



American option pricing under GARCH with non-normal innovations

Jean-Guy Simonato¹

Received: 20 June 2018 / Revised: 5 January 2019 / Accepted: 5 January 2019 /
Published online: 14 January 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

As it is well known from the time-series literature, GARCH processes with non-normal shocks provide better descriptions of stock returns than GARCH processes with normal shocks. However, in the derivatives literature, American option pricing algorithms under GARCH are typically designed to deal with normal shocks. We thus develop here an approach capable of pricing American options with non-normal shocks. The approach uses an equilibrium pricing model with shocks characterized by a Johnson S_u distribution and a simple algorithm inspired from the quadrature approaches recently proposed in the option pricing literature. Numerical experiments calibrated to stock index return data show that this method provides accurate option prices under GARCH for non-normal and normal cases.

Keywords American options · GARCH · Johnson distribution · Quadrature

JEL Classification C63 · G13

1 Introduction

The variance clustering phenomenon observed in stock returns is a well-known empirical regularity. ARCH and GARCH processes introduced in Engle (1982) and Bollerslev (1986), have now become standard models capable of capturing this critical characteristic. Another salient feature of stock returns is the non-normality of the innovations filtered with these models. For example, as shown in Bollerslev (1987), Hsieh (1989) and Hansen (1994), the estimation of various GARCH specifications on stock returns with distributions allowing for skewness and kurtosis different from those of a normal distribution provide more adequate descriptions of the data dynamics. Building on these empirical observations, the option pricing literature has

✉ Jean-Guy Simonato
jean-guy.simonato@hec.ca

¹ Department of Finance, HEC Montréal, 3000 Chemin-Côte-Sainte-Catherine, Montréal, QC, Canada

recently examined models incorporating these two key characteristics.¹ In this literature, most of the papers examine European-style options. However, the majority of exchange-traded options are of American-style and require numerical algorithms that can handle the early exercise features associated with these securities.

In the GARCH option pricing literature, Stentoft (2008) examines American option pricing with non-normal shocks. Using least-squares Monte Carlo simulations, he computes American option values with shocks parameterized by a normal Inverse Gaussian distribution. Although the least-squares Monte Carlo approach represents a viable and flexible pricing tool, as outlined in Glasserman (2004), such an algorithm is not free of concerns. For example, the accuracy of the computed prices depends on the choice of the basis functions representing the independent variables in the regressions. Also, in finite samples, because these algorithms produce a suboptimal early exercise policy, they provide a lower bound of the correct option prices. A numerical method less prone to such caveats would be a welcomed addition to the financial analyst toolbox. With this objective in mind, we develop here a pricing algorithm for options under GARCH with non-normal shocks inspired by the recent literature on quadrature option pricing algorithms.

This literature, developed in the early 2000s, consists in a series of papers which sequentially solved the different problems associated with the application of quadrature techniques to option pricing. In a nutshell, these quadrature approaches compute option values recursively from the maturity of the instrument on a grid of stock prices. At each time step, the early exercise and continuation values are computed and compared to obtain the option value for a stock price on the grid. A critical step in this recursive procedure is the computation of the continuation value. In this step, the integral associated with the expected value is computed using quadrature techniques such as the trapezoid, Simpson or Gauss–Legendre. In Sullivan (2000), a first attempt was put forward with Gaussian quadratures and function approximations using Chebyshev polynomials. In Andricopoulos et al. (2003), the quadrature approach is put on more solid grounds with a more widely applicable algorithm which bypassed the need of Chebyshev polynomials. For underlying securities processes with transition densities available in closed form, their algorithm allows fast and accurate computations of option prices for path dependent cases on one underlying asset, together with American-style options. In Andricopoulos et al. (2007), the method developed in their earlier paper is extended to cover more complex and challenging problems in option valuations involving one or more underlying. More recently, Chen et al. (2014) lift the requirement of underlying securities with closed-form transition densities by using precise density approximations, which results in a genuinely universal method capable of pricing derivative securities in a wide variety of contexts. Other contributions to this literature include Chung et al. (2010) who apply a fast quadrature engine to the algorithm of Andricopoulos et al. (2003). Finally, Cosma et al. (2016), Simonato (2016), and Su et al. (2017) show how the quadrature approaches can be used with a time-homogenous and equally spaced

¹ See, for example, the recent survey in Christoffersen et al. (2013).

price grid to provide substantial savings in computational efforts without losses in precision.

Here, we contribute to the literature on American options in the GARCH context by showing how a quadrature approach can be developed to handle GARCH processes with non-normal shocks. Such an extension is non-trivial as it requires a different pricing and distributional framework. For this purpose, we use the equilibrium pricing model proposed in Duan (1999) coupled with random shocks characterized by a Johnson S_u distribution (Johnson 1949). As discussed in Simonato and Stentoft (2015), Johnson S_u error terms fit the data well and offer a range of attainable skewness and kurtosis similar to that of commonly used alternatives such as the non-central t . Importantly, in the context of the equilibrium pricing model used here, they represent a natural choice for which it becomes possible to obtain the distribution of the random shocks under the risk-neutral measure. We also contribute to the quadrature literature by extending the simplified quadrature approach examined in Simonato (2016) to the more complex GARCH case. While the typical quadrature scheme uses a time-varying grid, the procedure examined here uses an evenly spaced, time-homogenous price grid which results in a simple and efficient algorithm.

Compared to the other numerical techniques that have been developed to compute American options under GARCH for the Normal case, the algorithm presented here shows important differences. A first difference is the use of the density function instead of the distribution function. For example, in Duan and Simonato (2001) and Ben-Ameur et al. (2009), the repeated use of the Gaussian distribution function is a crucial requirement. Here, as in other quadrature schemes, the use of the density function leads to a more flexible algorithm since these are typically easier to compute than distribution functions which must be approximated numerically, even in the simplest case. Second, the algorithm we propose is also different from the lattice techniques developed in Ritchken and Trevor (1999), Cakici and Topyan (2000) and Lyuu and Wu (2005). Here the actual dynamics of the process is used directly while these lattices only provide an approximate GARCH dynamics. Third, the algorithm is also different from the quadrature algorithm of Chung et al. (2010) which can only handle the Gaussian GARCH case, and that must use interpolations in the variance dimension. Their procedure is also more complex and computationally involved than the one we propose as it requires recomputing the densities at every time steps. Finally, we note that the algorithm we develop shows some similarities to the quadrature based approaches independently developed in Cosma et al. (2016) and Simonato (2016). Although they do not tackle the GARCH case, their algorithms also use a time-homogenous grid of prices in the quadrature context. The algorithm we present extends these papers by developing a simple and efficient implementation strategy designed explicitly for the GARCH case which brings numerous advantages regarding computing speed and memory requirements.

The remainder of the paper is as follows. The physical and risk-neutral processes are described in Sect. 2. Section 3 presents the quadrature algorithm developed here while Sect. 4 discuss some key implementation issues. Section 5 examines numerical results under the Gaussian and Johnson S_u cases. In that section, we show the relevance of our non-normal time series model with a maximum likelihood analysis.

Using the parameter estimates from this analysis, we examine the performance of the pricing algorithm for the Gaussian and non-Gaussian case. We then compare the results of our algorithm with those from the Markov chain approach of Duan and Simonato (2001). Section 6 concludes.

2 The physical and risk neutral processes

We assume the following dynamics for the stock price under the physical probability measure:

$$p_{t+1} - p_t = \alpha - \gamma_{t+1} + \sigma_{t+1}\varepsilon_{t+1}, \quad (1)$$

$$\sigma_{t+1}^2 = \beta_0 + \beta_1\sigma_t^2 + \beta_2\sigma_t^2(\varepsilon_t - \theta)^2, \quad (2)$$

$$\varepsilon_{t+1} | \mathcal{F}_t \sim J_{S_u}(a, b) \quad (3)$$

where \mathcal{F}_t is the set containing all information up to and including time t , $p_{t+1} = \ln P_{t+1}$ is the log of the stock price at $t + 1$, and σ_{t+1}^2 is the known conditional variance of the continuously compounded return. Here, α is the expected return over the next period while γ_{t+1} is mean correction factor defined as $e^{\gamma_{t+1}} = E_t(e^{\sigma_{t+1}\varepsilon_{t+1}})$ which implies $E_t[P_{t+1}] = P_t e^\alpha$. The conditional variance follows a standard NGARCH(1,1) process with the usual parameter interpretation (see for example Duan and Simonato 2001). This GARCH specification is used here as an example. Other GARCH specifications are possible with only minor changes. The non-normal random shocks $\varepsilon_{t+1} | \mathcal{F}_t$ are assumed to be distributed according to a Johnson S_u distribution (introduced in Johnson 1949), with parameters a and b , denoted as $J_{S_u}(a, b)$. Formally, ε_{t+1} is given by

$$\varepsilon_{t+1} = c + d \times \sinh\left(\frac{z_{t+1} - a}{b}\right), \quad (4)$$

where z_{t+1} is a standard normal random variable, and c and d are location and scale parameters that are functions of a and b . Computing these parameters can be done with

$$c = -M(a, b) / \sqrt{V(a, b)} \quad \text{and} \quad d = 1 / \sqrt{V(a, b)}$$

where $M(\cdot)$ and $V(\cdot)$ are simple functions described in ‘‘Appendix 1’’. Such shocks have a mean of zero and a standard deviation of one. The skewness and kurtosis of these shocks are different from the normal case, and the values taken by a and b jointly determine the magnitude of these moments. A positive (negative) value of a induces a negative (positive) skewness while smaller (larger) values of b lead to larger (smaller) kurtosis values. However, these relationships are not independent since the skewness and kurtosis are simultaneous functions of a and b .

