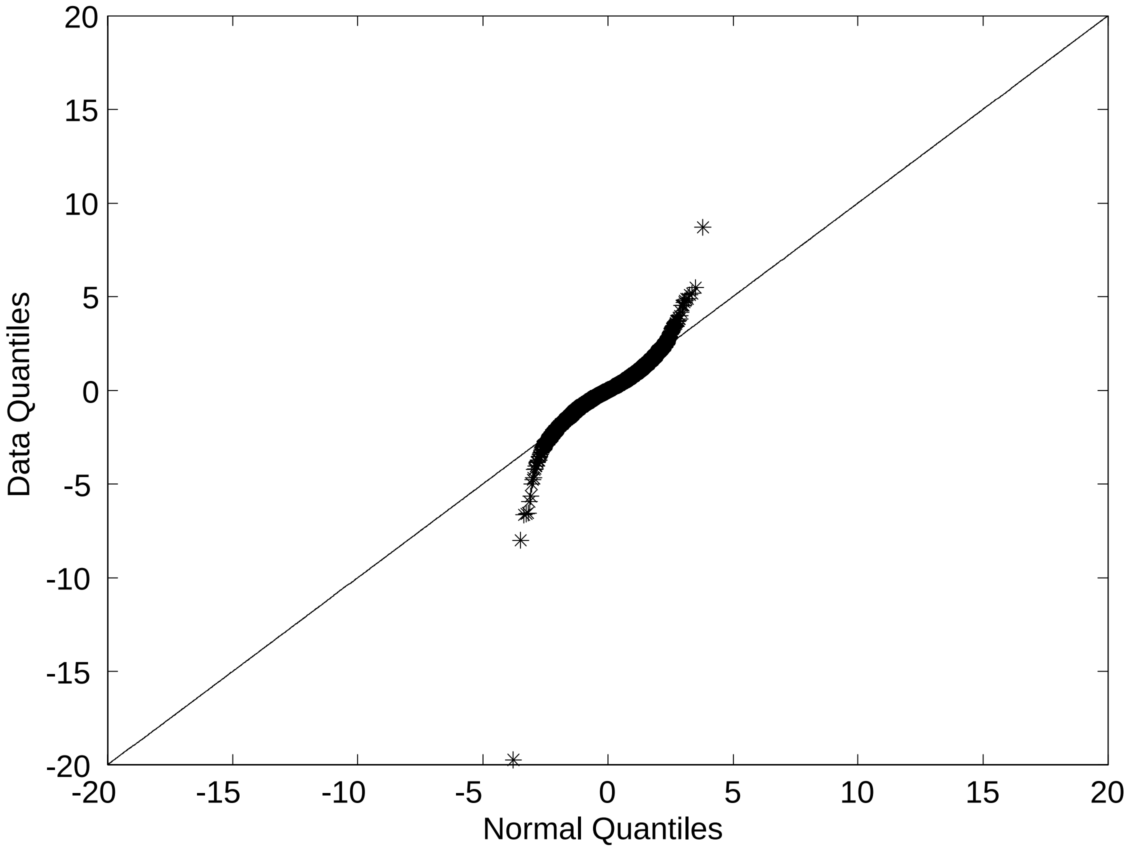

Notes to Figure: We take daily returns on the S&P500 from January 2, 1980 to December 30, 2005 and standardize them by the sample mean and sample standard deviation. The quantiles of the standardized returns are plotted against the quantiles from the standard normal distribution.

