

Recursive Marginal Quantization of Higher-Order Schemes

Thomas A. McWalter^{*1,2}, Jörg Kienitz^{1,3}, Eckhard Platen^{1,4} and Ralph Rudd¹

¹Department of Actuarial Science and the African Collaboration for Quantitative Finance and Risk Research, University of Cape Town

²Department of Finance and Investment Management, University of Johannesburg

³Fachbereich Mathematik und Naturwissenschaften, Bergische Universität Wuppertal

⁴School of Finance and Economics and School of Mathematical Sciences, University of Technology Sydney

December 10, 2016

Abstract

Recursive Marginal Quantization of the Euler scheme has recently been proposed by [Pagès and Sagna \[2015\]](#) as an efficient numerical method for evaluating functionals of solutions of stochastic differential equations. This method involves recursively quantizing the conditional marginals of the discrete-time Euler approximation of the underlying process. By generalizing their approach, we show that it is possible to perform recursive marginal quantization for two higher-order schemes: the Milstein scheme and a simplified weak order 2.0 scheme. As part of this generalization a simple matrix formulation is presented, allowing efficient implementation. We further extend the applicability of recursive marginal quantization by showing how absorption and reflection at the zero boundary may be incorporated, when this is necessary. To illustrate the improved accuracy of the higher order schemes, various computations are performed using geometric Brownian motion and its generalization, the constant elasticity of variance model. For both processes, we show numerical evidence of improved weak order convergence and we compare the marginal distributions implied by the three schemes to the known analytical distributions. By pricing European, Bermudan and Barrier options, further evidence of improved accuracy of the higher order schemes is demonstrated.

1 Introduction

Quantization is a lossy compression technique that produces a discrete representation of a signal using less information than the original. The technique originated in the field of signal compression, but has found application in fields as far-reaching as signal processing, pattern recognition, data mining, integration theory and, more recently, numerical probability. For general overviews of the mathematics and applications of quantization see [Du et al. \[1999\]](#) and [Pagès \[2014\]](#).

Vector quantization of probability distributions was formalized in the work of [Graf and Luschgy \[2000\]](#) and has been applied to the field of mathematical finance since its inception. It

^{*}tom@analytical.co.za

is a technique for optimally representing a continuous distribution by a discrete distribution, where a measure of the ‘distance’ between the two, called the distortion, is minimized. The distortion is most commonly measured using the squared Euclidean error.

The application of vector quantization to the solution of finance-related problems generally proceeds by discretising time and then quantizing the corresponding marginal distributions of the system of stochastic differential equations specific to the problem. The quantized grids and their associated weights are then used to compute the expectations required in pricing contingent claims (including claims with early exercise) or for performing the optimizations required in stochastic control problems [Pagès et al., 2004]. Due to the reliance on Lloyd’s Algorithm [Lloyd, 1982], or variants thereof, these approaches generally incur a heavy computational burden.

A more efficient approach for the single-factor case has recently been proposed by Pagès and Sagna [2015]. Known as recursive marginal quantization (RMQ), it makes use of a Newton-Raphson iteration to quantize the Euler-Maruyama updates of the underlying SDE. The technique has been used to provide fast calibration of a local volatility model by Callegaro et al. [2014, 2015a], and extended for use with two factor SDEs and applied to stochastic volatility models [Callegaro et al., 2015b].

In the present work, the RMQ algorithm is generalized, allowing the implementation of schemes of higher order than the Euler-Maruyama scheme. We now provide an overview of the rest of the paper.

Section Two provides a review of vector quantization (VQ) as applied to probability distributions. In particular, we strive to simultaneously provide a precise, concise and intuitive description of the methodology. The resulting algorithm is presented using a matrix formulation, allowing for efficient implementation. The section concludes by showing examples of vector quantization applied to the Gaussian and noncentral chi-squared distributions.

The third section introduces recursive marginal quantization applied to stochastic differential equations. Our formulation of the problem is presented in more generality than the original formulation by Pagès and Sagna [2015]. A matrix formulation, which demonstrates the connection with Markov chains, is provided, allowing easy and efficient implementation.

In the section that follows we extend the RMQ algorithm to higher-order updates, specifically the Milstein scheme and a simplified weak order 2.0 Taylor scheme of Kloeden and Platen [1999]. Geometric Brownian motion (GBM) and the CEV process serve as examples to illustrate the improved error in the quantized marginal distributions and the improved weak order convergence.

When performing a Monte Carlo simulation of a discrete-time approximation of a process, non-negativity of the solution is usually enforced by implementing absorption or reflection. Under certain circumstances this is also required when using RMQ. In Section Five we present the modifications of the RMQ algorithm necessary to ensure an absorbing or reflecting boundary at zero. These modifications allow the RMQ algorithm to be applied to the CEV process for parameter sets that would otherwise be problematic under the original formulation.

Section Six presents numerical results of the application of the three RMQ schemes to option pricing. European, barrier and Bermudan options are priced under the GBM and CEV models. Where possible, the results are compared to available closed-form solutions, otherwise they are compared to high-resolution finite difference or Monte Carlo methods.

One of the goals of this paper has been to produce a more general and easily accessible introduction to the theory of VQ and RMQ. The primary contribution is, however, the more general formulation of RMQ that enables the systematic extension of the work of Pagès and

Sagna [2015]. The paper concludes with a discussion of ongoing work.

2 Vector Quantization

Vector quantization is a lossy compression technique that provides a way to encode a vector space using a discrete subspace. While the technique is applicable more generally, we shall only consider the quantization of one dimensional distributions. The vector quantization problem we aim to address in this section may be specified intuitively as follows:

Find the discrete distribution that “best” represents the continuous distribution function associated with a random variable X .

This is depicted in the Figure 1, which shows a density function of a continuous random variable and its corresponding quantized version. Here, for ease of visualization, we have chosen to plot the probability density function (of the continuous random variable) and the probability mass function of the quantizer instead of the continuous and discrete distribution functions.

We now provide a more rigorous specification of the problem. Let X be a continuous random variable, taking values in \mathbb{R} , and defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The above question may be reposed as:

How does one optimally approximate X , in a least-squares sense, by a discrete random variable $\hat{X} : \Omega \rightarrow \Gamma$, where Γ is a finite set of elements in \mathbb{R} ?

The reason that quantization is useful is that it allows efficient approximations of expectations of functionals $H(X)$ of X , that is,

$$\mathbb{E}[H(X)] = \int_{\mathbb{R}} H(x) d\mathbb{P}(X \leq x) \approx \sum_{\gamma \in \Gamma} H(\gamma) \mathbb{P}(\hat{X} = \gamma).$$

Here, for example, H may be the discounted payoff of financial claim and \mathbb{P} the risk-neutral probability measure.

Consider the approximation of X given by \hat{X} , a discrete random vector defined as the nearest-neighbour projection of X onto $\Gamma = \{\gamma^1, \gamma^2, \dots, \gamma^N\}$, a set of distinct points in \mathbb{R} with finite cardinality $N \in \mathbb{N}^+$. We shall refer to Γ as the *quantizer* and its elements as *codewords*. The *nearest neighbor projection operator* $\pi_{\Gamma} : \mathbb{R} \rightarrow \Gamma$ is defined as

$$\pi_{\Gamma}(X) = \{\gamma^i \in \Gamma \mid \|X - \gamma^i\| \leq \|X - \gamma^j\| \text{ for } j = 1, \dots, N; \text{ where equality holds only for } i < j\}.$$




Figure 1: The continuous probability density function on the left is quantized on the right, with probabilities represented on the vertical axis.




Figure 2: Representation of the region $R_i(\Gamma)$ associated with codeword γ^i .

Here, $\|\cdot\|$ is the Euclidean norm. Associated with the quantizer, the *regions* $R_i(\Gamma) \subset \mathbb{R}$ are the subset of values of X that are mapped to each codeword γ^i :

$$R_i(\Gamma) = \{x \in \mathbb{R} \mid \pi_\Gamma(x) = \gamma^i\}.$$

These are also known as *Voronoi regions*. For the sake of brevity we shall use R_i to refer to $R_i(\Gamma)$ when it is clear that the quantizer we are referring to is Γ . The set of regions $\{R_i\}_{i=1}^N$ is called a *tessellation* of \mathbb{R} , and has the following properties:

$$R_i \cap R_j = \emptyset \text{ for } i \neq j \quad \text{and} \quad \cup_{i=1}^N R_i = \mathbb{R}.$$

Since we are working in one dimension, the regions R_i may be defined directly as $R_i = \{x \mid r^{i-} < x \leq r^{i+}\}$ with

$$r^{i-} = \frac{\gamma^{i-1} + \gamma^i}{2} \quad \text{and} \quad r^{i+} = \frac{\gamma^i + \gamma^{i+1}}{2},$$

where $r^{1-} = -\infty$ and $r^{N+} = \infty$. If the distribution under consideration is not defined over the whole real line, then r^{1-} and r^{N+} are adjusted to reflect the interval of support. Figure 2 shows a simple graphical representation of these regions.

