Quantized calibration in local volatility

Pricing of a derivative should be fast and accurate, otherwise it cannot be calibrated efficiently. Here, Giorgia Callegaro, Lucio Fiorin and Martino Grasselli apply a fast quantization methodology, in a local volatility context, to the pricing of vanilla and barrier options that overcomes the numerical problems in existing methods

Quantization is a widely used tool in information theory, cluster analysis, pattern and speech recognition, numerical integration, data mining and, as in our case, numerical probability. The birth of optimal quantization dates back to the 1950s, when the necessity to optimise signal transmission, by appropriate discretisation procedures, arose.

Quantization consists of approximating a signal that admits a continuum of possible values by a signal that takes values in a discrete set. Vector quantization deals with signals that are finite dimensional, such as random variables, while functional quantization extends the concepts to the infinite-dimensional setting, as it is in the case of stochastic processes. Quantization of random vectors can be considered a discretisation of the probability space, providing in some sense the best approximation to the original distribution. It is therefore crucial for a given distribution to optimise the geometric location of these points and to evaluate the resulting error. Some numerical procedures have been developed to get optimal quadratic quantization of the Gaussian (and even non-Gaussian) distribution in high dimension, mostly based on stochastic optimisation algorithms. Over the years, many other application fields have been discovered, such as, in the 1990s, numerical integration. This opened the door, especially in France and Germany, to new research perspectives in numerical probability and applications to mathematical finance.

For a comprehensive introduction to optimal vector quantization and its applications, we refer the reader to the recent paper of Pagès (2014) and the references therein.

While theoretically sound and deeply investigated, optimal quantization typically suffers from the numerical burden that the algorithms involve (see, for example, the numerical results in Pagès & Printems (2005)). The main reason is related to the highly time-consuming procedure required by the determination of the optimal grid, especially in the multi-dimensional case, where stochastic algorithms are necessary. Recently, a very promising type of quantization, called recursive marginal quantization, has been introduced by Pagès & Sagna (2014) and applied to the Euler scheme of a pseudo-CEV local volatility model in a pricing context. This new approach provides sub-optimal quantization grids in a very precise and fast way.

Following the lines of Pagès & Sagna (2014), in our paper we apply recursive marginal quantization to a special local volatility model, namely the quadratic normal volatility (QNV) model, that has been investigated by Blacher (2001), Ingersoll (1997), Lipton (2002), Zühlsdorff (2002) and lately revisited by Andersen (2011) and Carr, Fischer & Ruf (2013). We find stationary quantizers via a Newton-Raphson method, in order to efficiently price vanilla and exotic derivatives. Indeed, the Newton-Raphson procedure, being deterministic, is very fast and it allows us to provide the first example of calibration based on quantization. The recursive marginal quantization is competitive even when closed-form formulas for vanillas are available (as in the case of call and put prices for the QNV model). Finally, we show the flexibility and the efficiency of the recursive marginal quantization in the pricing of non-vanilla contracts, when compared with the classic Monte Carlo simulation. Our numerical algorithms have performed quite well (with regard to Monte Carlo), so that in this paper no speed-up procedure has been tested. As a consequence, this paper does not provide the fastest possible numerical method, but a procedure that is competitive enough if compared with Monte Carlo.

The paper is organised as follows. In the next section we give a quick application-oriented overview of the vector quantization methodology. We then extend the vector quantization method to the class of Markov diffusion processes, leading to the recursive marginal quantization. Then we introduce the QNV model together with the well-known results on closed-form formulas for vanilla option prices. Moreover, we apply the recursive marginal quantization approach to the pricing of barrier options. The next section illustrates our numerical results, with particular emphasis on the calibration exercise on real data, and the final section concludes. Some technical details are given in the appendix of the extended version, available at http://ssrn.com/abstract=2495829.