Figure 2: Autocorrelation Function of Absolute S&P500 Returns

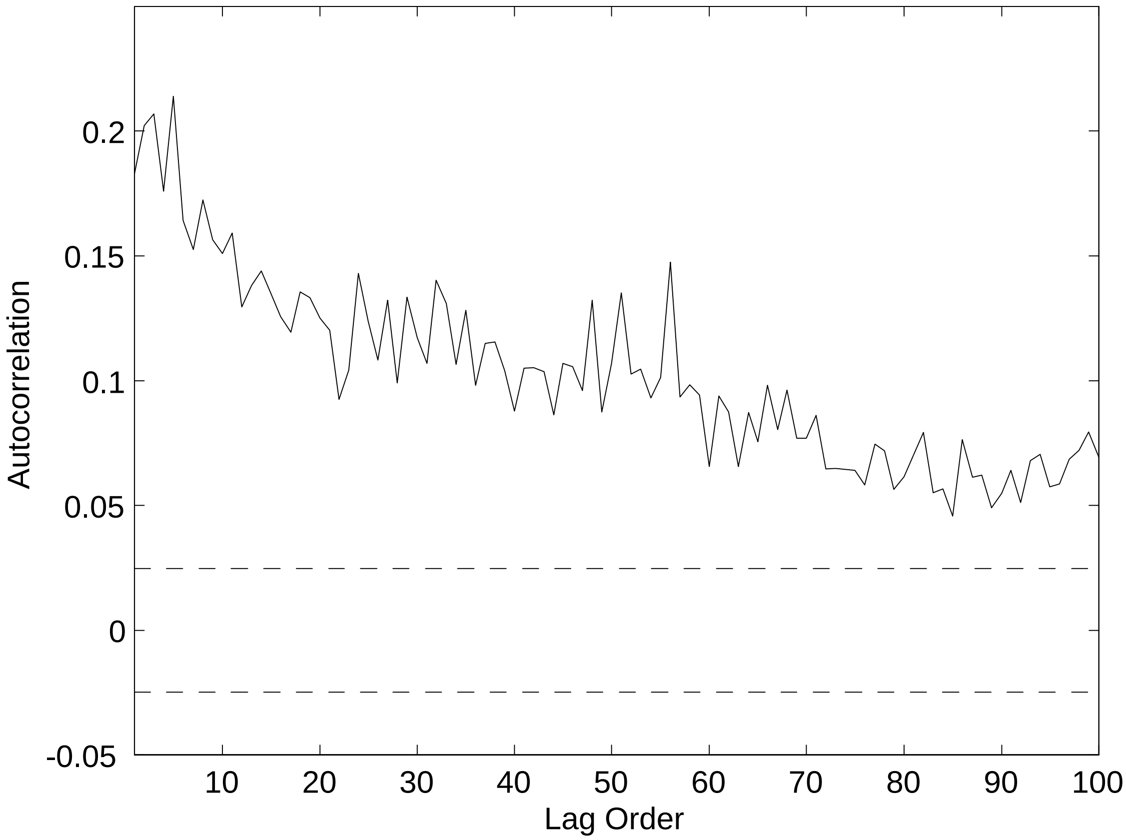

Notes to Figure: From daily absolute returns on the S&P500 from January 2, 1980 to December 30, 2005 we compute and plot the sample autocorrelations for lags one through 100 days. The horizontal dashed lines denote 95% Bartlett confidence intervals around zero.