As shown in Simonato and Stentoft (2015), using the above physical process with the equilibrium pricing model of Duan (1999) and the assumption of a constant

pricing parameter λ , we obtain the following risk-neutral system which can be used to compute the prices of derivative securities:

$$p_{t+1}^* - p_t^* = \phi(\sigma_{t+1}^*) + \sigma_{t+1}^* \varepsilon_{t+1}^*, \tag{5}$$

$$\sigma_{t+1}^{*2} = \beta_0 + \beta_1 \sigma_t^{*2} + \beta_2 \sigma_t^{*2} (\varepsilon_t^* - \theta)^2, \tag{6}$$

$$\varepsilon_t^* = c + d \times \sinh\left(\frac{z_t^* - a^*}{b}\right), \tag{7}$$

and

$$\phi(\sigma_{t+1}^*) = r - \ln\left[1 + \sum_{i=1}^4 \sigma_{t+1}^{*i} \mu_i^*\right] \tag{8}$$

where p_{t+1}^* is the log of the stock price under the risk-neutral measure \mathcal{Q} , σ_t^* is the risk-neutral volatility, r is the periodic continuously compounded risk-free rate and z_t^* is a standard Gaussian random variable under measure \mathcal{Q} . Here, λ is a constant pricing parameter which can be estimated using a maximum likelihood approach as in Stentoft (2008) and Simonato and Stentoft (2015). Although this parameter can be time-varying as in Duan (1999), Simonato and Stentoft (2015) show that option pricing with a time-varying or constant pricing parameter obtain very similar prices.² With this pricing parameter, the risk-neutral random shock ε_t^* becomes a four parameter Johnson S_u random variable, with parameters a^* , b , c and d with moments $\mu_i^* = E_t^{\mathcal{Q}}[\varepsilon_{t+1}^{*i}]$ and $a^* = a + \lambda$. These expected values are available in closed form and are described in ‘‘Appendix 1’’. We note that the above risk-neutral process uses an approximation which is required to compute some key expected values. The approximation uses a fourth order Taylor series which is, as shown in Simonato and Stentoft (2015), precise enough for all practical purposes.

3 American option prices: a quadrature algorithm

For the purpose of pricing American options in the above context, we use the following well-known recursive dynamic programming formulation:

$$v_t(p_t^*, \sigma_{t+1}^{*2}) = \max\left[g(e^{p_t^*}, K), e^{-r} E_t^{\mathcal{Q}}\{v_{t+1}(p_{t+1}^*, \sigma_{t+2}^{*2})\}\right] \tag{9}$$

with

$$g(e^{p_t^*}, K) = \max[\rho(e^{p_t^*} - K), 0]$$

² Simonato and Stentoft (2015) also find that the equilibrium approach with a constant or time-varying pricing parameter obtains option prices very close to those obtained by the no-arbitrage pricing approach developed in Christoffersen et al. (2010).

where $v_t(p_t^*, \sigma_{t+1}^{2*})$ is the option price at t , which is a function of the current log of stock price level, and the known variance level σ_{t+1}^{2*} . Inside the bracket of Eq. (9), the discounted expected value denotes the continuation value, while $g(e^{p_t^*}, K)$ is the payoff function of the option with strike price K when $t = T$, or the early exercise value when $t < T$, with $\rho = 1$ for a call option and $\rho = -1$ for a put. Given the risk-neutral process specified in Sect. 2, the continuation value can be written more explicitly as:

$$\begin{aligned}
 & e^{-r} E_t^Q \{ v_{t+1}(p_{t+1}^*, \sigma_{t+2}^{2*}) \} \\
 & = e^{-r} \int_{-\infty}^{+\infty} v_{t+1}(p_{t+1}^*, \sigma_{t+2}^{2*}) \times f(p_{t+1}^* | p_t^*, \sigma_{t+1}^{2*}) \times dp_{t+1}^*
 \end{aligned}
 \tag{10}$$

where $f(p_{t+1}^* | p_t^*, \sigma_{t+1}^{2*})$ is the risk-neutral conditional density of the log of stock price next period, given the current price and volatility level. In the context of Johnson S_u error terms, this density function is given by Eq. (20) in ‘‘Appendix 1’’. The next subsection examines how a simple quadrature algorithm can be used to compute this dynamic programming problem on a grid of log of stock prices and variances.

3.1 The quadrature algorithm

Denote the time-homogenous grid of m equally spaced log of stock prices as $\bar{\mathbf{p}} = [\bar{p}_1, \bar{p}_2, \dots, \bar{p}_m]'$ with m an odd number and the currently observed log of the stock price as $\bar{p}_{\frac{m+1}{2}} = p_0$. The time-homogeneity will allow substantial computational savings since the transition densities required in the computations will be the same at all time steps, unlike the typical quadrature approaches. The prices on the grid must be equally spaced. As seen in a later section, substantial savings in memory will be achieved with such a spacing.

A time-homogenous grid of n values for the variance, denoted as $\bar{\sigma}^2 = [\bar{\sigma}_1^2, \bar{\sigma}_2^2, \dots, \bar{\sigma}_n^2]'$, is also required. As for the stock price grid, one of the elements of the variance grid must be set equal to the known variance level prevailing at $t = 0$. Unlike the log of stock price, this grid need not be equally spaced. Details about how to build these grids are in the next section. Using the Cartesian product of sets $\bar{\mathbf{p}}$ and $\bar{\sigma}^2$, we obtain a two-dimensional grid with $m \times n$ points, whose elements (i, j) represent the log of the stock price level in state i and the variance level in state j . A bar over a letter denotes a quantity associated with a point on this grid.

As noticed in Broadie and Detemple (1996), when performing numerical computations, one period before maturity, the continuation value of an American option can be replaced by a European option value. For the GARCH case, the option has a constant variance at this time step. A quasi-analytical formula can be derived and used to compute the European option price at time $T - 1$ on a point (i, j) on the grid. More specifically, we compute

$$\bar{v}_{i,j,T-1} = \max \left[\rho \left(e^{\bar{p}_i} - K \right), A \left(\bar{p}_i, \bar{\sigma}_j, \tau, K, r \right) \right]
 \tag{11}$$

for $i = 1$ to m and $j = 1$ to n , where $\bar{v}_{i,j,T-1}$ is the option price in state (i, j) and time $T - 1$, $A(\cdot)$ is a generic function denoting the quasi-analytical formula, $\tau = 1$ is the maturity of the option in number of periods, with r and $\bar{\sigma}_j$ being periodic values. “Appendix 2” provides the quasi-analytical formula to compute the option value for the model presented in the above section. This formula is quasi-analytical since it involves an integral over a standard normal which cannot be computed analytically. As indicated in the appendix, a Gauss–Legendre approach can be used to compute this integral in a fraction of a second with high precision.

For time-steps $T - 2$ to 0 , the discounted continuation value given by Eq. (10) is computed numerically using a trapezoid rule (see for example Judd (1998)).³ For the case of a GARCH process, when performing these computations, care must be taken to pick the prices in the appropriate variance state. “Appendix 3” shows that, given a stock price and variance in state (i, j) , the continuation values can be computed recursively with the following formula:

$$\bar{c}_{i,j,t} = e^{-r} \sum_{k=1}^m \bar{v}_{k,l_{i \rightarrow k},t+1} \times f(\bar{p}_k \mid \bar{p}_i, \bar{\sigma}_j^2) \times w_k \tag{12}$$

for $i = 1$ to m and $j = 1$ to n . Here $l_{i \rightarrow k}$ represents the index value of the variance on the k -th line of the grid which is the nearest to a quantity that we describe as the *inherited* variance level. We denote this quantity as $\hat{\sigma}^2$ (notice the “hat”). We require this concept since, for GARCH processes, the variance next period is an inherited value, i.e., a function of the current shock to the stock price. This inherited variance is computed as

$$\hat{\sigma}^2 = \beta_0 + \beta_1 \bar{\sigma}_j^2 + \beta_2 \bar{\sigma}_j^2 (\hat{\varepsilon}_{i \rightarrow k} - \theta - \lambda)^2 \tag{13}$$

where

$$\hat{\varepsilon}_{i \rightarrow k} = \frac{\bar{p}_k - \bar{p}_i - \phi(\bar{\sigma}_j)}{\bar{\sigma}_j} \tag{14}$$

is the random shock that would induce a transition from state i to state k . The notation used to describe the index $l_{i \rightarrow k}$, although a bit tedious, emphasizes that this value is a function of the prices transiting from state i to state k . This notation will prove useful in the next section where an efficient algorithm for the identification of these $l_{i \rightarrow k}$ indexes will be proposed. This algorithm will allow substantial savings in memory and computing efforts. The Johnson density value is denoted by $f(\bar{p}_k \mid \bar{p}_i, \bar{\sigma}_j^2)$ and can be computed with Eq. (20) from “Appendix 1”, and w_j is the weight associated to the trapezoid integration scheme. Denoting by $\Delta \bar{p}$ the

³ Other quadrature rules could also be used with minimal changes. For the GARCH context examined here, it is found that the simple trapezoid rule works slightly better than the more sophisticated Simpson rule.

difference between the equidistant points on the grid of log of stock prices, these weights are computed as $w_1 = w_m = \Delta\bar{p}/2$ and $w_j = \Delta\bar{p}$ otherwise. With the above continuation values, we compute the option prices on the grid with

$$\bar{v}_{i,j,t} = \max \left(\rho \left(e^{\bar{p}_i} - K \right), \bar{c}_{i,j,t} \right) \tag{15}$$

for $i = 1$ to m and $j = 1$ to n . At time $t = 0$, the American option price is $\bar{v}_{\frac{m+1}{2},j_0,0}$ where j_0 is the index value of the variance grid which was set equal to the known variance at $t = 0$.