Brief overview on vector quantization

We first provide some more technical details on vector quantization of a random variable (see, for example, Graf & Luschgy 2000; Pagès & Printems 2005; Pagès, Pham & Printems 2003; Pagès 2014). Consider an $\mathbb{R}^d$-valued random variable X defined on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with finite rth moment and probability distribution $\mathbb{P}_X$. As already mentioned, quantization can be considered as a discretisation of the probability space by at most N values, providing in some sense the best approximation to the original distribution. In other words, N-quantizing the random variable X, taking infinitely many values, boils down to approximating it by a discrete random variable $\hat{X}$ valued in a set of cardinality N, $\Gamma = \{x_1, \dots, x_N\}$. As a consequence, in view of our application to quantitative finance, integrals of the form $\mathbb{E}[h(X)]$ (for a given Borel function $h\colon \mathbb{R}^d \to \mathbb{R}$) can be approximated by the finite sum below:

$$\mathbb{E}[h(X)] \cong \mathbb{E}[h(\hat{X})] = \sum_{i=1}^{N} h(x_i)\mathbb{P}(\hat{X} = x_i) \tag{1}$$

Clearly it still remains to clarify how to get the optimal or at least a ‘good’ grid Γ and the associated weights $\mathbb{P}(\hat{X} = x_i)$, $i = 1, \dots, N$, and to estimate the error. An example of optimal 50-dimensional quantization grid for the bivariate Gaussian distribution is given in figure 1.

More rigorously, quantizing X on a given grid $\Gamma = \{x_1, \dots, x_N\}$ consists of projecting X on the grid Γ following the closest neighbour rule. An N-quantizer is a Borel function $f_N : \mathbb{R}^d \to \Gamma \subset \mathbb{R}^d$ projecting X on Γ. The induced mean L^r-error (for $r > 0$) is called

L^r-mean quantization error and is given by:

$$\|X - f_N(X)\|_r = \Big\| \min_{1 \leqslant i \leqslant N} |X - x_i| \Big\|_r$$

where $\|X\|_r := [\mathbb{E}(|X|^r)]^{1/r}$ is the usual norm in L^r. The projection of X on Γ, $f_N(X)$, is called the quantization of X (in the sequel, we will alternatively use $f_N(X)$ or $\mathrm{Proj}_\Gamma(X)$ to indicate the quantization of X). As a function of the grid Γ, the L^r-mean quantization error is continuous and reaches a minimum over all the grids with size at most N. A grid $\Gamma^\star$ minimising the L^r-mean quantization error over all the grids with size at most N is called an L^r-optimal quantizer.

An optimal quantizer is then associated to an optimal grid of points $\Gamma^\star$ and to an optimal Borel partition of the space $\mathbb{R}^d$, $(C_i(\Gamma^\star))_{1 \leqslant i \leqslant N}$, and vice versa, so that the quantizer is defined as follows:

$$f_N(X) = \sum_{i=1}^{N} x_i \mathbb{1}_{C_i(\Gamma^\star)}(X)$$

where the above partition $\{C_i(\Gamma^\star)\}_{i=1,\dots,N}$, with $C_i(\Gamma^\star) \subset \{\xi \in \mathbb{R}^d : \|\xi - x_i\| = \min_{1 \leqslant j \leqslant N} \|\xi - x_j\|\}$, is called the Voronoi partition, or tessellation induced by $\Gamma^\star$. Moreover, the L^r-mean quantization error vanishes as the grid size $N \to +\infty$ and the convergence rate has been computed in the celebrated Zador theorem (see Graf & Luschgy 2000):

$$\min_{\Gamma,\, |\Gamma| = N} \|X - \mathrm{Proj}_\Gamma(X)\|_r = Q_r(\mathbb{P}_X) N^{-1/d} + o(N^{-1/d})$$

where $Q_r(\mathbb{P}_X)$ is a non-negative constant ($r = 2$ of course will be of particular interest, with the corresponding quadratic optimal quantizer). From a numerical point of view, finding an optimal quantizer may be a very challenging task. This motivates the introduction of sub-optimal criteria, mostly because one is typically interested in quantizations that are close to X in distribution. We then introduce the notion of stationary quantizers.