With these definitions in place, we are now in a position to precisely define the optimization problem. We wish to find the quantizer Γ such that $\widehat{X} = \pi_\Gamma(X)$ best approximates X . The sense in which the quantizer Γ is “best” is specified by the *distortion function*

$$\begin{aligned} D(\Gamma) &= \mathbb{E} \left[\|X - \widehat{X}\|^2 \right] \\ &= \int_{\mathbb{R}} \|x - \pi_\Gamma(x)\|^2 d\mathbb{P}(X \leq x) \\ &= \sum_{i=1}^N \int_{R_i(\Gamma)} \|x - \gamma^i\|^2 d\mathbb{P}(X \leq x). \end{aligned} \tag{1}$$

We require the Γ that minimizes $D(\Gamma)$. The probability weights then follow directly as a result of the nearest neighbor projection operator, i.e., $\mathbb{P}(\widehat{X} = \gamma^i) = \mathbb{P}(X \in R_i)$. A necessary condition on the optimal Γ is that the gradient of the distortion function is zero, that is, $\nabla D(\Gamma) = \bar{0}$, where the elements of $\nabla D(\Gamma)$ are given by

$$\frac{\partial D(\Gamma)}{\partial \gamma^i} = 2 \left(\int_{R_i(\Gamma)} (\gamma^i - x) d\mathbb{P}(X \leq x) \right),$$

for $1 \leq i \leq N$. Intuitively, this means that the first moment conditioned on the outcomes of a region equals the respective codeword. Thus, one way to solve this system of equations is to set up a fixed-point iteration using the above expression for the gradient. This is the basis for Lloyd’s algorithm which starts with an initial guess for the quantizer, $\Gamma^{(0)}$, and generates

successive updates, $\Gamma^{(n+1)}$, with the new codewords, γ_{n+1}^i , computed as the centroids of the regions associated with the previous quantizer $\Gamma^{(n)}$:

$$\gamma_{n+1}^i = \frac{\int_{R_i(\Gamma^{(n)})} x d\mathbb{P}(X \leq x)}{\int_{R_i(\Gamma^{(n)})} d\mathbb{P}(X \leq x)},$$

for $1 \leq i \leq N$ and $0 \leq n < n_{\max}$.

Another approach is to represent the quantizer by the column vector $\mathbf{\Gamma}$, derive the Hessian, $\nabla^2 D(\mathbf{\Gamma})$, and compute the updated estimates of the quantizer using an iterative Newton-Raphson method

$$\mathbf{\Gamma}^{(n+1)} = \mathbf{\Gamma}^{(n)} - \left[\nabla^2 D(\mathbf{\Gamma}^{(n)}) \right]^{-1} \nabla D(\mathbf{\Gamma}^{(n)})$$

for $0 \leq n < n_{\max}$. We now develop this approach further.

Suppose f_X and F_X are the PDF and CDF of X . We define the p -th lower partial expectation as

$$M_X^p(x) = \mathbb{E} [X^p \mathbb{I}_{\{X < x\}}],$$

where $M_X^0(X) = F_X(x)$ represents the distribution function of X . Then, direct integration of the distortion function (1) gives

$$\begin{aligned} D(\mathbf{\Gamma}) &= \sum_{i=1}^N \int_{r^{i-}}^{r^{i+}} \|x - \gamma^i\|^2 f_X(x) dx \\ &= \sum_{i=1}^N \left[M_X^2(r^{i+}) - M_X^2(r^{i-}) - 2\gamma^i (M_X^1(r^{i+}) - M_X^1(r^{i-})) \right. \\ &\quad \left. + (\gamma^i)^2 (F_X(r^{i+}) - F_X(r^{i-})) \right]. \end{aligned}$$

Consequently, the elements of the vector $\nabla D(\mathbf{\Gamma})$ are given by

$$\frac{\partial D(\mathbf{\Gamma})}{\partial \gamma^i} = 2\gamma^i (F_X(r^{i+}) - F_X(r^{i-})) - 2 (M_X^1(r^{i+}) - M_X^1(r^{i-})),$$

for $1 \leq i \leq N$.

Similarly, the tridiagonal Hessian matrix, $\nabla^2 D(\mathbf{\Gamma})$, may be computed. It has diagonal elements given by

$$\frac{\partial^2 D(\mathbf{\Gamma})}{\partial (\gamma^i)^2} = 2 (F_X(r^{i+}) - F_X(r^{i-})) + \frac{1}{2} (f_X(r^{i+})(\gamma^i - \gamma^{i+1}) + f_X(r^{i-})(\gamma^{i-1} - \gamma^i)),$$

and super- and sub-diagonal elements given by

$$\frac{\partial^2 D(\mathbf{\Gamma})}{\partial \gamma^i \partial \gamma^{i+1}} = \frac{1}{2} f_X(r^{i+})(\gamma^i - \gamma^{i+1})$$

and

$$\frac{\partial^2 D(\mathbf{\Gamma})}{\partial \gamma^i \partial \gamma^{i-1}} = \frac{1}{2} f_X(r^{i-})(\gamma^{i-1} - \gamma^i).$$

Note that the quantities required to compute a Newton-Raphson iteration (i.e., the gradient and Hessian) only require the PDF, CDF and first lower partial expectation to be known. The second lower partial expectation is required only if one wishes to compute a numerical estimate of the distortion.

2.1 Efficient Implementation

We now provide a matrix formulation of the above Newton iteration intended to aid efficient implementation. As stated previously, the quantizer is represented by a column vector $\mathbf{\Gamma}$. This vector and three other column vectors required are defined by

$$\begin{aligned} [\mathbf{\Gamma}]_i &= \gamma^i, & [\mathbf{M}]_i &= M_X^1(r^{i+}) - M_X^1(r^{i-}), & 1 \leq i \leq N, \\ [\mathbf{f}]_i &= f_X(r^{i+}), & [\Delta\mathbf{\Gamma}]_i &= \gamma^{i+1} - \gamma^i, & 1 \leq i \leq N-1. \end{aligned}$$

Note that the last two vectors are one element shorter than the first two. The row vector of probabilities \mathbf{p} is defined as

$$\begin{aligned} [\mathbf{p}]_i &= \mathbb{P}(\widehat{X} = \gamma^i) \\ &= \mathbb{P}(X \in R_i(\mathbf{\Gamma})) \\ &= F_X(r^{i+}) - F_X(r^{i-}), & 1 \leq i \leq N. \end{aligned}$$

Defining \mathbf{p} as a row vector is convenient since the expectation of a functional H applied to the quantizer is

$$\mathbb{E}[H(X)] = \sum_{i=1}^N H(\gamma^i) \mathbb{P}(\widehat{X} = \gamma^i) = \mathbf{p}H(\mathbf{\Gamma}), \quad (2)$$

where H is applied element-wise to $\mathbf{\Gamma}$. Moreover, this will be compatible with a Markov chain formulation of the recursive marginal quantization technique presented later.

Using these vectors the gradient of the distortion function is then

$$\nabla D(\mathbf{\Gamma}) = 2\mathbf{\Gamma} \circ \mathbf{p}^\top - 2\mathbf{M},$$

where \circ indicates the element-wise Hadamard product.

The super- and sub-diagonal (or off-diagonal) entries of the Hessian matrix $\nabla^2 D(\mathbf{\Gamma})$ are given by the length- $(N-1)$ row vector

$$\mathbf{h}_{\text{off}} = -\frac{1}{2}[\mathbf{f} \circ \Delta\mathbf{\Gamma}]^\top,$$

with the main diagonal given by

$$\mathbf{h}_{\text{main}} = 2\mathbf{p} + [\mathbf{h}_{\text{off}}|0] + [0|\mathbf{h}_{\text{off}}],$$

where the copies of the \mathbf{h}_{off} vector are appended and prepended with a zero. It is now straightforward to set up the Newton-Raphson iteration in terms of the quantities.

2.2 Examples

In this section we apply the above theory to the Gaussian distribution and the noncentral chi-squared distribution with one degree of freedom. The latter is important for the higher-order recursive marginal quantization schemes that we explore later in the paper.




Figure 3: Vector Quantization of the Standard Normal Distribution

2.2.1 The Standard Normal Distribution

When X is a standard normal random variable we have

$$\begin{aligned} f_X(x) &= \phi(x) \\ F_X(x) &= \Phi(x) \\ M_X^1(x) &= -\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}} = -\phi(x), \end{aligned}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal PDF and CDF respectively. Here, a good guess for the initial quantizer $\mathbf{\Gamma}^{(0)}$ is

$$\gamma^n = \frac{5.5n}{N+1} - 2.75,$$

for $1 \leq n \leq N$. Figure 3 shows the quantizer of cardinality $N = 50$, using this initial guess, after $n_{\max} = 20$ Newton-Raphson iterations. At first glance, it is tempting to think of the quantizer (represented by bars) as one might think of a histogram and suspect that it does not adequately capture the features of the density because it has the “incorrect shape”. This is, however, a misleading analogy since a histogram represents the probability of realising a random variable in equally sized intervals on the x -axis. The quantizer, however, accumulates the probability mass over the Voronoi regions associated with each of the codewords and, as these regions are larger in the tails than in the body of the distribution, proportionally more probability mass is accumulated for codewords in the tails than for codewords in the body.

2.2.2 The Noncentral Chi-squared Distribution

While, in general, the noncentral chi-squared distribution must be specified using Bessel functions, this is not the case when the degree of freedom equals one. In particular, consider the random variable $X = (Z + \mu)^2$, where $Z \sim \mathcal{N}(0, 1)$. Then $X \sim \chi'^2(1, \lambda)$, is a noncentral chi-squared distributed with one degree of freedom and noncentrality parameter $\lambda = \mu^2$.