3.2 Computing the grid

In order to compute the grid of stock prices and variances, we require an interval of representative values at the maturity of the option. For example, in the Gaussian GARCH case examined by Duan and Simonato (2001), these grids are obtained using the known standard deviations of the log of stock prices and variances at maturity. In the case of non-Gaussian shocks, there are no available formulas for such quantities. We therefore rely on a Monte Carlo simulation to obtain the required values. More specifically, for the log of stock price, the minimum and maximum values are computed as

$$\bar{p}_1 = p_0 - q \times \text{std} \left(\{ \tilde{p}_{j,T} \}_{j=1}^M \right) \text{ and } \bar{p}_m = p_0 + q \times \text{std} \left(\{ \tilde{p}_{j,T} \}_{j=1}^M \right) \tag{16}$$

where $\text{std} \left(\{ \tilde{p}_{j,T} \}_{j=1}^M \right)$ denotes the standard deviation of a sample of M simulated risk-neutral paths of log of stock prices at T , and q is a multiplication factor for the standard deviations. We compute the equally spaced grid in the log of stock price dimension with

$$\bar{p}_i = \bar{p}_1 + (i - 1)\delta_p$$

for $i = 1$ to m and $\delta_p = (\bar{p}_m - \bar{p}_1)/(m - 1)$. For the variance, we also rely on simulated values to determine the interval computed as

$$\ln \left(\bar{\sigma}_1^2 \right) = \min \left(\left\{ \ln \left(\tilde{\sigma}_{j,T}^2 \right) \right\}_{j=1}^M \right) \text{ and } \ln \left(\bar{\sigma}_n^2 \right) = \max \left(\left\{ \ln \left(\tilde{\sigma}_{j,T}^2 \right) \right\}_{j=1}^M \right)$$

where $\left\{ \ln \left(\tilde{\sigma}_{j,T}^2 \right) \right\}_{j=1}^M$ denotes a sample of M simulated risk-neutral paths of log of variances at T . We use the maximum and minimum values since the distribution is much less symmetric than the one of the log from the stock price. We then compute the grid with

$$\ln \left(\bar{\sigma}_i^2 \right) = \ln \left(\bar{\sigma}_1^2 \right) + (i - 1)\delta_\sigma$$

for $i = 1$ to n and $\delta_\sigma = \left(\ln \left(\bar{\sigma}_n^2 \right) - \ln \left(\bar{\sigma}_1^2 \right) \right)/(n - 1)$. These log values are then converted into levels, resulting in an unevenly spaced grid in the variance dimension.

For options under GARCH, the initial variance (the variance next period) is one of the parameters required to compute the price. Because this variance is most probably not a part of this grid, we fix the element of the grid which is the closest to the initial variance, equal to this initial variance.

As it is the case with many numerical algorithms, the design of an efficient implementation procedure is a crucial step. The next section provides the details about how the continuation values can be efficiently computed by exploiting some regularities that are inherent to the above algorithm.

4 Implementation

The time-homogeneous and equally spaced grid of stock prices can bring significant reductions in computing efforts by exploiting regularities associated with the density values. Using the simplified representation explained below, “Appendix 4” presents a compact and efficient algorithm for computing the $m \times 1$ vector of continuation values in a variance state $\bar{\sigma}_j^2$. The algorithm allows for the following advantages in computation and storage when compared to a naive implementation: (1) computation and storage of $(2m - 1) \times n$ density values instead of $(m \times m) \times n$ values; (2) computation and storage of $(2m - 1) \times n$ index values of inherited variance level positions, instead of $(m \times m) \times n$ index values; (3) avoids multiplications by negligible density values. The fundamental regularities leading to the above advantages are explained below in the rest of this section. Readers wishing to skip this section can jump directly to Sect. 5 which presents numerical results about the performance of the algorithm.

A first set of critical elements leading to significant simplification are the conditional density values. Given a state of variance $\bar{\sigma}_j^2$, the conditional density values required to compute the continuation value in Eq. (12) are the same for all time-steps, unlike the typical quadrature scheme for which it will be time-varying. Hence, avoiding the repeated computation of this set of quantities results in substantial computational savings. We can store these values in a $m \times m$ matrix written

$$\mathbf{F}(\bar{\sigma}_j^2) = \begin{bmatrix} f(\bar{p}_1 | \bar{p}_1, \bar{\sigma}_j^2) & \cdots & f(\bar{p}_m | \bar{p}_1, \bar{\sigma}_j^2) \\ \vdots & \ddots & \vdots \\ f(\bar{p}_1 | \bar{p}_m, \bar{\sigma}_j^2) & \cdots & f(\bar{p}_m | \bar{p}_m, \bar{\sigma}_j^2) \end{bmatrix}.$$

It is important to notice that, because the points on the log of stock price grid are equally spaced, we can obtain all the elements of matrix $\mathbf{F}(\bar{\sigma}_j^2)$ that are above the diagonal from the elements of the last line. Conversely, we can obtain all the elements below the diagonal from the elements of the last line. All the elements on the diagonal are equal to $f(\bar{p}_1 | \bar{p}_1, \bar{\sigma}_j^2)$. These regularities can be summarized more formally as:

$$F_{i,k}(\bar{\sigma}_j^2) = \begin{cases} f(\bar{p}_{k-i+1} \mid \bar{p}_1, \bar{\sigma}_j^2) & \text{if } k \geq i, \\ f(\bar{p}_{i-k+1} \mid \bar{p}_m, \bar{\sigma}_j^2) & \text{if } k < i, \end{cases}$$

where $F_{i,k}(\bar{\sigma}_j^2)$ is element (i, k) of matrix $\mathbf{F}(\bar{\sigma}_j^2)$. Given such a structure, we reduce the storage and computational costs from $(m \times m) \times n$ elements, to $(m \times 2 - 1) \times n$ elements. Another important point to notice is about the magnitude of the elements. The elements on the first line are monotonically decreasing towards zero, while the elements on the last line are monotonically increasing towards $f(\bar{p}_m \mid \bar{p}_m, \bar{\sigma}_j^2)$. Using a threshold value below which the density value is considered too small to be relevant, we can easily predict which elements need to be multiplied when performing the expected value computations.

A second set of elements leading to important simplifications is the set of indexes of nearest variance values. These are required to pick the appropriate variance state when computing the continuation values. Similarly to the conditional density values, these index can be stored in a matrix depicted as

$$\mathbf{L}(\bar{\sigma}_j^2) = \begin{bmatrix} l_{1 \rightarrow 1} & \cdots & l_{1 \rightarrow m} \\ \vdots & \ddots & \vdots \\ l_{m \rightarrow 1} & \cdots & l_{m \rightarrow m} \end{bmatrix}$$

where $l_{i \rightarrow j}$ is the index value of the variance on line i of the grid which is the nearest to $\hat{\sigma}^2$, the inherited variance level determined by the transition of the stock price from state i to j . We observe that all the elements of this matrix of indexes can also be obtained from the first and last line using the relations summarized as:

$$L_{i,k}(\bar{\sigma}_j^2) = \begin{cases} l_{1 \rightarrow (k-i+1)} & \text{if } k \geq i, \\ l_{m \rightarrow (i-k+1)} & \text{if } k < i, \end{cases}$$

where $L_{i,k}(\bar{\sigma}_j^2)$ is element i, k of matrix $\mathbf{L}(\bar{\sigma}_j^2)$. Again, the storage and computational costs can be reduced from $(m \times m) \times n$ elements, to $(2m - 1) \times n$ elements. All in all, these above compact representations for the transition densities and nearest variance indexes represent a reduction in storage and computation of $(m \times m) \times 2n$ elements, to $(2m - 1) \times 2n$ elements.

The simple and compact algorithm presented in ‘‘Appendix 4’’ uses the above simplifications and monotonicity properties of density values to computes a $m \times 1$ vector of non-discounted continuation values for a given variance state j . Discounting the continuation values given in output by this algorithm, a vector of option prices in variance state j at time t can then be computed with Eq. (15).

5 Numerical results

In this section, we first present some maximum likelihood analysis of the NGARCH model with the assumption of a constant pricing parameter under the Gaussian and Johnson distribution hypothesis. The numerical performance of the algorithm for computing the values of European and American style options is then examined using the estimated parameters.

5.1 Parameter estimates

As in Stentoft (2008) and Simonato and Stentoft (2015), under the assumption of a constant pricing parameter, the inputs required to compute the option prices can be obtained from the following system:

$$R_{t+1} = g(\lambda) + \sigma_{t+1} \varepsilon_{t+1}, \quad (17)$$

$$\sigma_{t+1}^2 = \beta_0 + \beta_1 \sigma_t^2 + \beta_2 \sigma_t^2 (\varepsilon_t - \theta)^2, \quad (18)$$

where

$$g(\lambda) = \begin{cases} r + \lambda \sigma_{t+1} - \frac{1}{2} \sigma_{t+1}^2 & \text{if } \varepsilon_{t+1} \sim \text{Normal}(0, 1) \\ \phi(\sigma_{t+1}^*) & \text{if } \varepsilon_{t+1} \sim J_{s_u}(a, b) \end{cases}.$$

For the $\text{Normal}(0, 1)$ case, which corresponds to the model in Duan (1995), we obtain the likelihood function from the Gaussian density. For the $J_{s_u}(a, b)$ case, we use the Johnson S_u density given by Eq. (20). Table 1 presents the parameters obtained by maximizing the log-likelihood function for both cases with daily S&P 500 value-weighted returns (including dividends) from CRSP for the beginning of 1990 to the end of 2015. For both cases, the estimated parameters are all significant and very close to each other. Such a result is anticipated since the maximum likelihood estimator of the Gaussian density can be seen as a quasi-maximum likelihood estimator capable of estimating the parameters of a GARCH process with non-Gaussian shocks. Although the two models are non-nested, it is interesting to notice that the likelihood value obtained with the Johnson S_u density shows a much higher value with only two additional parameters, which are both statistically significant. Since a Johnson S_u random variable is built from an invertible function of a standard normal random variable, using the filtered ε_t from the MLE approach, it is possible to obtain the corresponding standard normal shocks. A normality test about these standard shocks can then be used to assess the relevance of the assumed distribution. The Jarque–Berra test shows that the hypothesis of the standard normal residual at the source of the Johnson shocks cannot be rejected. A visual inspection with a quantile-to-quantile plot also confirms this result.