Definition 1   An N-quantizer $\Gamma^N = \{x_1, \dots, x_N\}$ inducing the quantization f_N of X is said to be stationary if:

$$\mathbb{E}[X \mid f_N(X)] = f_N(X)$$

In particular, if we introduce the distortion function associated with Γ^N:

$$D(\Gamma^N) := \sum_{i=1}^{N} \int_{C_i(\Gamma^N)} |z - x_i|^2 \,\mathrm{d}\mathbb{P}_X(z) \tag{2}$$

then it turns out that stationary quantizers are critical points of the distortion function (that is, a stationary quantizer Γ^N satisfies $\nabla D(\Gamma^N) = 0$). Computing the quadratic optimal quantizers, or L^r-optimal (or stationary) quantizers in general, together with finding the associated weights and L^r-mean quantization errors, are important issues. Several algorithms are used in practice. In the one-dimensional framework, the L^r-optimal quantizers are unique up to the grid size as soon as the density of X is strictly log-concave. In this case, the Newton algorithm is commonly used to carry out the L^r-optimal quantizers when closed or semi-closed formulas are available for the gradient (and the Hessian matrix). From a numerical point of view, stationary quantizers are interesting insofar as they can be found through zero search recursive procedures such as Newton's algorithm, which can be efficiently performed.

1 Example of optimal grid for a bivariate (standard) Gaussian distribution with $N = 50$

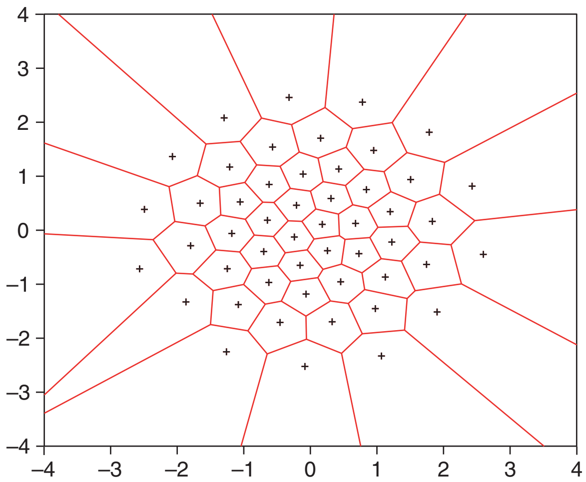

This is essential for vector quantization and we need this in order to proceed. For a thorough treatment of this topic, we refer to Graf & Luschgy (2000).[1]

Recursive marginal quantization

In this section, we consider the quantization of a continuous-time diffusive Markov process Y, whose evolution is specified by the following SDE:

$$\mathrm{d}Y_t = b(t, Y_t)\,\mathrm{d}t + a(t, Y_t)\,\mathrm{d}W_t, \quad Y_0 = y_0 > 0 \tag{3}$$

where W is a standard Brownian motion and the functions a and b satisfy the usual conditions ensuring the existence of a strong solution to the SDE. Following the approach presented in Pagès & Sagna (2014), we work on the Euler scheme of Y and we discretise the process by exploiting its Markov property via vector quantization.