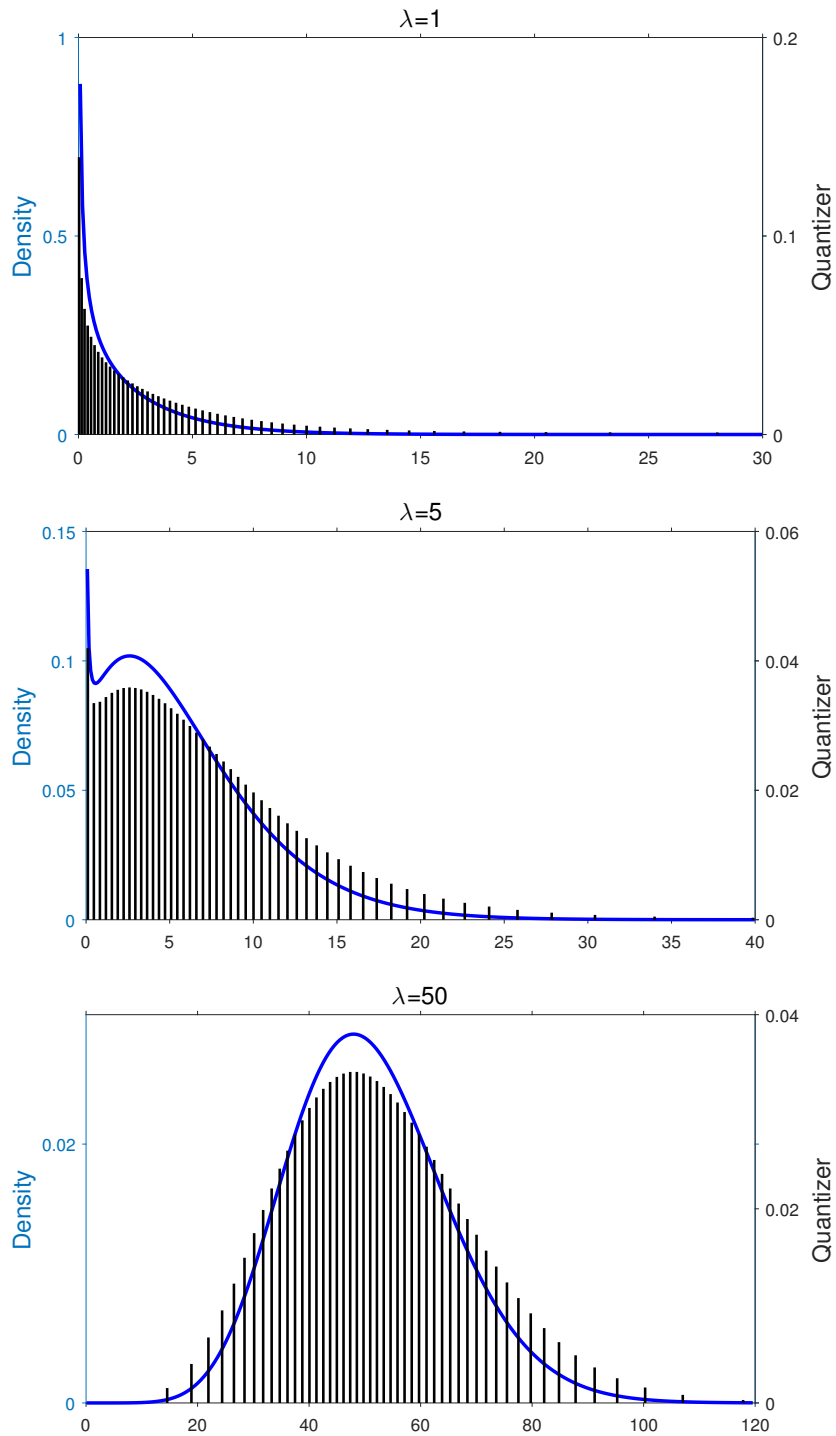


Figure 4: Three examples of the noncentral chi-squared distribution with one degree of freedom for different noncentrality parameters.

Moreover, we have

$$\begin{aligned} f_X(x) &= \frac{1}{2\sqrt{x}} (\phi(x^+) + \phi(x^-)) \\ F_X(x) &= \Phi(x^+) - \Phi(x^-) \\ M_X^1(x) &= (\lambda + 1) (\Phi(x^+) - \Phi(x^-)) + \phi(x^+)x^- - \phi(x^-)x^+, \end{aligned}$$

with

$$x^\pm = \pm\sqrt{x} - \sqrt{\lambda}.$$

This means that we may express the noncentral chi-squared distribution with one degree of freedom using the standard normal PDF and CDF, thus allowing efficient computation of a quantization scheme. This will be important for computational efficiency when we implement higher-order RMQ schemes later in the paper. A good initial guess for $\mathbf{\Gamma}^{(0)}$ is given by

$$\gamma^n = \begin{cases} \left(\frac{(3+\sqrt{\lambda})n}{N}\right)^2 & \text{for } \sqrt{\lambda} < 2.5 \\ \left(\frac{5n}{N+1} - 2.5 + \sqrt{\lambda}\right)^2 & \text{for } \sqrt{\lambda} \geq 2.5, \end{cases}$$

for $1 \leq n \leq N$.

Figure 4 shows three examples of quantizers of cardinality $N = 50$ for the noncentral chi-squared distribution with one degree of freedom for a range of noncentrality parameters. Note that for certain values of the noncentrality parameter (e.g. the central panel) the distribution is bimodal while for larger values it resembles a Gaussian distribution.

3 Recursive Marginal Quantization

Consider the continuous-time diffusion specified by the stochastic differential equation

$$dX_t = a(X_t) dt + b(X_t) dW_t, \quad X_0 = x_0 \in \mathbb{R}, \quad (3)$$

defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$ with a and b sufficiently smooth and regular functions to ensure the existence of a weak solution. The question of interest is:

How does one optimally approximate $X_{t_k} : \Omega \rightarrow \mathbb{R}$, for some time discretisation $t_k \in [0, T]$, when the distribution of each X_{t_k} is unknown?

Usually this is achieved by performing a Monte Carlo experiment using a discrete-time approximation scheme for the SDE, the simplest scheme being the Euler-Maruyama [Maruyama, 1955] update

$$\begin{aligned} \tilde{X}_{k+1} &= \tilde{X}_k + a(\tilde{X}_k)\Delta t + b(\tilde{X}_k)\sqrt{\Delta t}Z_{k+1} \\ &=: \mathcal{U}(\tilde{X}_k, Z_{k+1}), \end{aligned}$$

for $0 \leq k < K$, where $\Delta t = T/K$ and $Z_{k+1} \sim \mathcal{N}(0, 1)$, with initial value $\tilde{X}_0 = x_0$. The innovation of Pagès and Sagna [2015] was to show that a recursive procedure based on quantizing these updates is possible.

Since \tilde{X}_1 has a Gaussian distribution, it is possible to use vector quantization to obtain Γ_1 , an optimal quantization grid for the first step of the above scheme. One must, however, find a way to quantize the successive (marginal) distributions of \tilde{X}_{k+1} . Given knowledge of the distribution of \tilde{X}_k , the distortion of the quantizer Γ_{k+1} may be written as

$$\begin{aligned}\tilde{D}(\Gamma_{k+1}) &= \mathbb{E} \left[\|\tilde{X}_{k+1} - \hat{X}_{k+1}\|^2 \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\|\tilde{X}_{k+1} - \hat{X}_{k+1}\|^2 \mid \tilde{X}_k \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\|\mathcal{U}(\tilde{X}_k, Z_{k+1}) - \hat{X}_{k+1}\|^2 \mid \tilde{X}_k \right] \right] \\ &= \int_{\mathbb{R}} \mathbb{E} \left[\|\mathcal{U}(x, Z_{k+1}) - \hat{X}_{k+1}\|^2 \right] d\mathbb{P}(\tilde{X}_k \leq x).\end{aligned}\tag{4}$$

Unfortunately, we do not know the exact distribution of \tilde{X}_k for $k > 1$. The main result of [Pagès and Sagna \[2015\]](#) shows that if one uses the previously quantized distribution of \hat{X}_k , instead of the continuous distribution of \tilde{X}_k , the resultant procedure converges. Furthermore, the error associated with this procedure is bounded by a constant, which is dependant on the parameters used. As a result, the integral in (4) may be rewritten as a sum over the codewords in quantizer Γ_k and their associated probabilities.

Then an approximate value for the distortion may be computed as

$$\tilde{D}(\Gamma_{k+1}) \approx D(\Gamma_{k+1}) := \sum_{i=1}^{N_k} \mathbb{E} \left[\|\mathcal{U}(\gamma_k^i, Z_{k+1}) - \hat{X}_{k+1}\|^2 \right] \mathbb{P}(\hat{X}_k = \gamma_k^i).$$

Here, N_k is the cardinality of the quantizer Γ_k at time step k , which is allowed to vary. With this definition of $D(\Gamma_{k+1})$ we may now specify a Newton-Raphson iteration in order to compute the quantizer at time step $k + 1$, which minimizes the distortion.

Given the quantizer at time t_k , represented as a column vector $\mathbf{\Gamma}_k$, and the associated probabilities, $\mathbb{P}(\hat{X}_k = \gamma_k^i)$ for $1 \leq i \leq N_k$, the Newton-Raphson iteration for the quantizer $\mathbf{\Gamma}_{k+1}$, at time t_{k+1} , is given by

$$\mathbf{\Gamma}_{k+1}^{(n+1)} = \mathbf{\Gamma}_{k+1}^{(n)} - \left[\nabla^2 D \left(\mathbf{\Gamma}_{k+1}^{(n)} \right) \right]^{-1} \nabla D \left(\mathbf{\Gamma}_{k+1}^{(n)} \right),\tag{5}$$

for $0 \leq n < n_{\max}$.