Figure 3: Autocorrelation Function of Absolute GARCH Innovations

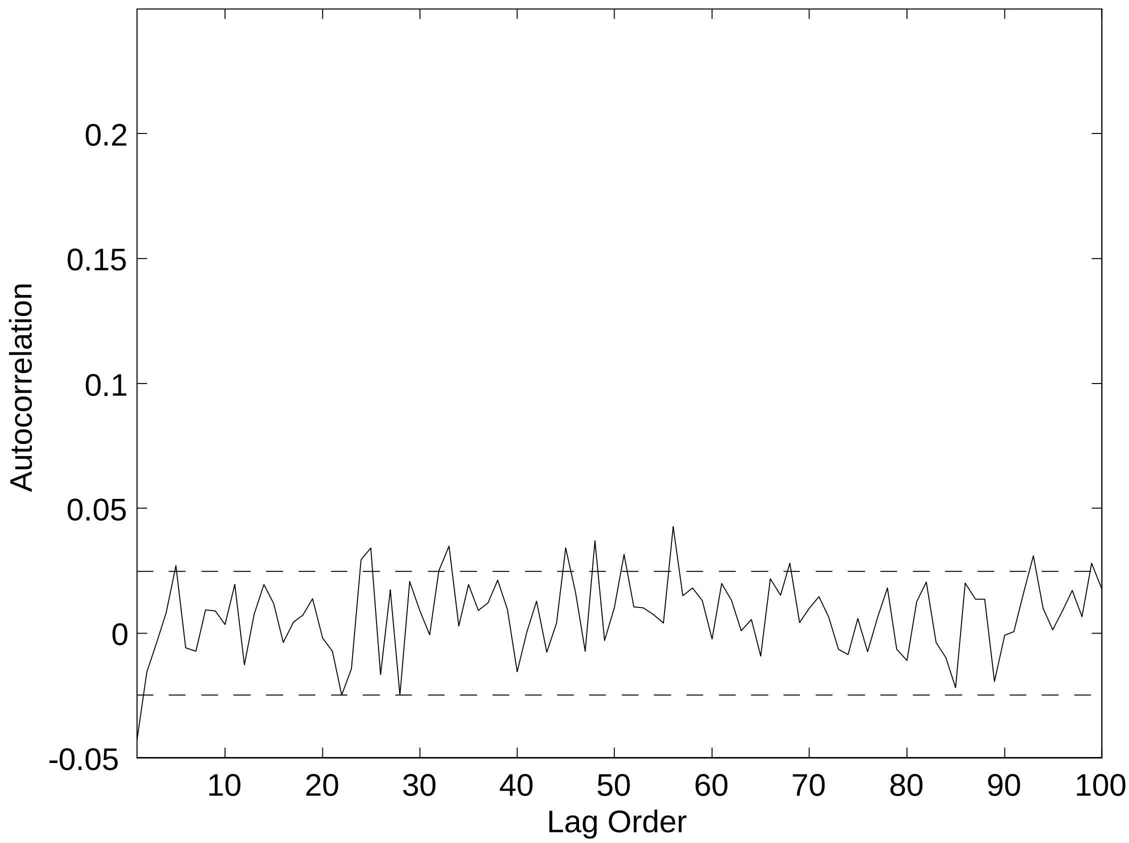

Notes to Figure: From the estimated GARCH model in Table 1 we construct the absolute standardized sequence of shocks and plot the sample autocorrelations for lags one through 100 days. The horizontal dashed lines denote 95% Bartlett confidence intervals around zero.

Figure 4: Quantile-Quantile Plots of GARCH Innovations Against the Normal Distribution

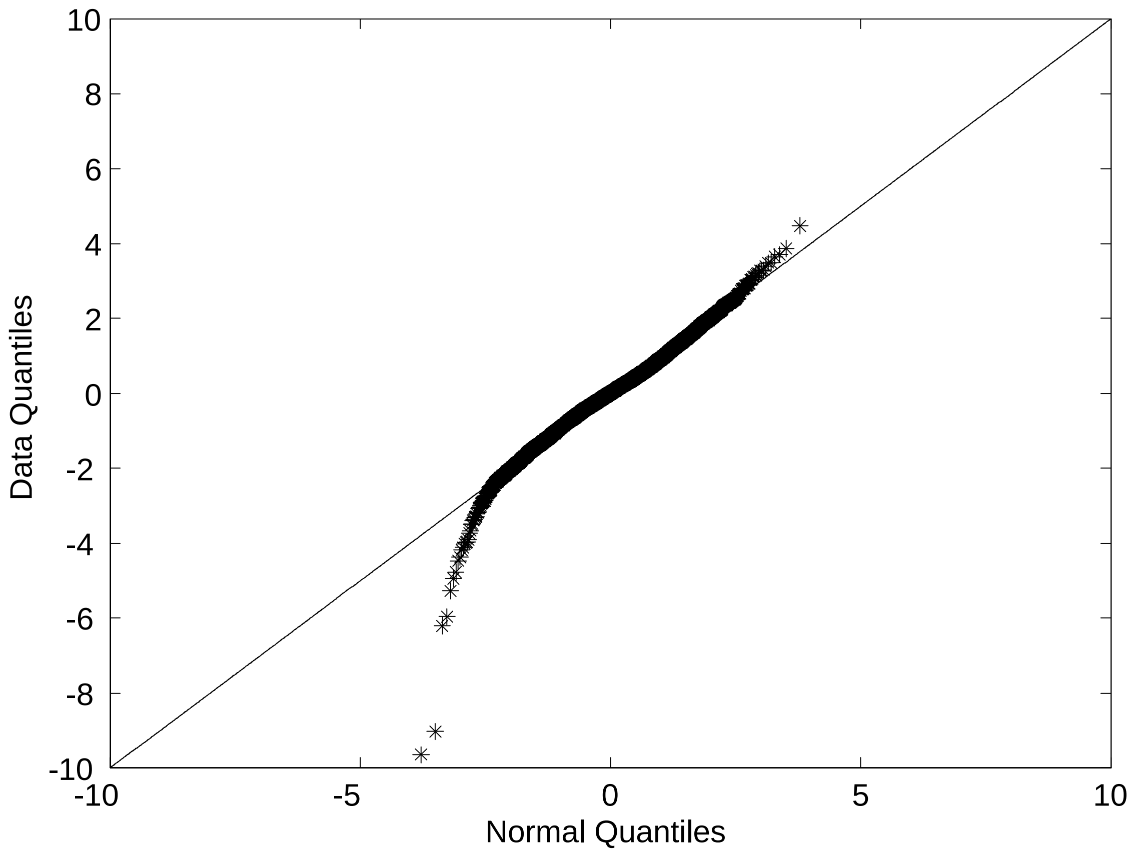

Notes to Figure: From the estimated GARCH models in Table 1 we compute the time series of dynamically standardized S&P500 returns. The quantiles of these GARCH innovations are plotted against the quantiles from the standard normal distribution.