Having fixed a time horizon $T > 0$ and a time discretisation grid $\{0 = t_0, t_1, \dots, t_M = T\}$, with constant step size $\Delta_k = t_k - t_{k-1}$, $k \geqslant 1$, such that $t_k = (kT)/M$, the Euler scheme for the process Y is given by:

$$\begin{aligned} \tilde{Y}_{t_k} &= \tilde{Y}_{t_{k-1}} + b(t_{k-1}, \tilde{Y}_{t_{k-1}})\Delta_k + a(t_{k-1}, \tilde{Y}_{t_{k-1}})\Delta W_k \\ \tilde{Y}_{t_0} &= \tilde{Y}_0 = y_0 \end{aligned}$$

where $\Delta W_k := W_{t_k} - W_{t_{k-1}}$ is a centred normal random variable with variance Δ_k, so that we have the following equality in distribution:

$$(\tilde{Y}_{t_k} \mid \tilde{Y}_{t_{k-1}} = x) \overset{\text{Law}}{=} \mathcal{N}(m_{k-1}(x), \sigma^2_{k-1}(x)) \tag{4}$$

[1] We also mention the optimal quantization website www.quantize.maths-fi.com, where one can download the optimised quadratic quantization grids of the d-dimensional Gaussian distributions $\mathcal{N}(0; I_d)$, for $N = 1$ up to 10^4 and for $d = 1, \dots, 10$. Moreover, at the same link one can also find functional quantization grids of the standard Brownian motion over the interval [0; 1], of the Brownian bridge, as well as a detailed procedure to compute grids for the (normalized) Ornstein-Uhlenbeck process and its bridge.

where:

$$m_{k-1}(x) = x + b(t_{k-1}, x)\Delta_k$$
$$\sigma^2_{k-1}(x) = [a(t_{k-1}, x)]^2\Delta_k$$

Our intention now is to use the vector quantization applied to every (one-dimensional) random variable $\tilde{Y}_{t_k}$, $k \geqslant 1$, since we know its marginal distribution conditional on $\tilde{Y}_{t_{k-1}}$. This explains the term 'marginal' of this quantization method. It can be seen in Pagès, Pham & Printems (2003) that the error made by quantizing the Euler scheme can be easily controlled, under some mild regularity assumptions on the process. The distortion function relative to $\tilde{Y}_{t_{k+1}}$, denoted D_{k+1} (recall (2)), reads:

$$D_{k+1}(\boldsymbol{x}^{k+1}) = \sum_{i=1}^{N} \int_{C_i(\boldsymbol{x}^{k+1})} (y_{k+1} - x_i^{k+1})^2 \mathbb{P}(\tilde{Y}_{t_{k+1}} \in dy_{k+1}) \quad (5)$$

where $\boldsymbol{x}^{k+1} = \{x_1^{k+1}, x_2^{k+1}, \dots, x_N^{k+1}\}$ is the quantizer at time t_{k+1} and N is the (fixed) size of the quantizer at every time step. The delicate point here is that, in order to quantize $\tilde{Y}_{t_{k+1}}$ we have to apply the Newton-Raphson method without knowing its distribution. However, by using the conditional distribution in (4) we can rewrite the distortion function (5) in terms of $\tilde{Y}_{t_k}$, thereby obtaining a recursive formula to compute the stationary quantizer. In fact, the distribution function of $\tilde{Y}_{t_{k+1}}$ can be written as follows:

$$\begin{aligned}\mathbb{P}(\tilde{Y}_{t_{k+1}} \in dy_{k+1}) &= dy_{k+1} \int_{\mathbb{R}} \phi_{m_k(y_k),\sigma_k(y_k)}(y_{k+1}) \mathbb{P}(\tilde{Y}_{t_k} \in dy_k) \\ &= dy_{k+1} \mathbb{E}[\phi_{m_k(\tilde{Y}_k),\sigma_k(\tilde{Y}_k)}(y_{k+1})]\end{aligned}$$

where $\phi_{m,\sigma}$ denotes the density function associated with a normal distribution $\mathcal{N}(m, \sigma^2)$. With this result, it is possible to compute the Hessian matrix of the distortion function. Note that we are interested in the quantization of the Euler scheme $\tilde{Y}$ that we denote by $\hat{Y}_{t_k}$, $k \geqslant 0$, so that we substitute $\tilde{Y}_{t_k}$ with $\hat{Y}_{t_k}$ in (5). Due to the discrete nature of the quantizer, the integral in (5) becomes a finite sum, thus leading to extremely fast computations. In the sequel, we will apply the recursive marginal quantization to a special local volatility model, namely the QNV model. We refer the interested reader to Pagès & Sagna (2014) for a complete background, including the analysis of the errors generated by the recursive quantization method.