Table 1 Parameter estimates

	Gaussian	Johnson S_u
β_0	1.9e-06 (1.5e-07)	1.6e-06 (1.8e-07)
β_1	0.8314 (0.0074)	0.8267 (0.0109)
β_2	0.0720 (0.0047)	0.0722 (0.0066)
θ	1.0749 (0.0719)	1.1351 (0.0994)
λ	0.0308 (0.0126)	0.0314 (0.0125)
a	– –	0.3783 (0.0801)
b	– –	2.1625 (0.1219)
Loglik	21,550	21,661
JB	662.6670 [0.0000]	0.2779 [0.8900]
$Q(20)$	32.0892 [0.0424]	31.6514 [0.0472]
$Q^2(20)$	23.4866 [0.2655]	21.7662 [0.3533]

This table presents the maximum likelihood parameter estimates for the physical NGARCH processes with Gaussian and Johnson S_u shocks under the assumption of a constant pricing parameter. Data: S&P 500 daily log returns from January 1990 to December 2015 and a constant risk-free rate set to $0.03 \times 1/252$. Standard errors are reported in parenthesis. “Loglik” is the log-likelihood value and “JB” is the Jarque–Bera normality test for the standard normal residuals. $Q(20)$ is the Ljung–Box portmanteau test for up to 20th-order serial correlation in the standardized residuals, whereas $Q^2(20)$ is the same test for the squared standardized residuals. The p-values for these tests are reported in square brackets

5.2 Option prices

In this section, we examine the performance of the algorithm for pricing European and American put options. For the European case, we benchmark our analysis with precise Monte Carlo simulated prices. For the American case, there are no benchmarks for GARCH with non-normal shocks. Nevertheless, the convergence results obtained for European options should translate to the American case.

Although the goal of the paper is to examine American option pricing in the non-normal GARCH context, it is interesting to first examine the performance of the algorithm in the more straightforward Gaussian case. For this case, we use the quadrature algorithm of Sect. 3.1 implemented with the analytical formula at the penultimate time step given by a Black–Scholes price and the normal density. More specifically, we compute the analytical formula in Eq. (11) with

$$A(\bar{p}_i, \bar{\sigma}_j, \tau, K, r) = BSput(\exp(\bar{p}_i), \bar{\sigma}_j, 1, K, r)$$

where $BSput(\cdot)$ is the Black and Scholes (1973) formula for a put option computed with a daily variance and interest rate.

For the subsequent computations with Eq. (12), we use the normal density

$$f(\bar{p}_k | \bar{p}_i, \bar{\sigma}_j^2) = \frac{1}{\sqrt{2\pi\bar{\sigma}_j}} e^{-\frac{1}{2} \frac{(\bar{p}_k - \mu_{ij})^2}{\bar{\sigma}_j^2}}$$

where $\mu_{ij} = \bar{p}_i + r - \frac{1}{2}\bar{\sigma}_j^2$. In these computations, it is important to notice that the weights that are applied to the value function in Eq. (12) are the product of a density and a trapezoid weight $\Delta\bar{p}$. We use a target value for this product in order to determine the threshold value used in the algorithm presented in ‘‘Appendix 4’’ and below which the density value is considered negligible. More specifically, we compute the algorithm threshold with

$$\text{algorithm threshold} = (1/\Delta\bar{p}) \times \text{target threshold}.$$

For all the results presented in this paper, we use a target threshold of 1×10^{-7} . The use of a smaller threshold produces results that are virtually identical. Finally, to determine the minimum and maximum values of the log of stock prices, we set $q = 20$ in Eq. (16), i.e., an interval with 20 standard deviations for the log of stock prices.

Table 2 looks at European and American option prices replicating the examples found in Duan and Simonato (2001) and Ben-Ameur et al. (2009). The parameters from these examples present a low unconditional level and a low persistence level of the variance process. The persistence level indicates how close is the variance process from being non-stationary (integrated). A persistence smaller (larger) than one indicates a stationary (non-stationary) variance process. A persistence level smaller but very close to one indicates a case with high persistence. In the GARCH context, this is an important feature since such variance processes can generate larger variance levels, which in turn can affect the precision of the numerical approach. For this example, we obtain an unconditional variance and a persistence level of 0.200 and 0.909 for the physical process, and 0.221 and 0.925 for the risk-neutral process.⁴ For the European option case, the benchmarks are accurate Monte Carlo European put prices computed with 10 million sample paths and a Black–Scholes control variate. As shown by the numbers in the table, as we increase the values of m and n , the quadrature prices become remarkably close to the benchmarks, with differences at the fourth decimal place. Shorter maturities converge more quickly since the range

⁴ GARCH parameters are typically obtained with time series with 252 trading days. Hence, for the physical process, we compute the annual unconditional volatility as $\sqrt{\beta_0/(1 - \beta_1 - \beta_2(1 + \theta^2))} \times 252$. The persistence level is given by $\beta_1 + \beta_2(1 + \theta^2)$. The corresponding values for the risk-neutral case are obtained with the same formula but with θ replaced by $\theta + \lambda$.

Table 2 Option prices in the Gaussian GARCH context: a low variance persistence case

<i>m, n</i>	<i>T</i> = 30 days			<i>T</i> = 90 days			<i>T</i> = 270 days		
	<i>K</i> = 45	<i>K</i> = 50	<i>K</i> = 55	<i>K</i> = 45	<i>K</i> = 50	<i>K</i> = 55	<i>K</i> = 45	<i>K</i> = 50	<i>K</i> = 55
<i>Monte Carlo European put prices</i>									
	0.0773	1.0879	4.8396	0.4152	1.8220	4.9543	1.1953	2.8430	5.4780
<i>Quadrature European put prices</i>									
201, 201	0.0773	1.0877	4.8417	64.6811	6.06e+03	8.38e+04	8.71e+69	3.08e+76	1.01e+78
401, 401	0.0773	1.0876	4.8394	0.4155	1.8221	4.9545	9.19e+02	5.56e+03	2.24e+04
601, 601	0.0774	1.0880	4.8396	0.4154	1.8221	4.9541	1.2041	2.8683	5.5328
801, 801	0.0774	1.0880	4.8396	0.4154	1.8220	4.9540	1.1957	2.8437	5.4785
1001, 1001	0.0774	1.0878	4.8395	0.4155	1.8220	4.9541	1.1954	2.8433	5.4781
1501, 1501	0.0773	1.0878	4.8395	0.4155	1.8222	4.9541	1.1955	2.8434	5.4781
<i>Benchmark American put prices</i>									
	0.0779	1.0992	5.0000	0.4236	1.8727	5.1816	1.2613	3.0397	5.9627
<i>Quadrature American put prices</i>									
201, 201	0.0776	1.0990	5.0000	64.8401	6.08e+03	8.40e+04	8.71e+69	3.08e+76	1.01e+78
401, 401	0.0776	1.0989	5.0000	0.4230	1.8727	5.1818	9.22e+02	5.57e+03	2.25e+04
601, 601	0.0777	1.0993	5.0000	0.4229	1.8727	5.1817	1.2660	3.0538	5.9875
801, 801	0.0777	1.0993	5.0000	0.4229	1.8726	5.1816	1.2604	3.0393	5.9631
1001, 1001	0.0777	1.0992	5.0000	0.4230	1.8726	5.1817	1.2601	3.0389	5.9627
1501, 1501	0.0777	1.0991	5.0000	0.4230	1.8728	5.1817	1.2602	3.0390	5.9628

This table presents the computed prices for European and American put options in the Gaussian GARCH context for a low variance persistence case. Parameters: $P_0 = 50, \sigma_1 = 0.0105, r = 1.3698 \times 10^{-4}, \beta_0 = 1.0 \times 10^{-5}, \beta_1 = 0.80, \beta_2 = 0.10, \theta = 0.30, \lambda = 0.20$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. “Quadrature European put prices” and “Quadrature American put prices” are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 20$. “Benchmark American put prices” are prices taken from Ben-Ameur et al. (2009)

of possible values is smaller. We observe a similar behavior for the American put prices with only minimal differences from the benchmark values.

The second example presented in Table 3 also examines a Gaussian case. It uses the maximum likelihood parameters obtained for the normal case in Table 1 with an initial variance equal to the unconditional level of the physical variance process. This case presents a higher unconditional level and a higher persistence of the variance process when compared to the one examined in the previous table. The annualized unconditional physical variance and the persistence level are computed as 0.237 and 0.991 for the risk-neutral process. Table 3 shows results which are qualitatively similar to those we find in the previous table, with a small loss in precision for larger maturities. We attribute these differences to the higher persistence level of the risk-neutral variance process.