Figure 5: Quantile-Quantile Plots of GARCH Innovations Against the SVG Distribution

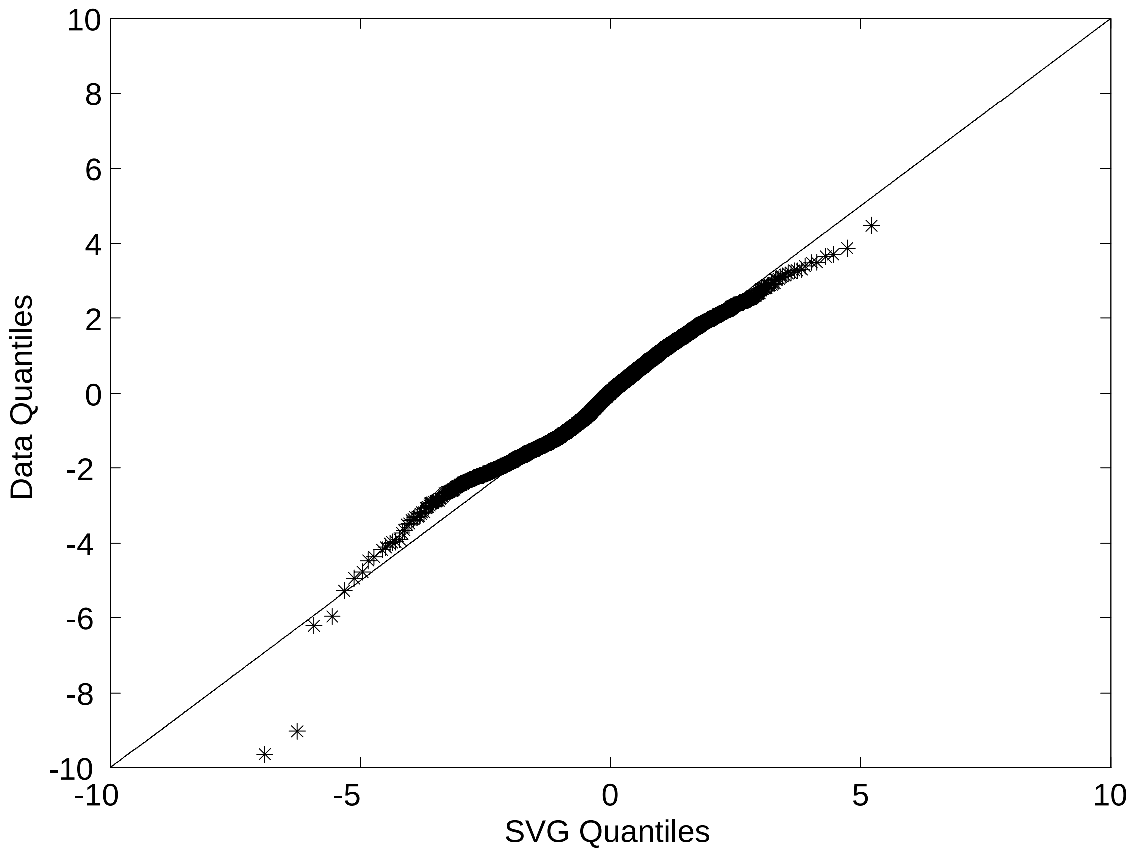

Notes to Figure: From the estimated GARCH models in Table 1 we compute the time series of dynamically standardized S&P500 returns. The quantiles of these GARCH innovations are plotted against the quantiles from the skewed variance gamma (SVG) distribution.

Figure 6: Implied Volatility Functions for Normal and SVG Independent Return Models