The quadratic normal volatility model

The class of QNV models has drawn much attention in the financial industry due to its analytic tractability and flexibility. We will refer to the works of Blacher (2001), Ingersoll (1997), Lipton (2002) and Andersen (2011).

A QNV model is associated to an asset Y evolving as follows:

$$dY_t = (e_1 Y_t^2 + e_2 Y_t + e_3)\, dW_t, \quad Y_0 = y_0 > 0 \quad (6)$$

for some $e_1, e_2, e_3 \in \mathbb{R}$, where the Brownian motion W is taken under the risk-neutral measure. This corresponds to the SDE (3) where $b(t, y) = 0$ (that is, we consider the forward-price process) and $a(t, y) = e_1 y^2 + e_2 y + e_3$. Note that (6) includes, as special cases, Brownian motion (for $e_1 = e_2 = 0$), geometric Brownian motion (for $e_1 = e_3 = 0$) and the inverse of a three-dimensional Bessel process (for $e_2 = e_3 = 0$), which leads to a strict local martingale (we refer to Andersen (2011) and Carr, Fischer & Ruf (2013) for other technical properties of the model). Apart from technicalities, the intuition underlying (6) is that mimicking a quadratic spot volatility gives some chances to get an implied volatility curve that is able to reproduce the smile and skew effects using a parsimonious number of parameters. This is more evident in the following parameterisation (taken by Andersen (2011)):

$$dY(t) = \sigma\left(qY(t) + (1-q)y_0 + \tfrac{1}{2}s\frac{(Y(t) - y_0)^2}{y_0}\right) dW(t), \quad Y_0 = y_0 > 0 \quad (7)$$

Here $\sigma > 0$ is a proxy for the at-the-money (ATM) volatility level, q is related to the implied volatility slope (that is, q is the skew parameter) and s is a measure of the convexity of the quadratic volatility function (the vol-of-vol parameter).

■ **Vanilla options pricing** The QNV model allows for closed-form solutions for the prices of vanilla options (see also the technical appendix A.2 at http://ssrn.com/abstract=2495829, taken from Andersen (2011)). The corresponding formulas depend on the roots of the polynomial in (6). Note that, even if closed-form formulas are available for vanillas, their implementation is time-consuming and requires some care, especially in the truncation of the trigonometric series. Moreover, a calibration procedure based on these formulas should allow for the possibility of switching from the first (real roots) to the second case (complex roots) without constraints. We will see in the calibration exercise that this is a real issue. On the contrary, in the recursive marginal quantization approach, one never deals with this problem. Following the steps illustrated in the previous section, one easily computes the critical points of the distortion function together with its Hessian. In appendix A.2 of the extended version, we present the formulas for the gradient, the Hessian matrix and the weights of the quantized random variable $\hat{Y}_T$.

■ **Barrier options pricing** We focus now on barrier options. More precisely, on discrete time barrier options with daily monitoring. Indeed, although most models in the literature assume continuous monitoring of the barrier (which can lead to analytic solutions as in the Black-Scholes model), in practice most barrier options are discretely monitored. Unfortunately, this realistic setting in general does not allow for closed-form solutions. We refer the interested reader to the pricing of discrete barrier options in, for example, Kou (2003) (for an introduction to the so-called continuity correction) and to Lipton & McGhee (2002) (for a PDE approach in a universal volatility model, that leads, in some benchmark cases, to analytic solutions). We also refer to Lipton, Gal & Lasis (2014) for a survey on pricing of barrier options in local-stochastic volatility models.