The first panel of Table 4 presents the computing times (in seconds) for American options computed with the parameters from the second example and various grid sizes. These computing times for this example are typically larger than those of the

Table 3 Option prices in the Gaussian GARCH context: a high variance persistence case

m, n	$T = 21$ days			$T = 63$ days			$T = 189$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.0871	1.0018	4.8759	0.4620	1.6427	4.8190	1.2746	2.6613	5.2027
<i>Quadrature European put prices</i>									
201, 201	0.0873	1.0057	4.8781	3.2842	5.91e+04	8.98e+06	2.75e+65	1.53e+79	6.72e+80
401, 401	0.0872	1.0032	4.8760	0.4563	1.6399	4.8347	1.07e+24	9.68e+28	1.63e+31
601, 601	0.0871	1.0012	4.8758	0.4595	1.6366	4.8159	34.8098	3.21e+02	1.15e+04
801, 801	0.0870	1.0023	4.8759	0.4620	1.6435	4.8191	1.2989	2.7247	5.4020
1001, 1001	0.0870	1.0013	4.8758	0.4612	1.6412	4.8181	1.2648	2.6496	5.1979
1501, 1501	0.0872	1.0021	4.8759	0.4619	1.6418	4.8183	1.2699	2.6557	5.1967
2001, 2001	0.0871	1.0018	4.8758	0.4618	1.6418	4.8181	1.2726	2.6598	5.2004
<i>Quadrature American put prices</i>									
201, 201	0.0874	1.0104	5.0000	3.2874	5.91e+04	8.99e+06	2.77e+65	1.53e+79	6.73e+80
401, 401	0.0873	1.0079	5.0000	0.4588	1.6612	5.0206	1.07e+24	9.69e+28	1.63e+31
601, 601	0.0872	1.0059	5.0000	0.4620	1.6575	5.0167	34.9019	3.22e+02	1.16e+04
801, 801	0.0871	1.0070	5.0000	0.4645	1.6643	5.0179	1.3311	2.8145	5.6415
1001, 1001	0.0870	1.0060	5.0000	0.4637	1.6620	5.0175	1.2980	2.7500	5.5187
1501, 1501	0.0872	1.0068	5.0000	0.4644	1.6626	5.0175	1.3035	2.7569	5.5201
2001, 2001	0.0872	1.0065	5.0000	0.4643	1.6626	5.0174	1.3062	2.7609	5.5231

This table presents the computed prices for European and American put options in the Gaussian GARCH context for a high variance persistence case with parameters obtained from the MLE estimates of Table 1. Parameters: $P_0 = 50, \sigma_1 = 0.0119, r = 0.03 \times 1/252$ and GARCH parameters from the Gaussian estimates of Table 1. "Monte Carlo European put prices" are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. "Quadrature European put prices" and "Quadrature American put prices" are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 20$

low persistence level example. The reported times do not include the time to compute the Monte Carlo inputs (which is the same for all cases with similar maturities). The prices are computed with Matlab and compiled C coded functions automatically generated with the Matlab Coder. For values of m smaller or equal to 1001, the computing times are quite low and range from a few seconds to a couple of minutes for longer maturities. These times grow quickly for larger values of m and n and can become excessive for long maturities and large values of m and n such as 2, 001. Fortunately, these large computation times can be cut down by computing the prices with $n < m$. This is possible since the magnitude of daily variances are much smaller (and easier to cover with a smaller grid) than the magnitude of log of prices. Values of n that are smaller than m by a factor of three or four provide answers converging to the benchmarks. Table 4 shows the computation times that are obtained with $n \approx m/3$, while Table 5 presents the option prices computed with this modification. Significant gains are made in computing efforts, with few losses in terms of precision, as the computing times are (roughly) divided by three.

Table 4 Computation times

	$T = 21$	$T = 63$	$T = 189$
$n = m$			
$m = 201$	0.6	1.1	2.8
$m = 401$	3	8	22
$m = 601$	10	27	73
$m = 801$	23	63	175
$m = 1001$	45	125	344
$m = 1501$	152	423	1173
$m = 2001$	362	1008	2778
$n \approx m/3$			
$m = 201$	0.1	0.4	1.0
$m = 401$	1	2	7
$m = 601$	3	8	24
$m = 801$	7	20	56
$m = 1001$	13	38	107
$m = 1501$	43	127	361
$m = 2001$	101	300	852

This table presents the computing times (in seconds) for at-the-money American put options in the Gaussian GARCH context for the high variance persistence case

The third and fourth examples are presented in Tables 6 and 7. These tables examine the Johnson S_u cases for the low and high variance persistence level. The computations are implemented with the algorithm described in Sect. 3. Table 6 uses the same parameters as those in Table 2 with the additional Johnson parameters set to $a = 0.5$ and $b = 2.5$. Such parameter values generate shocks with skewness and kurtosis of -0.27 and 3.93 . The physical unconditional variance and persistence level for the physical variance process are thus identical to those of the example in Table 2.⁵ As in the Gaussian case, we see that the convergence towards the Monte Carlo benchmark is very good with perhaps slightly larger pricing errors for deep in-the-money options. For the American case, there are no benchmark values, but as in the Gaussian case, the numbers stabilize nicely at different values of m and n , depending on the maturity. Table 7 examines the case where the parameters are obtained from the MLE analysis. For larger maturities, convergence is slower but the computed prices also reach the benchmark values with a good precision, again more quickly for smaller than longer maturities.

For the Johnson case, computing times are typically much larger than the Gaussian case. These larger computing times are caused by the larger probabilities of transiting to prices that are far away because of the negative skewness and

⁵ Unlike for the Gaussian case, there are no available formulas to compute the unconditional variance and persistence level for the risk-neutral process.

Table 5 Option prices in the Gaussian GARCH context: a high variance persistence case

m, n	$T = 21$ days			$T = 63$ days			$T = 189$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.0871	1.0018	4.8759	0.4620	1.6427	4.8190	1.2746	2.6613	5.2027
<i>Quadrature European put prices for $n \approx m/3$</i>									
201, 67	0.0734	0.9382	4.8775	0.5698	2.6520	10.7648	3.63e+23	1.82e+26	2.49e+27
401, 133	0.0829	0.9914	4.8754	0.4588	1.6007	4.8195	2.28e+13	2.07e+16	1.60e+18
601, 201	0.0865	1.0014	4.8759	0.4536	1.6085	4.8040	2.3395	6.5813	23.9652
801, 267	0.0870	1.0015	4.8758	0.4788	1.6858	4.8398	1.3307	2.7632	5.4021
1001, 333	0.0865	1.0001	4.8759	0.4636	1.6455	4.8204	1.2794	2.6904	5.2413
1501, 501	0.0870	1.0017	4.8758	0.4642	1.6475	4.8215	1.2669	2.6520	5.1936
2001, 667	0.0870	1.0015	4.8758	0.4603	1.6376	4.8158	1.2947	2.6907	5.2280
2501, 833	0.0871	1.0014	4.8758	0.4620	1.6429	4.8188	1.2838	2.6754	5.2140

This table presents the computed prices for European put options in the Gaussian GARCH context for a high variance persistence case. Parameters: $P_0 = 50, \sigma_1 = 0.0105, r = 1.3698 \times 10^{-4}, \beta_0 = 1.0 \times 10^{-5}, \beta_1 = 0.80, \beta_2 = 0.10, \theta = 0.30, \lambda = 0.20$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. “Quadrature European put prices” and “Quadrature American put prices” are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 20$

larger kurtosis values. Hence, for a same target threshold value, there are more computations to be performed by the algorithm in the Johnson case. The first panel of Table 8 reports the computing times that are achieved for at-the-money American put options using the parameters obtained from the MLE analysis presented above. As in the second Gaussian example, this case presents a higher unconditional level and a higher-persistence of the variance process. Again, it is possible to reduce the computing times to much lower levels by reducing the number of grid points in the variance dimension, as indicated in the second of the table. Table 9 presents the quadrature prices for this case with $n \approx m/3$. Again, the convergence to the Monte Carlo benchmarks is very good, with some discrepancies for the larger maturity and deep in-the-money options.

5.3 Markov-chain option prices

As stated in the introduction, pricing algorithms for American options in the Gaussian GARCH cases are available in the literature. In the GARCH model we consider, the Johnson S_{it} error terms are continuous and monotone transformations of standard

Table 6 Option prices in the Johnson S_u GARCH context: a low variance persistence case

m, n	$T = 30$ days			$T = 90$ days			$T = 270$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.1062	1.1132	4.8358	0.4893	1.8896	4.9735	1.3392	2.9916	5.5933
<i>Quadrature European put prices</i>									
201, 201	0.1122	1.3347	6.3094	1.38e+12	1.75e+15	3.63e+16	9.34e+108	4.14e+113	1.32e+115
401, 401	0.1063	1.1132	4.8208	0.5611	2.2968	6.2805	1.20e+33	5.17e+35	5.82e+36
601, 601	0.1061	1.1121	4.8214	0.4894	1.8864	4.9654	1.05e+02	3.44e+02	8.53e+02
801, 801	0.1063	1.1130	4.8216	0.4898	1.8905	4.9834	1.4440	3.2307	6.0315
1001, 1001	0.1063	1.1129	4.8215	0.4898	1.8905	4.9840	1.3193	2.9405	5.4828
1501, 1501	0.1063	1.1130	4.8215	0.4897	1.8904	4.9838	1.3401	2.9930	5.5927
2001, 2001	0.1062	1.1128	4.8214	0.4898	1.8905	4.9839	1.3402	2.9935	5.5936
<i>Quadrature American put prices</i>									
201, 201	0.1126	1.3424	6.3618	1.38e+12	1.76e+15	3.70e+16	9.34e+108	4.14e+113	1.35e+115
401, 401	0.1067	1.1249	5.0000	0.5662	2.3180	6.3313	1.20e+33	5.19e+35	5.89e+36
601, 601	0.1065	1.1238	5.0000	0.4984	1.9386	5.1889	1.05e+02	3.46e+02	8.58e+02
801, 801	0.1067	1.1247	5.0000	0.4986	1.9412	5.1965	1.4952	3.3642	6.3261
1001, 1001	0.1067	1.1246	5.0000	0.4986	1.9412	5.1967	1.3982	3.1625	6.0204
1501, 1501	0.1067	1.1247	5.0000	0.4986	1.9411	5.1966	1.4122	3.1932	6.0695
2001, 2001	0.1067	1.1245	5.0000	0.4986	1.9412	5.1967	1.4123	3.1935	6.0700

This table presents the computed prices for European and American put options in the Johnson S_u GARCH context for a low variance persistence case. Parameters: $P_0 = 50, \sigma_1 = 0.0105, r = 0.05 \times 1/365, \beta_0 = 1.0 \times 10^{-5}, \beta_1 = 0.80, \beta_2 = 0.10, \theta = 0.30, \lambda = 0.20, a = 0.5$ and $b = 2.5$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. “Quadrature European put prices” and “Quadrature American put prices” are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 25$

normal error terms, making it possible to adapt methods designed for the Gaussian case to this non-Gaussian framework. In this section, we examine how the Markov-chain approach of Duan and Simonato (2001) can be adapted to price American options under the risk-neutral dynamics examined in the above sections.