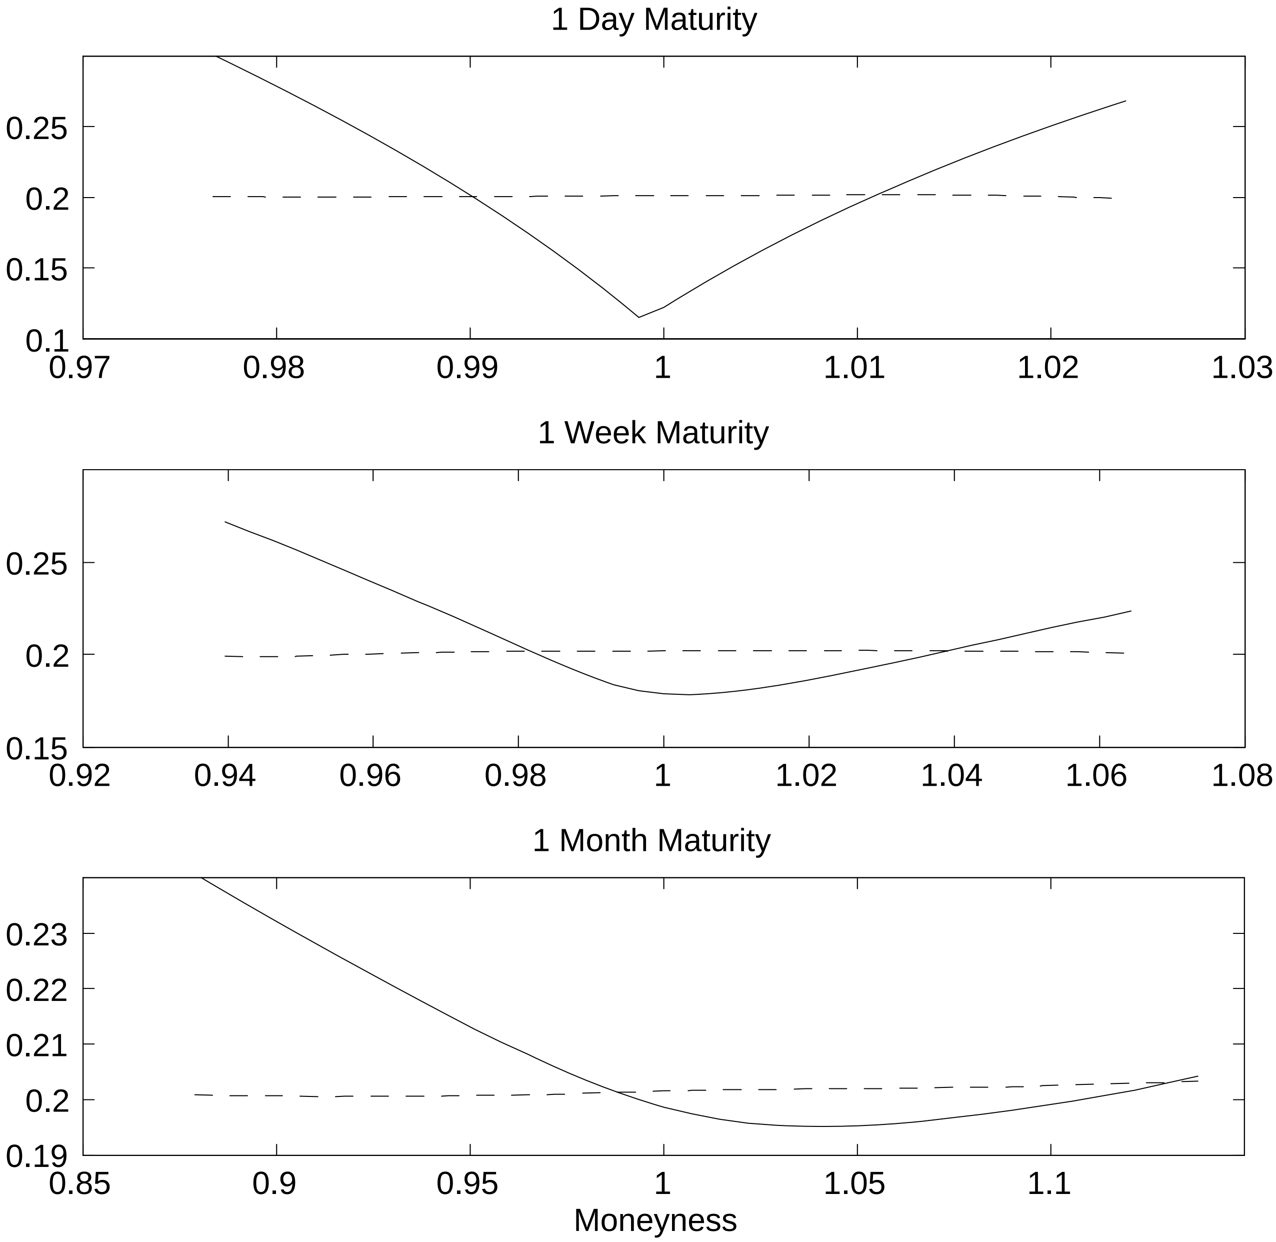

Notes to Figure: From the estimated independent return model in Table 1 we compute call option prices for various moneyness and maturities and we then compute implied Black-Scholes volatilities from the model option prices. Implied volatility is plotted against moneyness on the horizontal axis. The three panels correspond to maturities of 1 day, 1 week, and 1 month respectively. The solid lines show the i.i.d SVG model and the dashed lines the i.i.d. Normal models.