Before developing the mathematics further, we pause to provide an intuitive explanation of how the RMQ algorithm proceeds. [Figure 5](#) is a depiction of the process that occurs. The top panel shows the quantizer at time step k . Conditional on each codeword, a Gaussian Euler update is propagated (second panel). In panel three, these updates are weighted by the probability of the associated originating codeword and summed to produce the marginal density at time step $k + 1$, as shown in the final panel. The distribution associated with this marginal density is the distribution that is quantized to produce the quantizer at time step $k + 1$. This process is repeated until the quantizer at the final time is produced.

We now proceed to derive the quantities required for the Newton Raphson iterations. To summarize notation, we write the update in affine form as

$$\mathcal{U}(\gamma_k^i, Z_{k+1}) =: U_{k+1}^i = m_k^i Z_{k+1}^i + c_k^i,\tag{6}$$

where

$$m_k^i := b(\gamma_k^i) \sqrt{\Delta t} \quad \text{and} \quad c_k^i := \gamma_k^i + a(\gamma_k^i) \Delta t,$$




Figure 5: Illustration of the RMQ algorithm.

with $Z_{k+1}^i \sim \mathcal{N}(0, 1)$ identically distributed to Z_{k+1} . Here, we have introduced a new index i for the random variable Z_{k+1}^i anticipating that it may depend on γ_k^i . This is redundant in the case of the Euler update because Z_{k+1} is a standard normal random variate irrespective of starting point. This more general notation will become necessary when we analyse more general cases (see Section 4). We denote the corresponding density and distribution functions by $f_{Z_{k+1}^i}$ and $F_{Z_{k+1}^i}$ respectively.

With this notation in place, the approximate marginal distribution of \tilde{X}_{k+1} is

$$\begin{aligned}
 F_{\tilde{X}_{k+1}}(x) &= \int_{\mathbb{R}} \mathbb{P}(\mathcal{U}(y, Z_{k+1}) \leq x) d\mathbb{P}(\tilde{X}_k \leq y) \\
 &\approx \sum_{i=1}^{N_k} \mathbb{P}(U_{k+1}^i \leq x) \mathbb{P}(\hat{X}_k = \gamma_k^i) \\
 &= \sum_{i=1}^{N_k} \left[H(-m_k^i) + \text{sgn}(m_k^i) F_{Z_{k+1}^i} \left(\frac{x - c_k^i}{m_k^i} \right) \right] \mathbb{P}(\hat{X}_k = \gamma_k^i), \tag{7}
 \end{aligned}$$

where $H(\cdot)$ is the Heaviside step function and $\text{sgn}(\cdot)$ is the signum function. The last step follows due to the fact that the left-hand probability on the penultimate line may be written as

$$\mathbb{P}(U_{k+1}^i \leq x) = \begin{cases} \mathbb{P}\left(Z_{k+1}^i \leq \frac{x - c_k^i}{m_k^i}\right) & \text{for } m_k^i \geq 0 \\ 1 - \mathbb{P}\left(Z_{k+1}^i \leq \frac{x - c_k^i}{m_k^i}\right) & \text{for } m_k^i < 0. \end{cases}$$

It should be noted that, for the Euler update, m_k^i is a proxy for the volatility of the SDE and its positivity is usually guaranteed, in which case (7) may be simplified. However, we persist with this formulation because in the general case m_k^i may not be guaranteed to be positive.

The elements of the gradient of the distortion $\nabla D(\Gamma_{k+1})$ may then be written as

$$\begin{aligned} \frac{\partial D(\Gamma_{k+1})}{\partial \gamma_{k+1}^j} &= 2 \sum_{i=1}^{N_k} \mathbb{E} \left[\mathbb{I}_{\{U_{k+1}^i \in R_j(\Gamma_{k+1})\}} \left(\gamma_{k+1}^j - U_{k+1}^i \right) \right] \mathbb{P}(\widehat{X}_k = \gamma_k^i) \\ &= 2 \sum_{i=1}^{N_k} \int_{U_{k+1}^i \in R_j(\Gamma_{k+1})} \left(\gamma_{k+1}^j - U_{k+1}^i \right) d\mathbb{P}(Z_{k+1}^i \leq x) \mathbb{P}(\widehat{X}_k = \gamma_k^i), \end{aligned} \quad (8)$$

where $1 \leq j \leq N_{k+1}$ is the index tracking the elements of the t_{k+1} quantizer, and the index associated with the t_k quantizer is i . The integration bounds in (8) must now be expressed in terms of the variable of integration.

Here, as in Section 2, $U_{k+1}^i \in R_j(\Gamma_{k+1})$ is equivalent to the inequality

$$r_{k+1}^{j-} < U_{k+1}^i \leq r_{k+1}^{j+} \quad \text{with} \quad r_{k+1}^{j\pm} = \frac{1}{2}(\gamma_{k+1}^{j\pm 1} + \gamma_{k+1}^j).$$

Defining the conditionally normalized region boundaries,

$$r_{k+1}^{i,j\pm} = \frac{r_{k+1}^{j\pm} - c_k^i}{m_k^i},$$

allows the inequality to be written in terms of the random variable Z_{k+1}^i as

$$U_{k+1}^i \in R_j(\Gamma_{k+1}) = \begin{cases} r_{k+1}^{i,j-} < Z_{k+1}^i \leq r_{k+1}^{i,j+} & \text{for } m_k^i \geq 0 \\ r_{k+1}^{i,j-} > Z_{k+1}^i \geq r_{k+1}^{i,j+} & \text{for } m_k^i < 0, \end{cases}$$

which can now be used as the range over which the integration is taken.

Directly evaluating (8), each element of the gradient of the distortion at time t_{k+1} is given by

$$\begin{aligned} \frac{\partial D(\Gamma_{k+1})}{\partial \gamma_{k+1}^j} &= 2 \sum_{i=1}^{N_k} \left[(\gamma_{k+1}^j - c_k^i) \operatorname{sgn}(m_k^i) \left(F_{Z_{k+1}^i} (r_{k+1}^{i,j+}) - F_{Z_{k+1}^i} (r_{k+1}^{i,j-}) \right) \right. \\ &\quad \left. - |m_k^i| \left(M_{Z_{k+1}^i}^1 (r_{k+1}^{i,j+}) - M_{Z_{k+1}^i}^1 (r_{k+1}^{i,j-}) \right) \right] \mathbb{P}(\widehat{X}_k = \gamma_k^i). \end{aligned}$$

Furthermore, the diagonal of the tridiagonal Hessian, $\nabla^2 D(\Gamma_{k+1})$, is given by

$$\begin{aligned} \frac{\partial^2 D(\Gamma_{k+1})}{\partial (\gamma_{k+1}^j)^2} &= \sum_{i=1}^{N_k} \left[2 \operatorname{sgn}(m_k^i) \left(F_{Z_{k+1}^i} (r_{k+1}^{i,j+}) - F_{Z_{k+1}^i} (r_{k+1}^{i,j-}) \right) \right. \\ &\quad + \frac{1}{2|m_k^i|} f_{Z_{k+1}^i} (r_{k+1}^{i,j+}) (\gamma_{k+1}^j - \gamma_{k+1}^{j+1}) \\ &\quad \left. + \frac{1}{2|m_k^i|} f_{Z_{k+1}^i} (r_{k+1}^{i,j-}) (\gamma_{k+1}^{j-1} - \gamma_{k+1}^j) \right] \mathbb{P}(\widehat{X}_k = \gamma_k^i), \end{aligned}$$

with the super-diagonal and sub-diagonal elements given by

$$\frac{\partial^2 D(\Gamma_{k+1})}{\partial \gamma_{k+1}^j \partial \gamma_{k+1}^{j+1}} = \sum_{i=1}^{N_k} \frac{1}{2|m_k^i|} f_{Z_{k+1}^i}(r_{k+1}^{i,j+}) (\gamma_{k+1}^j - \gamma_{k+1}^{j+1}) \mathbb{P}(\widehat{X}_k = \gamma_k^i)$$

and

$$\frac{\partial^2 D(\Gamma_{k+1})}{\partial \gamma_{k+1}^j \partial \gamma_{k+1}^{j-1}} = \sum_{i=1}^{N_k} \frac{1}{2|m_k^i|} f_{Z_{k+1}^i}(r_{k+1}^{i,j-}) (\gamma_{k+1}^{j-1} - \gamma_{k+1}^j) \mathbb{P}(\widehat{X}_k = \gamma_k^i).$$

Although these expressions may appear complex, they are simply summations over the density function, cumulative distribution function and first lower partial expectation of the random variable, Z_{k+1}^i , and are thus easy to compute when these functions are known.

All the detail required to implement the Newton iteration (5) has now been provided with the exception of the initial guess. In all applications considered in this paper, we have assumed that $N_k = N$ for $1 \leq k \leq K$, and used the quantizer from the previous time step as the initial guess, i.e., $\mathbf{\Gamma}_{k+1}^{(0)} = \mathbf{\Gamma}_k$.