The idea behind the construction of an approximating Markov chain to compute options under GARCH is as follows. The real line, the support of the log of the stock price distribution, is first partitioned into m distinct cells. We build a similar partition for the log of the variance with n cells. We then assign a numerical value to each cell. These values yield an $m \times n$ grid of log of stock prices and variances. Given the values on point i, j of this grid, the probability of reaching a pair of price and volatility cells can then be computed analytically, using the distribution function of the random shocks associated to the stock price. Since in our case the error term is a Johnson S_u random variable, we can compute its distribution function with a standard normal distribution function. With these transition probabilities, it then becomes possible to compute the conditional expected values appearing in the Belman recursion.

Table 7 Option prices in the Johnson S_u GARCH context: a high variance persistence case

m, n	$T = 21$ days			$T = 63$ days			$T = 189$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.1690	1.1703	4.8949	0.6764	1.8898	4.9156	1.6168	2.9755	5.3977
<i>Quadrature European put prices</i>									
401, 401	0.1684	1.1713	4.9044	0.8674	17.2871	1.58e+04	1.08e+86	2.08e+92	1.04e+94
601, 601	0.1684	1.1695	4.9028	0.6883	1.9293	5.3756	6.62e+55	1.31e+61	1.44e+63
801, 801	0.1689	1.1699	4.9029	0.6757	1.8884	4.9375	2.22e+35	1.38e+39	1.26e+42
1001, 1001	0.1692	1.1711	4.9031	0.6742	1.8860	4.9115	2.19e+20	5.02e+22	1.17e+26
1501, 1501	0.1691	1.1707	4.9030	0.6753	1.8872	4.9100	16.1569	77.6308	9.25e+02
2001, 2001	0.1688	1.1699	4.9028	0.6753	1.8876	4.9101	1.6594	3.1226	5.9117
2501, 2501	0.1690	1.1703	4.9028	0.6755	1.8876	4.9099	1.6145	2.9824	5.4359
3001, 3001	0.1689	1.1703	4.9028	0.6755	1.8881	4.9104	1.6126	2.9733	5.3980
<i>Quadrature American put prices</i>									
401, 401	0.1685	1.1754	5.0000	0.8706	17.3207	1.59e+04	1.08e+86	2.08e+92	1.06e+94
601, 601	0.1686	1.1735	5.0000	0.6915	1.9470	5.4555	6.62e+55	1.31e+61	1.47e+63
801, 801	0.1690	1.1740	5.0000	0.6788	1.9069	5.0814	2.22e+35	1.38e+39	1.26e+42
1001, 1001	0.1693	1.1752	5.0000	0.6773	1.9046	5.0720	2.19e+20	5.03e+22	1.17e+26
1501, 1501	0.1692	1.1748	5.0000	0.6784	1.9058	5.0716	16.2027	77.8807	9.30e+02
2001, 2001	0.1690	1.1740	5.0000	0.6784	1.9062	5.0717	1.6918	3.1922	6.0522
2501, 2501	0.1691	1.1744	5.0000	0.6787	1.9062	5.0716	1.6517	3.0781	5.7107
3001, 3001	0.1691	1.1743	5.0000	0.6787	1.9068	5.0719	1.6502	3.0718	5.6914

This table presents the computed prices for European and American put options in the Johnson S_u GARCH context for a high variance persistence case. Parameters: $P_0 = 50, \sigma_1 = 0.0141, r = 0.03 \times 1/252, \beta_0 = 1.60e-06, \beta_1 = 0.8267, \beta_2 = 0.0722, \theta = 1.1351, \lambda = 0.0314, a = 0.3783$ and $b = 2.1625$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black-Scholes control variate. “Quadrature European put prices” and “Quadrature American put prices” are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 25$

As shown by this short description, the Markov chain and the simplified quadrature approach share many similarities. They both use a grid of stock prices and variances, and the option prices are computed recursively on each point of the grid. The main difference between the Markov chain approach and the simplified quadrature scheme is about the computation of the conditional expected value in the Bellman recursion. While the quadrature approach uses a numerical integration scheme such as the trapeze or Simpson rule, the Markov chain approach computes the conditional expected values using the transition probabilities of the approximating Markov chain. More specifically, the continuation values given by Eq. (12) for the quadrature case are computed as follows with the Markov chain approach:

$$\bar{c}_{i,j,t} = e^{-r} \sum_{k=1}^m \bar{v}_{k,t_i \rightarrow k,t+1} \times \Pr(\bar{p}_k | \bar{p}_i, \bar{\sigma}_j^2)$$

Table 8 Computation times

	$T = 21$	$T = 63$	$T = 189$
$n = m$			
$m = 201$	3.6	2.6	5.8
$m = 401$	9	19	43
$m = 601$	25	61	140
$m = 801$	56	135	327
$m = 1001$	103	263	643
$m = 1501$	316	839	2069
$m = 2001$	738	1944	4818
$n \approx m/3$			
$m = 201$	0.4	0.9	1.9
$m = 401$	2	6	14
$m = 601$	7	19	45
$m = 801$	17	42	104
$m = 1001$	30	80	198
$m = 1501$	94	260	653
$m = 2001$	215	600	1513

This table presents the computing times (in seconds) for at-the-money American put options in the Johnson GARCH context for the high variance persistence case

where $\Pr(\bar{p}_k | \bar{p}_i, \bar{\sigma}_j^2)$ is the transition probability of going from state $\{\bar{p}_i, \bar{\sigma}_j^2\}$ to the cells associated with state $\{\bar{p}_k, \bar{\sigma}_{i \rightarrow k}^2\}$.⁶ “Appendix 5” describes how these transition probabilities can be computed. Duan and Simonato (2001) store these probabilities into a square matrix of dimension $mn \times mn$, and prices are computed backward with simple matrix multiplications. Sparse matrices techniques are then used to save on storage and computation time. To compare the Markov chain and simplified quadrature approaches, we implement the Markov chain with the procedure outlined in Sect. 4. As in Simonato (2011), where the Markov chain approach is examined for the geometric Brownian motion and jump cases, we use a simple algorithm similar to the one presented in “Appendix 4” to save on computing effort and memory.

Table 10 presents the prices computed with the Markov chain algorithm for the Johnson S_u case with a high variance persistence level. To enable valid comparisons, we implement the Markov chain with the same grid as the one used for the simplified quadrature approach. Overall, for the small and medium maturity cases, the prices computed with the Markov chain converge to the target prices, but more slowly, and with more oscillation than those obtained with the simplified quadrature scheme in Table 7. The algorithm also encounters difficulties for the long maturity case of 189 days. While the quadrature algorithm was able to achieve accurate results with m and n

⁶ As explained in Sect. 3.1 and “Appendix 3”, the “ $i \rightarrow k$ ” notation used here emphasizes that the next period variance is an inherited value which depends on \bar{p}_i, \bar{p}_k and $\bar{\sigma}_j^2$.

Table 9 Option prices in the Johnson S_u GARCH context: a high variance persistence case

m, n	$T = 21$ days			$T = 63$ days			$T = 189$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.1690	1.1703	4.8949	0.6764	1.8898	4.9156	1.6168	2.9755	5.3977
<i>Quadrature European put prices $n \approx m/3$</i>									
201, 67	0.1233	1.1850	7.0099	1.05e+02	1.03e+03	5.72e+03	9.25e+78	2.86e+82	5.70e+83
401, 133	0.1739	1.1696	4.9037	0.9446	14.3651	6.89e+02	5.94e+69	3.66e+74	1.36e+76
601, 201	0.1658	1.1548	4.9006	0.7166	2.0200	5.1973	3.13e+46	6.99e+50	6.05e+52
801, 267	0.1706	1.1822	4.9045	0.6526	1.8613	4.9293	5.13e+28	5.08e+31	4.68e+34
1001, 333	0.1683	1.1677	4.9027	0.6846	1.8951	4.9126	3.94e+16	5.12e+18	4.06e+21
1501, 501	0.1694	1.1723	4.9033	0.6773	1.8966	4.9183	9.3642	38.8882	3.50e+02
2001, 667	0.1686	1.1685	4.9027	0.6773	1.8931	4.9137	1.6577	3.1159	5.9199
2501, 833	0.1691	1.1711	4.9030	0.6768	1.8895	4.9113	1.6274	2.9985	5.4497
3001, 1001	0.1688	1.1696	4.9027	0.6746	1.8860	4.9092	1.6164	2.9766	5.3985

This table presents the computed prices for European and American put options in the Johnson S_u GARCH context for a high variance persistence case. Parameters: $P_0 = 50, \sigma_1 = 0.0141, r = 0.03 \times 1/252, \beta_0 = 1.60e-06, \beta_1 = 0.8267, \beta_2 = 0.0722, \theta = 1.1351, \lambda = 0.0314, a = 0.3783$ and $b = 2.1625$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. “Quadrature European put prices” and “Quadrature American put prices” are option prices computed with the quadrature approach implemented with the algorithm described in Sect. 4, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 25$

around 2500, the Markov chain algorithm fails to converge to the benchmarks. Finally, we notice that computing times with the Markov chain are similar to those obtained for the quadrature.