Figure 7: Implied Volatility Functions for Normal and SVG GARCH Models

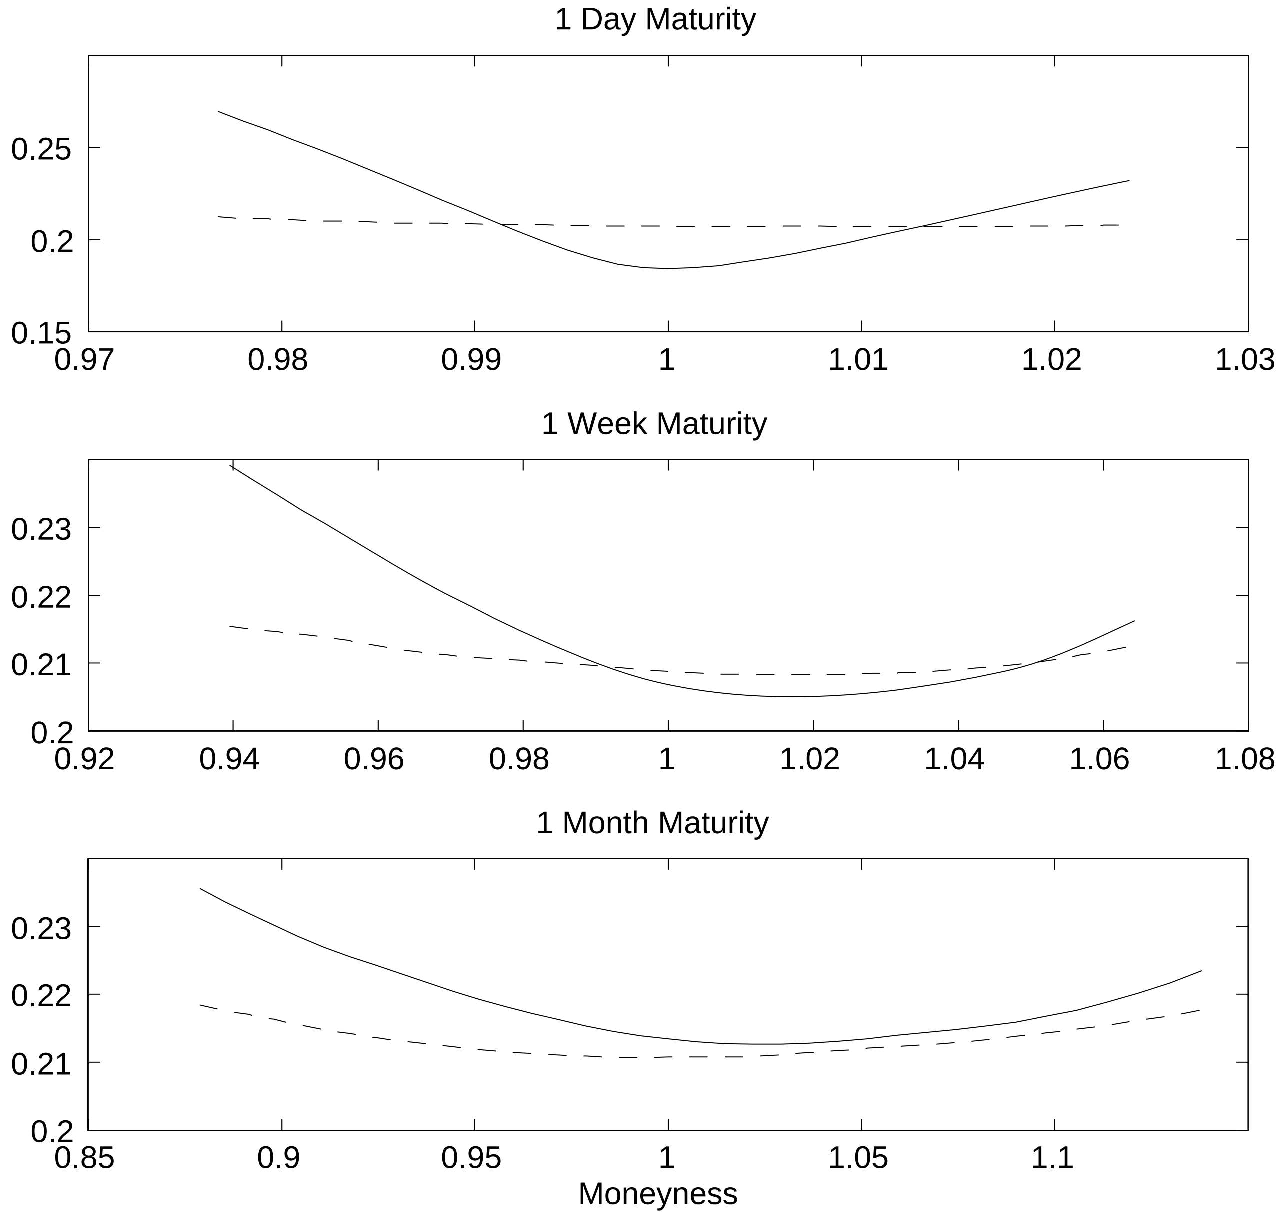

Notes to Figure: From the estimated GARCH model in Table 1 we compute call option prices for various moneyness and maturities and then we compute implied Black-Scholes volatilities from the model option prices. The implied volatilities are plotted with moneyness on the horizontal axis. The three panels correspond to maturities of 1 day, 1 week, and 1 month respectively. The solid lines show the SVG GARCH model and the dashed lines the Normal GARCH model.