To apply recursive marginal quantization to this setting, we follow the approach in Sagna (2010), where the author presents an algorithm based on optimal marginal quantization, to approximate the price of knock-out barrier options. We consider up-and-out put options. Pricing formulas in the other cases are just slight modifications of the ones we are going to present here.

Given the Euler scheme $\tilde{Y}$ for the process Y, the price of an up-and-out put option expiring at time T, with strike K and up-and-out barrier L can be approximated by:

$$P^{\text{LO}} := \mathrm{e}^{-rT}\mathbb{E}((K-\tilde{Y}_T)^+ \mathbb{1}_{\{\sup_{k=0,\dots,M}\tilde{Y}_{t_k}\leq L\}})$$
$$= \mathrm{e}^{-rT}\mathbb{E}\biggl((K-\tilde{Y}_T)^+ \prod_{k=1}^{M} G_{\tilde{Y}_{t_{k-1}},\tilde{Y}_{t_k}}(L)\biggr) \qquad (8)$$

where:

$$G_{x,y}(u) = \biggl(1-\exp\biggl\{-2M\frac{(x-u)(y-u)}{T\sigma^2(x)}\biggr\}\biggr)\mathbb{1}_{\{u\geq\max(x,y)\}}$$

and where $\sigma(\cdot)$ is the volatility function of Y. The last equality in the above equation can be obtained via an application of the so-called regular Brownian bridge method, which is connected to the knowledge of the distribution of the minimum (or the maximum) of the continuous Euler scheme $\tilde{Y}$ relative to a process Y over a time interval $[0, T]$, given its values at the discrete time observation points $0 = t_0 < t_1 < \cdots < t_M = T$ (see, for example, Glasserman 2003).

The expectation in (8) can be computed recursively, as soon as we have an approximation of the transition probability of $\tilde{Y}_{t_k}$ given $\tilde{Y}_{t_{k-1}}$. The idea now is to approximate this expectation using $\hat{Y}_{t_k}$ instead of $\tilde{Y}_{t_k}$, $k \geqslant 1$, and the transition matrix of $\hat{Y}_{t_k}$ given $\hat{Y}_{t_{k-1}}$. For all the detailed formulas we refer to appendix A.3 at http://ssrn.com/abstract=2495829.

Numerical results

In this section, we provide the first example of competitive and efficient calibration of a quantization-based method to real data and we then apply our result to the pricing of vanilla and non-vanilla derivatives. Note that, once we know the stationary grid for each time step, the pricing of a generic option becomes immediate. For example, the price at $t = 0$ of a European vanilla put option on Y with maturity T and strike K that we have N-quantized at $t = t_M = T$ with an optimal grid $\boldsymbol{y}^M = (y_1^M, \dots, y_N^M)$ and associated optimal quantizer $\hat{Y}_T$ is given by (recall (1)):

$$\mathbb{E}[(K-Y_T)^+] \cong \sum_{i=1}^{N}(K-y_i^M)^+\mathbb{P}(\hat{Y}_T = y_i^M)$$

which can be immediately computed. Note also that:

$$C_i(\boldsymbol{y}^M) = \biggl[\frac{y_{i-1}^M + y_i^M}{2}, \frac{y_i^M + y_{i+1}^M}{2}\biggr]$$

since we work in a one-dimensional setting.

The dimension of the quantization grids is taken to be constant over time, which is obviously not the optimal choice. Nevertheless, it represents a good trade-off between price precision and implementation cost. For more details on this aspect, as well as for an analysis on optimal dispatching, we refer to Pagès & Sagna (2014).

Calibration on vanillas We first test the goodness of the pricing via recursive marginal quantization. Here we use nine different strikes, equally spaced from 80% to 120% of the initial value of the underlying, and six different maturities, from two months to two years. As an error measure for this test we consider the residual norm, that is the sum of the squared differences between the model implied volatilities and the ones generated by the closed-form formulas of the previous section. We use 30-dimensional quantizers and 10 time steps for every maturity. Figure 2 shows the corresponding quantization grids. Computations are performed using MATLAB on a computer with a 2.4 GHz CPU and 8 Gb of memory. The inverse of the Hessian matrix, which is tridiagonal and symmetric, is calculated using the LU-decomposition.