3.1 Efficient Implementation

As in Section 2.1, where we provided an efficient matrix formulation for the Newton-Raphson iteration required for VQ, RMQ is also amenable to a matrix specification. This aids simple and computationally efficient implementation.

Aside from a guess for $\mathbf{\Gamma}_{k+1}$, we require the following time-indexed column vectors

$$[\mathbf{m}_k]_i = m_k^i, \quad [\mathbf{c}_k]_i = c_k^i, \quad 1 \leq i \leq N_k$$

and

$$[\Delta \mathbf{\Gamma}_{k+1}]_i = \gamma_{k+1}^{i+1} - \gamma_{k+1}^i, \quad 1 \leq i \leq N_{k+1} - 1.$$

The row vector of probabilities

$$\mathbf{p}_k = [\mathbb{P}(\widehat{X}_k = \gamma_k^1), \dots, \mathbb{P}(\widehat{X}_k = \gamma_k^{N_k})],$$

is retained and a row-vector of ones of length d is denoted as \mathbf{j}_d . With the exception of $\Delta \mathbf{\Gamma}_{k+1}$, which must be recomputed before each Newton iteration, the other vectors are computed once per time step.

Before each Newton iteration, three matrices must be computed in terms of the new estimate of $\mathbf{\Gamma}_{k+1}$: an $N_k \times N_{k+1}$ matrix of transition probabilities

$$\begin{aligned} [\mathbf{P}_{k+1}]_{i,j} &= \mathbb{P}(\widehat{X}_{k+1} = \gamma_{k+1}^j | \widehat{X}_k = \gamma_k^i) \\ &= \text{sgn}(m_k^i) \left[F_{Z_{k+1}^i}(r_{k+1}^{i,j+}) - F_{Z_{k+1}^i}(r_{k+1}^{i,j-}) \right], \end{aligned}$$

another matrix, of the same size, of lower partial moment values

$$[\mathbf{M}_{k+1}]_{i,j} = M_{Z_{k+1}^i}^1(r_{k+1}^{i,j+}) - M_{Z_{k+1}^i}^1(r_{k+1}^{i,j-}) \quad (9)$$

and an $N_k \times N_{k+1} - 1$ matrix of density values at the positive region boundaries

$$[\mathbf{f}_{k+1}]_{i,j} = f_{Z_{k+1}^i}(r_{k+1}^{i,j+}).$$

The gradient of the distortion function at time step $k + 1$ may now be written in terms of these vectors and matrices as

$$\nabla D(\mathbf{\Gamma}_{k+1})^\top = 2\mathbf{p}_k \left(((\mathbf{\Gamma}_{k+1}\mathbf{j}_{N_k})^\top - \mathbf{c}_k\mathbf{j}_{N_{k+1}}) \circ \mathbf{P}_{k+1} - (|\mathbf{m}_k|\mathbf{j}_{N_{k+1}}) \circ \mathbf{M}_{k+1} \right), \quad (10)$$

where \circ is the Hadamard (or element-wise) product.

The super and sub-diagonal elements of the (tridiagonal) Hessian matrix, $\nabla^2 D(\mathbf{\Gamma}_{k+1})$, are given by the vector

$$\mathbf{h}_{\text{off}} = -\frac{1}{2}\mathbf{p}_k \left((|\mathbf{m}_k|^{\circ-1}\mathbf{j}_{(N_{k+1}-1)}) \circ \mathbf{f}_{k+1} \circ (\Delta\mathbf{\Gamma}_{k+1}\mathbf{j}_{N_k})^\top \right), \quad (11)$$

while the main diagonal is given by

$$\mathbf{h}_{\text{main}} = 2\mathbf{p}_k\mathbf{P}_{k+1} + [\mathbf{h}_{\text{off}}|0] + [0|\mathbf{h}_{\text{off}}], \quad (12)$$

where $\circ - 1$ in the exponent refers to the element-wise inverse.

Equations (10), (11) and (12) provide the necessary components required for implementation of the Newton-Raphson iteration in (5). After the requisite number of iterations, the probabilities associated with the final quantizer $\mathbf{\Gamma}_{k+1}$ are computed using

$$\mathbf{p}_{k+1} = \mathbf{p}_k\mathbf{P}_{k+1},$$

where \mathbf{P}_{k+1} must be recomputed in terms of the final $\mathbf{\Gamma}_{k+1}$. Thus, the matrix formulation presented here allows RMQ to be interpreted as the propagation of an inhomogeneous Markov chain, where $\mathbf{\Gamma}_k$ represents the Markov states at time step k , the probability of being in those states is \mathbf{p}_k and the associated transition probability matrix is \mathbf{P}_{k+1} . Sometimes in the literature the transition probability matrix between time step k and $k + 1$ is represented as $\mathbf{P}_{k,k+1}$, we have chosen to omit the first index.

3.2 Example

As a first example of the RMQ algorithm, Figure 6 shows the evolution of the quantizers through time for the canonical (risk-neutral) geometric Brownian motion process

$$dS_t = rS_t dt + \sigma S_t dW_t, \quad (13)$$

using parameters $S_0 = 100$, $r = 5\%$ and $\sigma = 30\%$. The RMQ parameters used were $T = 1$, $K = 12$, $\Delta t = T/K$ and $N_k = 200$ for all k , with $n_{\text{max}} = 50$ for the VQ algorithm and $n_{\text{max}} = 5$ for the RMQ algorithm. Unless otherwise stated, these parameters are used wherever geometric Brownian motion is used for other calculations in the paper. In these plots, the colour of the quantizer indicates the associated time step, with plots in blue closer to initial time and plots in green closer to final time. This convention is kept throughout.




Figure 6: Time evolution of quantizers for the GBM process.

4 Higher-order RMQ Schemes

Given the general formulation of the RMQ algorithm in the previous section, we now explore two high-order extensions: the Milstein scheme and a simplified weak order 2.0 scheme. Any numerical scheme for an SDE that can be written in the affine form (6), may be used with the RMQ algorithm as long as the CDF, PDF and lower partial expectation of the random variable Z_{k+1}^i can be computed for all i and k .

4.1 The Milstein Scheme

The Milstein [1975] scheme for the SDE in (3) is given by

$$\tilde{X}_{k+1} = \tilde{X}_k + a(\tilde{X}_k)\Delta t + b(\tilde{X}_k)\sqrt{\Delta t}Z_{k+1} + \frac{1}{2}b(\tilde{X}_k)b'(\tilde{X}_k)\Delta t (Z_{k+1}^2 - \Delta t),$$

for $0 \leq k < K$, where $\Delta t = T/K$ and $Z_{k+1} \sim \mathcal{N}(0, 1)$ with initial value $\tilde{X}_0 = x_0$. By completion of the square, this may be written as

$$\begin{aligned} \tilde{X}_{k+1} = \tilde{X}_k + & \left(a(\tilde{X}_k) - \frac{1}{2}b(\tilde{X}_k)b'(\tilde{X}_k) \right) \Delta t - \frac{1}{2}b(\tilde{X}_k)b'(\tilde{X}_k)^{-1} \\ & + \frac{1}{2}b(\tilde{X}_k)b'(\tilde{X}_k)\Delta t \left(Z_{k+1} + \left(\sqrt{\Delta t}b'(\tilde{X}_k) \right)^{-1} \right)^2. \end{aligned}$$

Thus, the update (6) may be written in the affine form required as

$$U_{k+1}^i = m_k^i Z_{k+1}^i + c_k^i,$$

where

$$m_k^i = \frac{1}{2}b(\gamma_k^i)b'(\gamma_k^i)\Delta t$$

and

$$c_k^i = \gamma_k^i + \left(a(\gamma_k^i) - \frac{1}{2}b(\gamma_k^i)b'(\gamma_k^i) \right) \Delta t - \frac{1}{2}b(\gamma_k^i)b'(\gamma_k^i)^{-1}.$$

The random variable Z_{k+1}^i is now noncentral chi-squared distributed with one degree of freedom and noncentrality parameter

$$\lambda_{k+1}^i = \left(\sqrt{\Delta t} b'(\gamma_k^i) \right)^{-2}.$$

It is important to note that, unlike the Euler-Maruyama case, the distribution of the random variable $Z_{k+1}^i \sim \chi'^2(1, \lambda_{k+1}^i)$ now depends on the codeword γ_k^i .

Although the Milstein scheme possesses a strong order of convergence of 1, compared to the Euler scheme, which only has strong order of convergence of $\frac{1}{2}$, both schemes have a weak order of convergence of 1. Thus, while the Milstein scheme may be more accurate than the Euler scheme, especially for path dependent claims, we require a different update for higher weak order convergence. We now explore such a scheme.