Figure 8: Convergence of Homoskedastic Inverse Gaussian to Black-Scholes Option Price

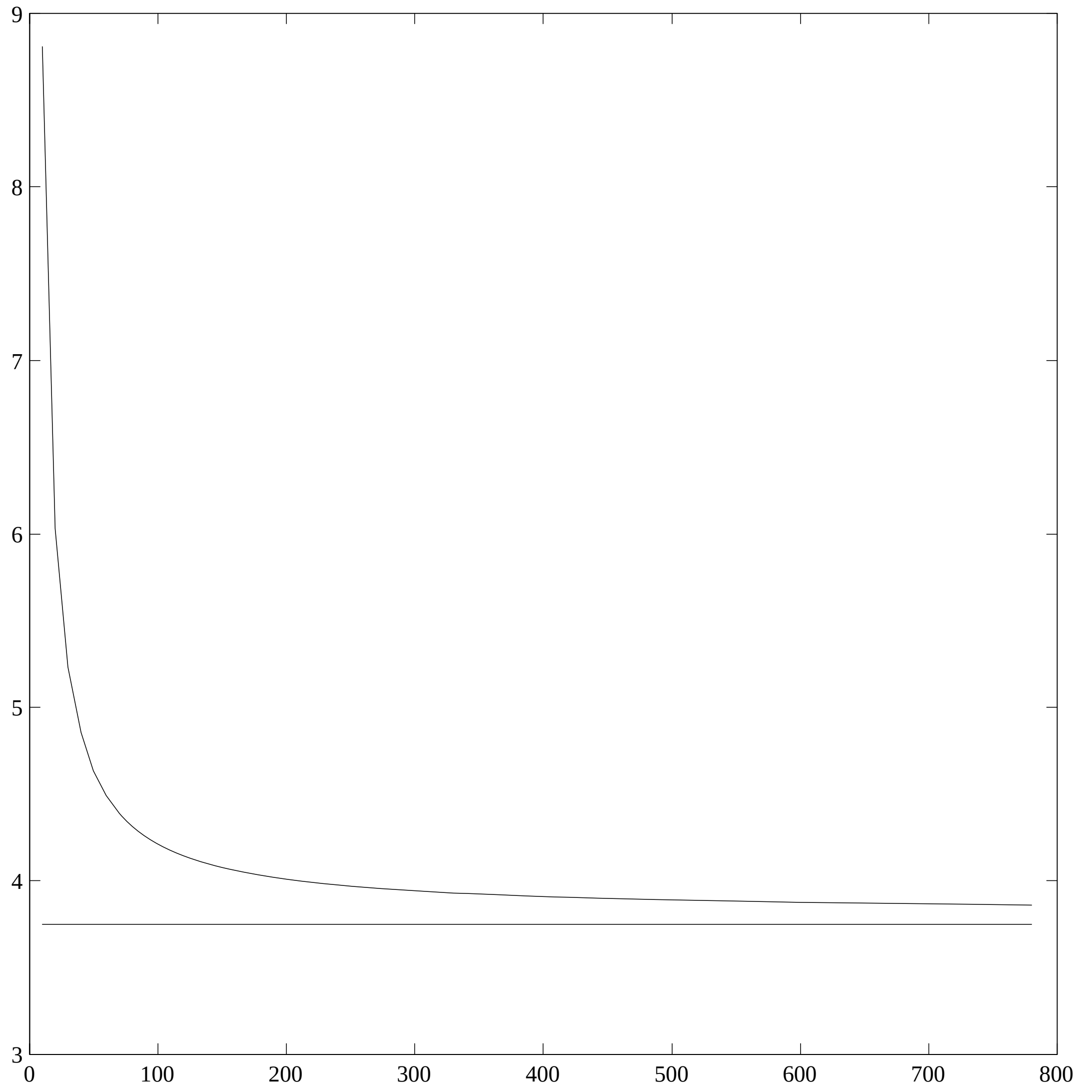

Notes to Figure: We illustrate the convergence of the homoskedastic IG option price to the BSM price when Δ goes to zero. We let $\eta(\Delta) = \eta\Delta$, $\sigma^2(\Delta) = \sigma^2\Delta$, and set $\lambda = 1$, $\eta = -0.001$ per day, and $\sigma = .10.$ per year.