6 Conclusion

GARCH models with non-normal errors are now commonly used by analysts for modeling financial returns. The choice of algorithms to compute American option prices in this context is however limited. We have proposed a simple quadrature approach to compute American and European option values in this context. The algorithm is shown to converge well to benchmark values obtained from accurate Monte Carlo simulations. Computing times can be long, but effective strategies are available to temper down this disadvantage with small impacts on the precision. Finally, we also show that, within the pricing context presented here, it is possible to adapt the Markov chain approach of Duan and Simonato (2001) which was explicitly designed to deal with Gaussian

Table 10 Option prices in the Johnson S_u GARCH context with a Markov chain algorithm: a high variance persistence case

m, n	$T = 21$ days			$T = 63$ days			$T = 189$ days		
	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$	$K = 45$	$K = 50$	$K = 55$
<i>Monte Carlo European put prices</i>									
	0.1690	1.1703	4.8949	0.6764	1.8898	4.9156	1.6168	2.9755	5.3977
<i>Markov chain European put prices</i>									
401, 401	0.1441	1.0849	4.8853	0.3649	1.5950	5.0346	3.9148	6.3824	9.5337
601, 601	0.1545	1.1188	4.8944	0.3907	1.4535	4.8141	2.1272	4.2684	7.3972
801, 801	0.1604	1.1397	4.8984	0.4546	1.4979	4.7572	1.3392	3.2131	6.3231
1001, 1001	0.1637	1.1514	4.9001	0.5134	1.5885	4.7622	0.9618	2.6199	5.6975
1501, 1501	0.1665	1.1617	4.9017	0.5930	1.7324	4.8175	0.6980	2.0040	4.9366
2001, 2001	0.1672	1.1643	4.9020	0.6268	1.7962	4.8526	0.7612	1.9213	4.6764
2501, 2501	0.1681	1.1670	4.9023	0.6433	1.8267	4.8709	0.8949	2.0243	4.6450
3001, 3001	0.1682	1.1679	4.9024	0.6530	1.8456	4.8829	1.0360	2.1880	4.7240
<i>Markov chain American put prices</i>									
401, 401	0.1443	1.0899	5.0000	0.3679	1.6155	5.1179	3.9286	6.4038	9.5638
601, 601	0.1546	1.1233	5.0000	0.3935	1.4830	5.0152	2.1389	4.2904	7.4322
801, 801	0.1605	1.1439	5.0000	0.4572	1.5264	5.0082	1.3513	3.2406	6.3730
1001, 1001	0.1638	1.1556	5.0000	0.5160	1.6135	5.0122	0.9764	2.6613	5.7925
1501, 1501	0.1666	1.1658	5.0000	0.5956	1.7531	5.0293	0.7213	2.0941	5.2443
2001, 2001	0.1673	1.1684	5.0000	0.6294	1.8157	5.0430	0.7904	2.0397	5.1435
2501, 2501	0.1682	1.1710	5.0000	0.6458	1.8456	5.0512	0.9263	2.1474	5.1510
3001, 3001	0.1684	1.1719	5.0000	0.6555	1.8642	5.0570	1.0678	2.3056	5.2011

This table presents the computed prices for European and American put options in the Johnson S_u GARCH context for a high variance persistence case. Parameters: $P_0 = 50$, $\sigma_1 = 0.0141$, $r = 0.03 \times 1/252$, $\beta_0 = 1.60e-06$, $\beta_1 = 0.8267$, $\beta_2 = 0.0722$, $\theta = 1.1351$, $\lambda = 0.0314$, $a = 0.3783$ and $b = 2.1625$. “Monte Carlo European put prices” are Monte Carlo option prices computed with 10 million sample paths and a Black–Scholes control variate. “Markov chain European put prices” and “Markov chain American put prices” are option prices computed with the Markov chain approach implemented with the algorithm described in Sect. 5.3, and minimum and maximum values of the grids determined with a Monte Carlo simulation with 500,000 paths and $q = 25$

GARCH cases. Numerical results show that the quadrature approach we propose here is more efficient than this Markov chain approach.

Acknowledgements This research was supported by the Social Sciences and Humanities Research Council of Canada. I would like to thank Lynn Mailloux for helpful comments

Appendix 1: Moments of Johnson S_u random variables

This appendix describes in more details the Johnson S_u random variables that have been introduced in Johnson (1949).

A zero-mean, unit variance random variable following a two-parameter Johnson S_u distribution with parameters a and b is written

$$\varepsilon = c + d \times x \tag{19}$$

where $x = \sinh\left(\frac{z-a}{b}\right)$, z is a standard normal random variable and $c = -M(a, b) / \sqrt{V(a, b)}$, $d = 1 / \sqrt{V(a, b)}$ with

$$M(a, b) = -w^{\frac{1}{2}} \sinh(\Omega),$$

$$V(a, b) = \frac{1}{2}(w - 1)(w \cosh(2\Omega) + 1),$$

and $w = e^{\frac{1}{b^2}}$ and $\Omega = \frac{a}{b}$. The parameters a and b control the skewness and kurtosis. The functions $\cosh(u) = (e^u + e^{-u})/2$ and $\sinh(u) = (e^u - e^{-u})/2$ are the hyperbolic cosine and sine functions. Here $M(\cdot)$ and $V(\cdot)$ are the mean and variance of x_t , the unstandardized two-parameter Johnson S_u random variable. The values c and d , which are functions of a and b , are used to change the location and scale of x in such a way that it becomes a standardized Johnson S_u random variable.

The first four moments of the standardized Johnson S_u random shock can be computed with:

$$E[\varepsilon^1] = c + dE[x],$$

$$E[\varepsilon^2] = c^2 + 2cdE[x] + d^2E[x^2],$$

$$E[\varepsilon^3] = c^3 + 3c^2dE[x] + 3cd^2E[x^2] + d^3E[x^3],$$

$$E[\varepsilon^4] = c^4 + 4c^3dE[x] + 6c^2d^2E[x^2] + 4cd^3E[x^3] + d^4E[x^4],$$

where the first four moments of x can be derived to be:

$$E[x^1] = -w^{\frac{1}{2}} \sinh(\Omega),$$

$$E[x^2] = \frac{1}{2} [w^2 \cosh(2\Omega) - 1],$$

$$E[x^3] = \frac{1}{4} [3w^{\frac{1}{2}} \sinh(\Omega) - w^{\frac{9}{2}} \sinh(3\Omega)],$$

$$E[x^4] = \frac{1}{8} [w^8 \cosh(4\Omega) - 4w^2 \cosh(2\Omega) + 3].$$

Because Johnson error terms are a monotone, continuous and invertible transformation of a standard normal, the conditional density of the stock price can be obtained from the standard normal density with a change of variable (see Simonato 2012). The conditional density of the risk-neutral price described by Eq. (5) can be written as:

$$f(p_{t+1}^* | p_t^*, \sigma_{t+1}^*) = b \sqrt{\frac{V(a, b)}{2\pi}} e^{-\frac{1}{2} \left(\frac{(a+\lambda) + b \sinh^{-1} \left(\frac{M(a, b) + \varepsilon_{t+1} \sqrt{V(a, b)}}{\sqrt{V(a, b)}} \right)}{\sqrt{\left(M(a, b) + \varepsilon_{t+1}^* \sqrt{V(a, b)} \right)^2 + 1}} \right)^2} \tag{20}$$

where

$$\varepsilon_{t+1}^* = \frac{P_{t+1}^* - P_t^* - \phi(\sigma_{t+1}^*)}{\sigma_{t+1}^*}.$$

Appendix 2: Quasi-analytical option price for the one period case

For a one-period to maturity case, pricing options in the GARCH context amounts to a constant variance case since the variance is a predictable process, i.e., the variance next period is a known quantity. In such a context, it is thus possible to obtain a quasi-analytical option pricing formula for the GARCH model outlined in Sect. 2. The solution is quasi-analytic since it involves an integral over a standard normal random variable that must be solved numerically. This integral can be computed numerically in a fraction of a second using a precise Gauss–Legendre approach. This Appendix derives this quasi-analytical solution.

Denote the logarithm of the risk-neutral price for the next period of time as⁷

$$p = \phi + \varepsilon\sigma$$

where the expressions for ϕ , ε and σ are given in Sect. 2. The value of a put option with a one-period to maturity and a strike price K is

$$V = e^{-r}E[\max(K - P, 0)] = e^{-r} \int_0^K (K - P) \times f(P) \times dP.$$

Using the change of variable $\ln(P) = \phi + \varepsilon\sigma$, the above integral can be rewritten in terms of the Johnson S_u random shock as

$$V = e^{-r}K \int_{-\infty}^k f(\varepsilon) \times d\varepsilon - e^{-r} \int_{-\infty}^k e^{\phi+\varepsilon\sigma} \times f(\varepsilon) \times d\varepsilon$$

where $k = \frac{\ln K - \phi}{\sigma}$. Using another change of variable given by $\varepsilon = c + d \times \sinh\left(\frac{z-a}{b}\right)$, we rewrite the above expression as a function of standard normal random shock given by

$$V = e^{-r}K \times N(q) - e^{-r} \int_{-\infty}^q e^{\phi+(c+d \times \sinh\left(\frac{z-a}{b}\right))\sigma} \times f(z) \times dz$$

where $N(\cdot)$ is the distribution function of a standard normal random variable, and $q = a + b \sinh^{-1}\left(\frac{k-c}{d}\right)$. The integral in this expression can be computed very efficiently using a Gauss–Legendre numerical integration technique (see Judd 1998 for a detailed description). For example, using parameter values from the empirical

⁷ In this Appendix, without loss of generality, we drop the time subscript and the stars denoting risk neutral quantities.

section, the Gauss–Legendre approach with 15 integration nodes obtains prices in 0.0003 s on a standard laptop computer. These prices are identical, up to the fourth digit, to prices obtained with a precise Monte Carlo simulation with 100 million sample paths.