4.2 A Weak Order 2.0 Taylor Scheme

While it is not possible to write an exact weak order 2.0 scheme in the affine form required, the simplified weak order 2.0 scheme of Kloeden and Platen [1999] is amenable. This scheme is given by

$$\begin{aligned} \tilde{X}_{k+1} &= \tilde{X}_k + a(\tilde{X}_k)\Delta t + b(\tilde{X}_k)\sqrt{\Delta t}Z_{k+1} + \frac{1}{2}b(\tilde{X}_k)b'(\tilde{X}_k)\Delta t(Z_{k+1}^2 - 1) \\ &\quad + \frac{1}{2} \left(a'(\tilde{X}_k)b(\tilde{X}_k) + a(\tilde{X}_k)b'(\tilde{X}_k) + \frac{1}{2}b''(\tilde{X}_k)b^2(\tilde{X}_k) \right) (\Delta t)^{\frac{3}{2}}Z_{k+1} \\ &\quad + \frac{1}{2} \left(a(\tilde{X}_k)a'(\tilde{X}_k) + \frac{1}{2}a''(\tilde{X}_k)b^2(\tilde{X}_k) \right) (\Delta t)^2, \end{aligned}$$

for $0 \leq k < K$, where $\Delta t = T/K$ and $Z_{k+1} \sim \mathcal{N}(0, 1)$ with initial value $\tilde{X}_0 = x_0$. Again, completion of squares is used to write this update in the required affine form,

$$U_{k+1}^i = m_k^i Z_{k+1}^i + c_k^i,$$

where

$$m_k^i = \frac{1}{2}b(\gamma_k^i)b'(\gamma_k^i)\Delta t$$

and

$$\begin{aligned} c_k^i &= \gamma_k^i + \left(a(\gamma_k^i) - \frac{1}{2}b(\gamma_k^i)b'(\gamma_k^i) \right) \Delta t + \frac{1}{2} \left(a(\gamma_k^i)a'(\gamma_k^i) + \frac{1}{2}a''(\gamma_k^i)b^2(\gamma_k^i) \right) (\Delta t)^2 \\ &\quad - \frac{\left(b(\gamma_k^i) + \frac{1}{2} \left(a'(\gamma_k^i)b(\gamma_k^i) + a(\gamma_k^i)b'(\gamma_k^i) + \frac{1}{2}b''(\gamma_k^i)b^2(\gamma_k^i) \right) \Delta t \right)^2}{2b(\gamma_k^i)b'(\gamma_k^i)}. \end{aligned}$$

Here, Z_{k+1}^i is again noncentral chi-squared distributed with one degree of freedom, with noncentrality parameter given by

$$\lambda_{k+1}^i = \left(\frac{b(\gamma_k^i) + \frac{1}{2} \left(a'(\gamma_k^i)b(\gamma_k^i) + a(\gamma_k^i)b'(\gamma_k^i) + \frac{1}{2}b''(\gamma_k^i)b^2(\gamma_k^i) \right) \Delta t}{b(\gamma_k^i)b'(\gamma_k^i)\sqrt{\Delta t}} \right)^2,$$

or, more succinctly, $Z_{k+1}^i \sim \chi'^2(1, \lambda_{k+1}^i)$.




Figure 7: Numerical illustration of the convergence of the first moment for the GBM and CEV processes.

4.3 Examples

To illustrate the increased accuracy of these high-order schemes, the RMQ algorithm is applied to geometric Brownian motion, as previously described by (13), and the constant elasticity of variance (CEV) process. The SDE for the CEV process is

$$dS_t = rS_t dt + \sigma S_t^\alpha dW_t,$$

and, in the examples that follow, the process-specific parameters chosen were $S_0 = 100$, $r = 5\%$ and $\alpha = 0.7$ and $\sigma = \sigma_{\text{LN}} S_0^{1-\alpha}$, with σ given in terms of the instantaneous log-normal volatility $\sigma_{\text{LN}} = 30\%$.

In Figure 7 we show numerical evidence for weak order convergence. Using both GBM and CEV, the first moment of the resulting terminal quantizer is compared against the true moment for a range of time step sizes. In the case of GBM the parameters in Section 3.2 were used. For both GBM and CEV, the RMQ specific parameters mentioned in that section were also used with the exception of the cardinality, which was set at $N_k = 1000$ for increased accuracy. The approximate gradient of the results, β , which is an indicator of the weak order convergence, is computed and displayed. As expected, the Euler and Milstein scheme have approximately the same weak order of convergence of close to one, whereas the simplified weak order 2.0 scheme is close to two. The latter scheme also produces results with substantially lower absolute error.

Since the true conditional distributions for both the GBM and CEV processes are known in closed form, we can compare the approximate marginal distributions at time step $k+1$, as implied by the quantizer at time step k and computed using (7), with the exact distributions at time step $k+1$. In the case of the CEV process we used the analytical expression for the distribution given by Lindsay and Brecher [2012] adjusted for the drift component in the SDE.

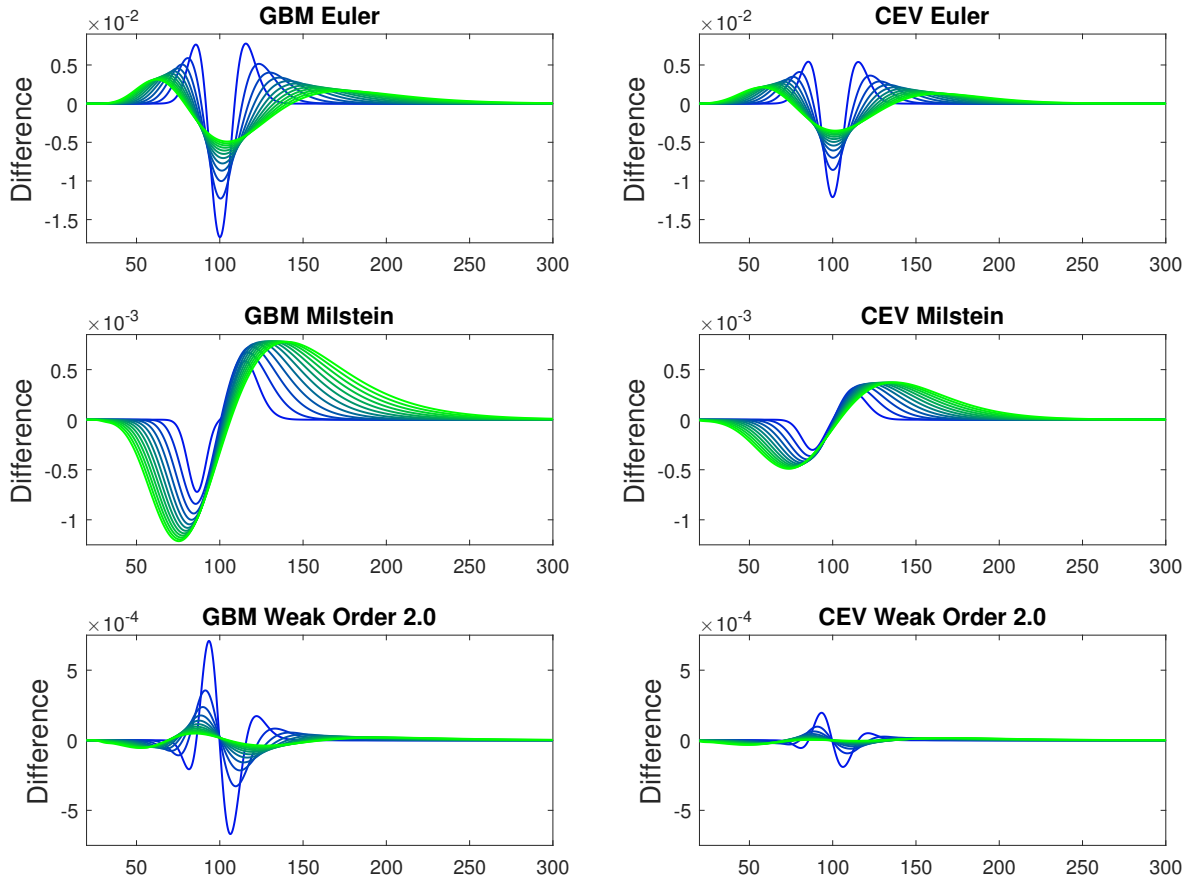


Figure 8: The true marginal distributions of the GBM and CEV processes compared to the marginal distribution implied by quantization.

The difference between the exact marginal distribution and the implied marginal distributions for each of the three schemes is displayed in Figure 8. Here we reverted to using a cardinality of $N_k = 200$. Note the scale of the y -axes of the graphs in the figure — from top to bottom, the magnitude of the error decreases by an order of magnitude in each successive row. This gives an indication of the improvement that can be expected when these higher order schemes are used to price contingent claims.

5 The Zero Boundary

Sometimes discrete-time approximations of an SDE may exhibit behaviour that is inconsistent with the true solution. For example, an Euler-Maruyama approximation of geometric Brownian motion, the square-root process or the CEV process can, under certain circumstances, generate negative values, even though the SDE specification guarantees non-negativity in each case. As a result, discrete-time Monte Carlo simulations are modified to allow for reflecting or absorbing behaviour at zero, see for example Lord et al. [2010]. In this section, we describe how the RMQ algorithm may be modified in a similar manner.

5.1 Absorbing Boundary

To model an absorbing boundary, the domain of the approximate marginal distribution of \tilde{X} , see (7), must be left-truncated at zero. This means that the probabilities associated with the quantizer at each step will not sum to unity. The probability that is not accounted for as a result of the domain truncation is the mass accumulated at the absorbing zero boundary. To compensate for this, the quantizer at each time step is augmented with an extra codeword, which has a value of zero and a probability equal to one minus the sum of the probabilities associated with all the other codewords at that time step. The transition probability matrix may also be augmented in a consistent manner by realising that once the process attains the zero state it must remain in that state indefinitely, i.e., the conditional probability of moving from the absorbing state to any other state is zero, and, correspondingly, the conditional probability of remaining in the absorbing state is one.