Figure 9: Convergence of GARCH to Stochastic Volatility Option Price

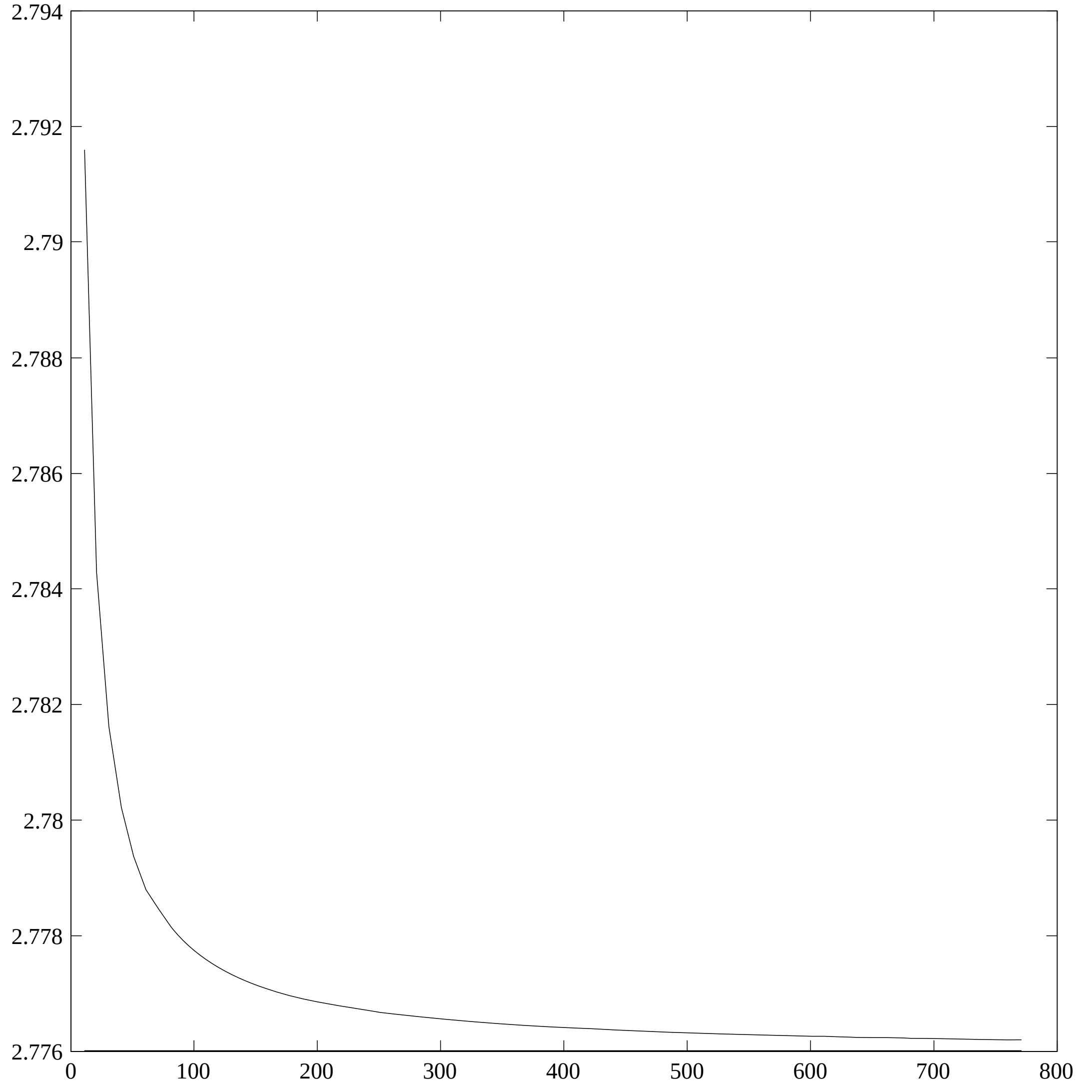

Notes to Figure: We show the convergence of the Heston and Nandi (2000) discrete time GARCH option price to the continuous time SV option price in Heston (1993). We use $r = 0$, $K = 100$, $S = 100$, $M\Delta = 180$, $\sigma = 0.1$, $V = 0.01$, $\kappa = 2$, $\theta = .01$, and shock correlation $\rho = -1$.

Table 1: Parameter Estimates and Model Properties
Daily S&P500 Returns. January 2, 1980 - December 30, 2005.

	Independent Returns		**GARCH Returns**	
Parameters	Estimate	Standard Error	Estimate	Standard Error
r	1.370E-04		1.370E-04	
λ	0.0313	0.0129	0.0312	0.0121
β_0	1.111E-04	9.9781E-06	1.516E-06	6.050E-07
β_1			0.8916	0.0274
β_2			0.0617	0.0154
β_3			0.7422	0.0808

Properties	**Independent Returns**	**GARCH Returns**
Log-Likelihood	20,615.00	21,586.28
Volatility Persistence	0	0.9873
Annual Volatility	0.1673	0.1734
Conditional Skewness	-1.2105	-0.4127
Conditional Kurtosis	27.3304	3.4935

Notes: We use quasi maximum likelihood to estimate an independent return and a GARCH return model on daily S&P500 returns from January 2, 1980 to December 30, 2005 for a total of 6,564 observations. We report various properties of the two models including conditional skewness and excess kurtosis which are later used as parameter estimates in the SVG models.