Appendix 3: Continuation values in the GARCH context

For time-steps $T - 2$ to 0, the discounted continuation value given by Eq. (10) is computed numerically on a point (i, j) of the grid using a trapezoid quadrature rule. For such computations, it is important to notice that, for GARCH processes, given that the stock price is presently in state i , a transition next period to a price in state k will automatically determine the variance level that will be inherited from the price movement. Such an inherited variance level is caused by the nature of the GARCH variance dynamics which specifies that next period variance is a function of the current shock leading to the stock price. More specifically, given that the stock price and variance are currently in state (i, j) , the shock that would be required to go next period to a price in state k can be computed from Eq. (5) with:

$$\hat{\varepsilon}_{i \rightarrow k} = \frac{\bar{p}_k - \bar{p}_i - \phi(\bar{\sigma}_j)}{\bar{\sigma}_j}.$$

This shock implies that the inherited variance level associated with price \bar{p}_k will be

$$\hat{\sigma}^2 = \beta_0 + \beta_1 \bar{\sigma}_j^2 + \beta_2 \bar{\sigma}_j^2 (\hat{\varepsilon}_{i \rightarrow k} - \theta - \lambda)^2.$$

Recall that for the k th price on the grid, there are n different variance levels. The inherited variance will almost never be equal to one of these points. We therefore choose the k th price with the variance level which is the closest to the inherited variance. Denoting by $l_{i \rightarrow k}$ the index of the closest variance level, the computation for the continuation value at time t on a point (i, j) of the grid can thus be written as:

$$\bar{c}_{i,j,t} = e^{-r} \sum_{k=1}^m \bar{v}_{k,l_{i \rightarrow k},t+1} \times f(\bar{p}_k | \bar{p}_i, \bar{\sigma}_j^2) \times w_k$$

where w_k is the trapezoid rule weight described in Sect. 3.

Appendix 4: Algorithm for continuation values computation

This appendix describes the pseudo-code which can be used to compute a vector of (non-discounted) continuation values in variance state j without the need to compute and store all the elements of matrix $\mathbf{F}(\bar{\sigma}_j)$. It also avoids multiplying densities that are too close to zero according to a pre-specified threshold value (1×10^{-10} for example).

The following inputs, which can be computed using the equations described in the main text, are required by the pseudo-code:

- F_{top} : an $m \times 1$ vector = $\left[f(\bar{p}_1 | \bar{p}_1, \bar{\sigma}_j^2), \dots, f(\bar{p}_m | \bar{p}_1, \bar{\sigma}_j^2) \right]'$
- F_{bot} : an $m \times 1$ vector = $\left[f(\bar{p}_1 | \bar{p}_m, \bar{\sigma}_j^2), \dots, f(\bar{p}_m | \bar{p}_m, \bar{\sigma}_j^2) \right]'$
- L_{top} : an $m \times 1$ vector = $[l_{1 \rightarrow 1}, \dots, l_{1 \rightarrow m}]'$
- L_{bot} : an $m \times 1$ vector = $[l_{m \rightarrow 1}, \dots, l_{m \rightarrow m}]'$
- $nzfirst$ and $nzlast$: the number of non-zero elements in F_{top} and F_{bot}
- $cbeg$: an $m \times 1$ vector with the i th element equal to $\max(i - nzlast + 1, 1)$
- $cend$: an $m \times 1$ vector with the i th element equal to $\min(i + nzfirst - 1, m)$
- $f11$: a scalar equal to $f(\bar{p}_1 | \bar{p}_1, \bar{\sigma}_j^2)$
- w : an $m \times 1$ vector of trapezoid rule weights

The pseudo-code below gives the output in the $m \times 1$ vector of non-discounted continuation values c :

```

%first line
for k=cbeg(1):cend(1)
    c(1)=c(1)+Ftop(k)*v(k,Ltop(k))*w(k)
end
%lines between first and last
for i=2:m-1
    for k=cbeg(i):i-1
        c(i)=c(i)+Fbot(m-i+k)*v(k,Lbot(m-i+k))*w(k)
    end
    c(i)=c(i)+f11*v(i,Ltop(1))*w(i)
    for k=i+1:cend(i)
        c(i)=c(i)+Ftop(k-i+1)*v(k,Ltop(k-i+1))*w(k)
    end
end
end
%last line
for k=cbeg(m):cend(m)
    c(m)=c(m)+Fbot(k)*v(k,Lbot(k))*w(k)
end

```

Appendix 5: Computing transition probabilities with the Markov-chain approach

The Markov chain algorithm is implemented as follows. First, an $m \times n$ grid of log stock prices and variances is built using the procedure described in Sect. 3.2. Using these points, the $m + 1$ boundaries of the intervals partitioning the support of the log of prices are build with

$$B_i = \frac{\bar{p}_i + \bar{p}_{i-1}}{2} \quad \text{for } i = 2 \text{ to } m$$

and $B_1 = -1.0 \times 10^6$ and $B_{m+1} = 1.0 \times 10^6$. With these, conditional on a point i, j on the grid, the probability of going from state $\{\bar{p}_i, \bar{\sigma}_j^2\}$ to the cells around state $\{\bar{p}_k, \bar{\sigma}_{i-k}^2\}$ is computed as

$$\Pr(\bar{p}_k | \bar{p}_i, \bar{\sigma}_j^2) = N(\bar{z}_{k+1}) - N(\bar{z}_k) \quad \text{for } k = 1 \text{ to } m$$

where $N(\cdot)$ is the distribution function of a standard normal random variable and

$$\bar{z}_k = a^* + b \times \sinh^{-1} \left(\frac{B_k - \bar{p}_i - \phi(\bar{\sigma}_j) - \bar{\sigma}_j c}{\bar{\sigma}_j d} \right).$$

References

- Andricopoulos A, Widdicks M, Newton D, Duck P (2007) Extending quadrature methods to value multi-asset and complex path dependent options. *J Financ Econ* 83:471–499
- Andricopoulos A, Widdicks M, Duck P, Newton D (2003) Universal option valuation using quadrature. *J Financ Econ* 67:447–471
- Ben-Ameur H, Breton M, Martinez JM (2009) Dynamic programming approach for valuing options in the GARCH Model. *Manag Sci* 55:252–266
- Black F, Scholes M (1973) The pricing of options and corporate liabilities. *J Polit Econ* 81:637–659
- Bollerslev T (1986) Generalized autoregressive conditional heteroskedasticity. *J Econ* 31:307–327
- Bollerslev T (1987) A Conditionally heteroskedastic time-series model for speculative prices and rates of return. *Rev Econ Stat* 69:542–547
- Broadie M, Detemple J (1996) American option valuation: new bounds, approximations, and a comparison of existing methods. *Rev Financ Stud* 9:1211–1250
- Cakici N, Topyan K (2000) The GARCH option pricing model: a lattice approach. *J Comput Finance* 3:71–85
- Chen D, Härkönen H, Newton D (2014) Advancing the universality of quadrature methods to any underlying process for option pricing. *J Financ Econ* 114:600–612
- Christoffersen P, Elkamhi R, Feunou B, Jacobs K (2010) Option valuation with conditional heteroskedasticity and non-normality. *Rev Financ Stud* 23:2139–2183
- Christoffersen P, Jacobs K, Ornthanalai C (2013) GARCH option valuation: theory and evidence. *J Deriv* 21:8–41
- Chung SL, Ko K, Shackleton M, Yeh CY (2010) Efficient quadrature and node positioning for exotic option valuation. *J Futures Mark* 30:1026–1057
- Cosma A, Galluccio S, Pederzoli P, Scaillet O (2016) Early exercise decision in american options with dividends, stochastic volatility and jumps, Swiss Finance Institute Research paper no. 16-73. <https://ssrn.com/abstract=2883202>
- Duan JC (1995) The GARCH option pricing model. *Math Finance* 5:13–32
- Duan JC (1999) Conditionally fat-tailed distributions and the volatility smile in options, working paper, Hong Kong University of Science and Technology
- Duan JC, Simonato JG (2001) American option pricing under GARCH by a Markov chain approximation. *J Econ Dyn Control* 25:1689–1718
- Engle R (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. *Econometrica* 50:987–1008
- Glasserman P (2004) *Monte Carlo Methods in Financial Engineering*. Springer, New York
- Hansen B (1994) Autoregressive conditional density estimation. *Int Econ Rev* 35:705–730
- Hsieh D (1989) Modeling heteroscedasticity in daily foreign-exchange rates. *J Bus Econ Stat* 7:307–317
- Johnson N (1949) Systems of frequency curves generated by methods of translation. *Biometrika* 36:149–176

- Judd K (1998) Numerical methods in economics. MIT Press, Cambridge
- Lyuu Y, Wu C (2005) On accurate and provably efficient GARCH option pricing algorithms. *Quant Finance* 5:181–198
- Ritchken P, Trevor R (1999) Pricing options under generalized GARCH and stochastic volatility processes. *J Finance* 54:377–402
- Simonato JG (2011) Computing American option prices in the lognormal jump-diffusion framework with a Markov chain. *Finance Res Lett* 8:220–226
- Simonato JG (2012) GARCH processes with skewed and leptokurtic innovations: revisiting the Johnson S_u case. *Finance Res Lett* 9:213–219
- Simonato JG (2016) A simplified quadrature approach for computing Bermudan option prices. *Int Rev Finance* 16:647–658
- Simonato JG, Stentoft L (2015) Which pricing approach for options under GARCH with non-normal innovations? CREATES research paper 2015-32
- Sullivan M (2000) Valuing American put options using Gaussian quadrature. *Rev Financ Stud* 13:75–94
- Stentoft L (2008) American option pricing using GARCH models and the Normal Inverse Gaussian distribution. *J Financ Econom* 6:540–582
- Su H, Chen D, Newton D (2017) Option pricing via QUAD: from Black–Scholes–Merton to Heston with jumps. *J Deriv* 24:9–27

Publisher Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.