Modifying the algorithm is simple and incurs no additional computational burden. Given that the elements of the previous quantizer $\mathbf{\Gamma}_k$ are all positive, the affine form of the update (6), will be negative when

$$Z_{k+1}^i < -\frac{c_k^i}{m_k^i}.$$

This implies that the domain of each Z_{k+1}^i must be left-truncated at $-\frac{c_k^i}{m_k^i}$ to ensure only positive codewords at time-step $k+1$. This is achieved by setting

$$r_{k+1}^{i,1-} = -\frac{c_k^i}{m_k^i}, \quad (14)$$

for $1 \leq i \leq N_k$, in the implementation described in Section 3.1. The rest of the algorithm proceeds without modification.

Of course, this all depends on the fact that the quantizer at the first time step, $\mathbf{\Gamma}_1$, also has positive elements. This is achieved using the analogous truncation in the vector quantization algorithm. The initial guess for the Newton iteration must also be truncated to ensure positivity.

5.2 Reflecting Boundary

Consider Figure 9, which shows $f(x)$, a density function — in this case a standard Gaussian density. The red line represents a reflecting boundary at $\bar{x} = -1.5$. The f values to the left of the boundary are reflected and depicted by the dashed yellow line in the figure. These values are given by $f(2\bar{x} - x)$ for $x > \bar{x}$.

Thus, restricting the domain to $[\bar{x}, \infty)$ the reflected density, denoted $\bar{f}(x)$, is given as the sum

$$\bar{f}(x) = f(x) + f(2\bar{x} - x).$$

Direct integration of this expression over the integration limits from \bar{x} to $x \in [\bar{x}, \infty)$ gives the reflected distribution function

$$\bar{F}(x) = F(x) - F(2\bar{x} - x)$$

and the first lower partial expectation function

$$\bar{M}^1(x) = M^1(x) + M^1(2\bar{x} - x) - 2\bar{x}F(2\bar{x} - x) - 2M^1(\bar{x}) + 2\bar{x}F(\bar{x}), \quad (15)$$




Figure 9: Illustration of the standard Gaussian density reflected around -1.5 .

where $F(x)$ and $M^1(x)$ are the un-reflected distribution and first lower expectation functions associated with f . This may be applied, not only to Gaussian case, but also to the noncentral chi-squared cases required for the higher order updates.

Modifying the RMQ algorithm to allow for a reflecting boundary at zero requires two changes to the implementation described in Section 3.1. Firstly, the lower bound for the integration, i.e., the domain of the Z_{k+1}^i random variable in each affine update, must be left-truncated by replacing the furthest left region boundary as in (14) above. Secondly, the density, distribution and first lower partial expectation functions associated with each random variable, must be replaced by their reflected counterparts

$$\begin{aligned}\bar{f}_{Z_{k+1}^i}(x) &= f_{Z_{k+1}^i}(x) + f_{Z_{k+1}^i}(2\bar{x}_k^i - x), \\ \bar{F}_{Z_{k+1}^i}(x) &= F_{Z_{k+1}^i}(x) - F_{Z_{k+1}^i}(2\bar{x}_k^i - x),\end{aligned}$$

and

$$\bar{M}_{Z_{k+1}^i}^1(x) = M_{Z_{k+1}^i}^1(x) + M_{Z_{k+1}^i}^1(2\bar{x}_k^i - x) - 2\bar{x}_k^i F_{Z_{k+1}^i}(2\bar{x}_k^i - x), \quad (16)$$

for $x \in [\bar{x}_k^i, \infty)$, where $\bar{x}_k^i = -\frac{c_k^i}{m_k^i}$. The rest of the algorithm proceeds as normal. The astute reader will have noticed that there are two terms missing in (16) when compared with (15). The reason for this omission is that these terms are constants for each i and that the RMQ algorithm always only requires differences of partial moment terms, as seen in (9). There is, therefore, a cancelation of the constant terms when this difference is taken and thus we may use a definition that excludes them.

As in the case of the absorbing boundary, the analogous reflection must be applied in the vector quantization algorithm to ensure that $\mathbf{\Gamma}_1$ is consistent.

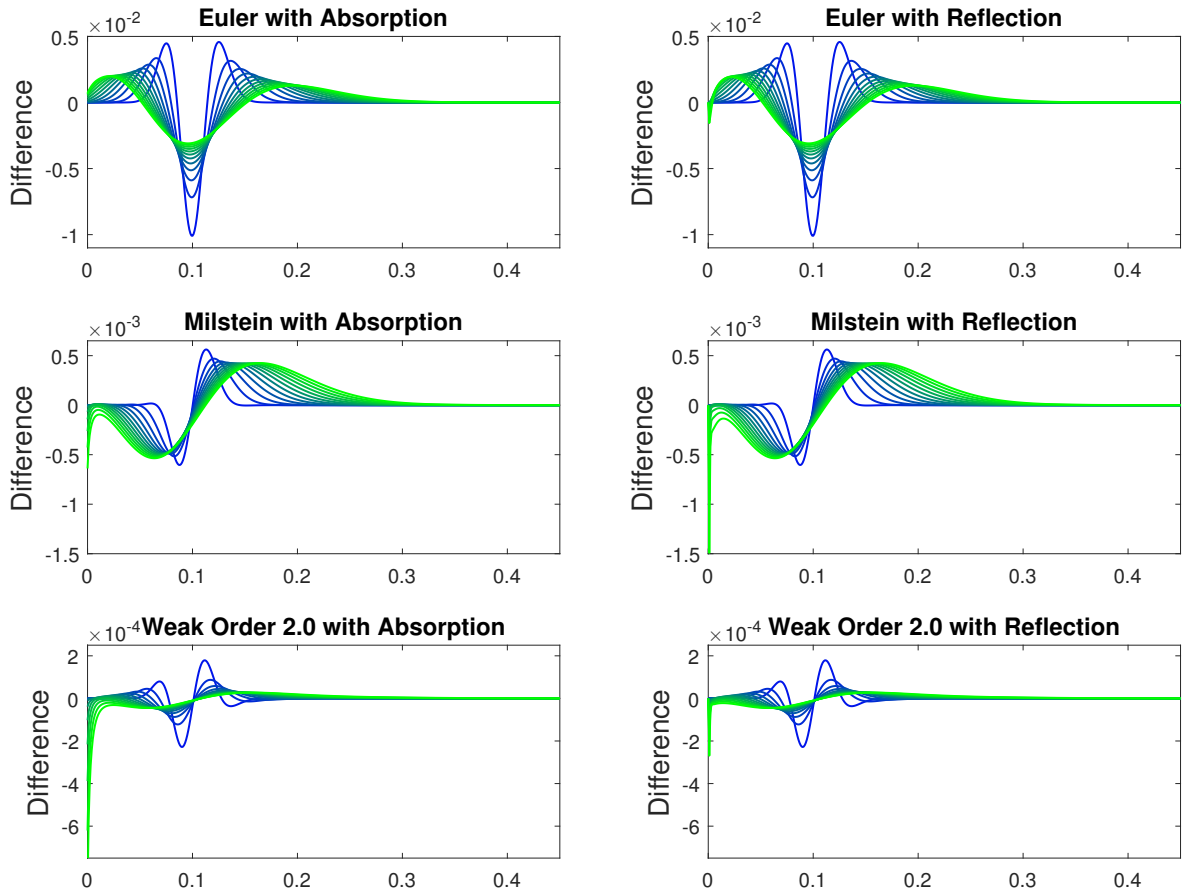


Figure 10: The difference between the true and approximate marginal distributions for the CEV process. The left column shows the case of absorption while the right shows reflection.

5.3 Examples

It is well known that when $0 < \alpha < 0.5$, the CEV process may reach zero and that this state may be either absorbing or reflecting. [Lindsay and Brecher \[2012\]](#) give the corresponding marginal distributions for both these cases (it is easy to adjust their formulations to account for the drift term in the SDE for CEV). In [Section 4.3](#) we considered the CEV process with $\alpha > 0.5$, which only allows absorption at zero. Now consider the case where $S_0 = 0.5$, $\alpha = 0.35$ and $\sigma_{LN} = 50\%$, with the rest of the parameters as before. [Figure 10](#) shows the difference between the exact marginal distribution and the marginal distribution implied by RMQ for the three schemes as modified to account for an absorbing boundary (left) and a reflecting boundary (right). As before, the scale of the graphs changes from top to bottom, indicating the improvement as a result of the higher order schemes.

Note that, under this choice of parameters for the CEV model, the standard RMQ formulation of [Section 3](#) fails. Without implementing the modifications for either absorption or reflection proposed in this section, some codewords become negative at a certain point in the execution of the RMQ algorithm, leading to discrete-time updates with imaginary values.




Figure 11: Accuracy of GBM and CEV European put prices computed using RMQ, as compared to analytical solutions.

6 Pricing

In this section, contingent claims are priced using the RMQ algorithm and the three update schemes are compared for accuracy. The claims priced include European, Bermudan and discretely-monitored barrier options under the dynamics of both geometric Brownian motion (GBM) and its generalization, the constant elasticity of variance (CEV) model.

The GBM model and its parameters are described in Section 3.2, whereas the specification for the CEV model can be found in Section 4.3. These parameters are used for pricing throughout this section. As is implied by these specifications, the continuously compounded interest rate is assumed constant with a value $r = 5\%$. All option maturities are one year and the RMQ algorithm is executed using $K = 12$, i.e., using monthly steps, with constant cardinality of $N_k = 200$ for all k .

6.1 European Option Pricing

Once a terminal quantizer has been obtained using the RMQ algorithm, a European option with payoff function $H(S, X)$ at maturity $T = t_K$, where S represents the asset process and X the strike, may be priced directly by using the expectation defined in (2). The price is given by

$$H_0 = e^{-rT} \mathbb{E} [H(S_T, X)] \approx e^{-rT} \mathbf{p}_K H(\mathbf{\Gamma}_K, X), \quad (17)$$

where H_0 is the value of the claim at initial time $t_0 = 0$ and $H(\mathbf{\Gamma}_K, X)$ is the function H applied element-wise to $\mathbf{\Gamma}_K$, which, as specified previously, is a column vector of length N_K .

Figure 11 shows the accuracy of put option prices for the GBM and CEV models for a wide range of strikes. The GBM option prices are compared against the Black-Scholes option pricing formula, whereas the CEV prices are compared against the analytical solution originally due to Schroder [1989] and reformulated in terms of the noncentral chi-squared




Figure 12: Accuracy of GBM and CEV Bermudan put option prices, as compared to a high resolution Crank-Nicholson finite difference scheme.

distribution by [Hsu et al. \[2008\]](#). In the graphs, the x -axis represents fixed-spot inverse moneyness, which is determined as the variable strike value over the initial asset price, S_0 .

Even though the Euler scheme is accurate to start with, the increased accuracy of the Milstein and the simplified weak order 2.0 schemes is evident. For certain strikes the error is reduced by an order of magnitude.

6.2 Bermudan Option Pricing

Bermudan option prices are computed using the standard Backward Dynamic Programming Principle (BDPP), an important result from discrete-time optimal stopping theory. [Pagès \[2014\]](#) reviews the use of the BDPP as applied to grids that result from a quantization.

Once quantization grids and corresponding transition probability matrices have been computed using the RMQ algorithm, the high-level algorithm for Bermudan option pricing may be specified as follows:

1. Initialize $\mathbf{h}_K = H(\mathbf{\Gamma}_K, X)$
2. For $k = K - 1, \dots, 1$
Set $\mathbf{h}_k = \max(H(\mathbf{\Gamma}_k, X), e^{-r\Delta t} \mathbf{P}_{k+1} \mathbf{h}_{k+1})$
3. Set $H_0 = e^{-r\Delta t} \mathbf{p}_1 \mathbf{h}_1$

Here the max function is applied element-wise with its second argument being the continuation value, which is easily computed as a conditional expectation due to availability of the transition probability matrix at each time step. The initial value of the Bermudan claim is given by H_0 .

In [Figure 12](#) the accuracy of a Bermudan put option with monthly exercise opportunities is shown for the GBM and CEV models. The reference price is computed using a high resolution




Figure 13: Accuracy of GBM and CEV discretely monitored up-and-out put option prices, as compared to a Monte Carlo simulation.

Crank-Nicholson finite difference scheme using 600 time steps and 800 stock increments, equally spaced between zero and $4 \times S_0$.

All three RMQ algorithms result in low absolute errors, with the simplified weak order 2.0 scheme again producing errors that are an order of magnitude smaller.

6.3 Barrier Option Pricing

The pricing of barrier options has previously been explored in the context of quantization by Sagna [2011]. In that work he showed that the barrier-crossing approach described in Section 6.4 of Glasserman [2003] may be applied to marginal quantization using a so-called transition kernel formulation. Using our notation, we now present this approach and apply it to discretely monitored barrier options.

Consider expression (17) for pricing European options, which may be re-written as

$$H_0 \approx e^{-rT} \left(\mathbf{P}_1 \prod_{k=1}^{K-1} \mathbf{P}_{k+1} \right) H(\mathbf{\Gamma}_K, X).$$

To price a knock-out barrier option the transition probability matrix at each time step in this expression must to be modified to take into account the possibility that the underlying process breaches the barrier. Thus, we rescale the transition probabilities by multiplying them by the probability of not having crossed the barrier.

Let $g(x, y)$ be the probability of transitioning between states x and y without crossing the barrier. If we form an N_k by N_{k+1} matrix of values

$$[\mathbf{G}_{k+1}]_{i,j} = g(\gamma_k^i, \gamma_{k+1}^j),$$

then $\mathbf{P}_{k+1} \circ \mathbf{G}_{k+1}$ defines the transition kernel. The barrier option may then be priced using

$$H_0 \approx e^{-rT} \left((\mathbf{p}_1 \circ \mathbf{g}_1) \prod_{k=1}^{K-1} (\mathbf{P}_{k+1} \circ \mathbf{G}_{k+1}) \right) H(\mathbf{\Gamma}_K, X),$$

where $\mathbf{g}_1 = [g(S_0, \gamma_1^1), \dots, g(S_0, \gamma_1^{N_1})]$ is a row vector.

In the case of discretely-monitored up-and-out barrier options with barrier level L , the function g is given simply as the indicator function

$$g(x, y) = \mathbb{I}_{\{\max(x, y) < L\}}.$$

See [Glasserman \[2003\]](#) and [Sagna \[2011\]](#) for the continuous monitoring case using the Rayleigh distribution.

In [Figure 13](#) the accuracy of discretely-monitored up-and-out put option prices generated using RMQ is compared to a Monte Carlo implementation under the GBM and CEV models. The barrier levels (x -axis) are expressed as multiples of the at-the-money strike. Since we have chosen $K = 12$ the barrier is monitored monthly.

The reference prices are provided by a million-path Monte Carlo experiment. The Monte Carlo paths are generated using Euler-Maruyama updates with 1 200 time steps, while ensuring that the barriers are only monitored at the monthly intervals. The reason that we used Euler updates is that it proved inefficient to compute Monte Carlo samples from the exact distribution for the CEV model. We did, however, use the exact transition density to generate Monte Carlo samples for GBM to confirm that results were consistent and generating the correct standard deviations.

The results show a similar pattern to the previous sections, with an important caveat: the simplified weak order 2.0 scheme produces prices that, for the majority of the barrier values considered, lie within the three standard deviation bound of the million-path Monte Carlo experiment. The other two RMQ schemes are producing results that are statistically significantly incorrect when compared to the Monte Carlo simulation.

7 Conclusion

In this work, the recursive marginal quantization algorithm of [Pagès and Sagna \[2015\]](#) has been extended from the standard Euler-Maruyama scheme to higher-order numerical schemes, specifically the Milstein scheme and a simplified weak order 2.0 scheme of [Kloeden and Platen \[1999\]](#). This entailed replacing the Gaussian updates with noncentral chi-squared updates and generalising the formulation of RMQ to account for these different distributions.

We have also shown how to augment the RMQ algorithm in order to implement absorption or reflection at the zero boundary, thus ensuring non-negativity of solutions. This allows RMQ to be applied in cases where the algorithm may previously have failed.

Improved approximation of the marginal distributions by the higher order updates has been demonstrated, using GBM and CEV dynamics. All the schemes were used to price European, Bermudan and discrete barrier options. The pricing results are encouraging and serve as justification for use of higher order schemes. Although only numerical evidence for the convergence of the new methods has been provided, further work to provide theoretical justification is ongoing.

References

- G. Callegaro, L. Fiorin, and M. Grasselli. Pricing and calibration in local volatility models via fast quantization. Available at SSRN 2495829, 2014.
- G. Callegaro, L. Fiorin, and M. Grasselli. Quantized calibration in local volatility. *Risk Magazine*, 28:62–67, 2015a.
- G. Callegaro, L. Fiorin, and M. Grasselli. Pricing via quantization in stochastic volatility models. Available at SSRN 2669734, 2015b.
- Q. Du, V. Faber, and M. Gunzburger. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
- P. Glasserman. *Monte Carlo methods in Financial Engineering*. Springer, 2003.
- S. Graf and H. Luschgy. *Foundations of Quantization for Probability Distributions*. Springer, 2000.
- Y.-L. Hsu, T. Lin, and C. Lee. Constant elasticity of variance (CEV) option pricing model: Integration and detailed derivation. *Mathematics and Computers in Simulation*, 79(1): 60–71, 2008.
- P. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, 1999.
- A. Lindsay and D. Brecher. Simulation of the CEV process and the local martingale property. *Mathematics and Computers in Simulation*, 82(5):868–878, 2012.
- S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- R. Lord, R. Koekkoek, and D. V. Dijk. A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2):177–194, 2010.
- G. Maruyama. Continuous markov processes and stochastic equations. *Rendiconti del Circolo Matematico di Palermo*, 4(1):48–90, 1955.
- G. Milstein. Approximate integration of stochastic differential equations. *Theory of Probability and Its Applications*, 19(3):557–562, 1975.
- G. Pagès. Introduction to optimal vector quantization and its applications for numerics. Technical report, July 2014. URL <https://hal.archives-ouvertes.fr/hal-01034196>. 54 pages.
- G. Pagès and A. Sagna. Recursive marginal quantization of the Euler scheme of a diffusion process. *Applied Mathematical Finance*, 22(5):463–498, 2015.
- G. Pagès, H. Pham, and J. Printems. Optimal quantization methods and applications to numerical problems in finance. In *Handbook of Computational and Numerical Methods in Finance*, pages 253–297. Springer, 2004.

- A. Sagna. Pricing of barrier options by marginal functional quantization. *Monte Carlo Methods and Applications*, 2011.
- M. Schroder. Computing the constant elasticity of variance option pricing formula. *Journal of Finance*, 44(1):211–219, 1989.