The results in table A confirm the precision of prices generated by the quantization. Note that in the case of complex roots, the quantization algorithm is faster than the computation based on closed-form formulas. This fact is relevant since market data calibration typically requires complex roots, as we are going to show.

Let us now turn to real market data. Calibration is done via a standard non-linear least-squares optimiser that minimises the total calibration error in terms of the difference between model and market-implied volatilities $\sum_n(\sigma_{n,\text{market}}^{\text{imp}} - \sigma_{n,\text{model}}^{\text{imp}})^2$. Using a major provider, we take prices of European vanilla call-put options on the Dax index, as of June 19, 2014. Using the closed-form formulas, it turns out that the implied volatility smile produced by the market is fitted better when the two roots are complex. As a consequence, quantization will be faster than

2 Quantization grids for the QNV model in the case: $\sigma = 0.2$, $q = 0.5$, $s = 0.1$ (in the (7) specification)

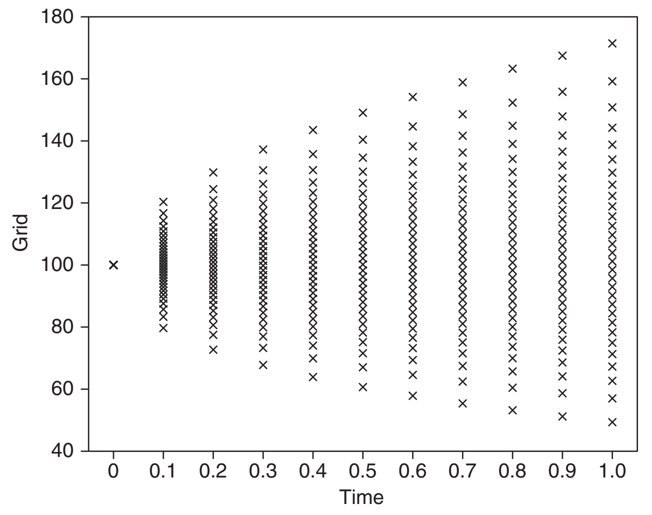

A. Comparison between pricing via closed-form formulas and quantization with $N = 30$ (CT stands for computational time)

	Analytic CT	Quantization CT	Residual norm
Real roots	0.03550 s	1.07239 s	1.34891×10^{-4}
Complex roots	10.25185 s	1.14839 s	1.79906×10^{-4}

In the case of two real roots, the parameters are taken from Andersen (2011): $\sigma = 0.2$, $q = 0.5$, $s = 0.1$, $y_0 = 100$ (in the (7) specification). We then perturb the s parameter to get two complex roots: $\sigma = 0.2$, $q = 0.5$, $s = 5$

B. Calibrated parameters of the quadratic normal volatility model

	σ	q	s
Exact formulas/long maturities	0.16019	−0.04380	26.69999
Quantization/long maturities	0.17451	0.00005	7.62015
Quantization/short maturities	0.14536	−4.67521	16.74793

Here $y_0 = 9837.63$

C. Computational times and calibration errors obtained via closed-form formulas and quantization

	Closed-form formulas	
	Computation time	Residual norm
Long maturities	339.15013 s	5.62922×10^{-4}
Short maturities	—	—
	Quantization	
	Computation time	Residual norm
Long maturities	221.15028 s	4.26904×10^{-4}
Short maturities	159.02147 s	4.00141×10^{-4}

For short maturities (from 2 to 5 months) we were not able to obtain meaningful results using the closed-form formulas

3 Implied volatility squared errors for the calibration via quantization

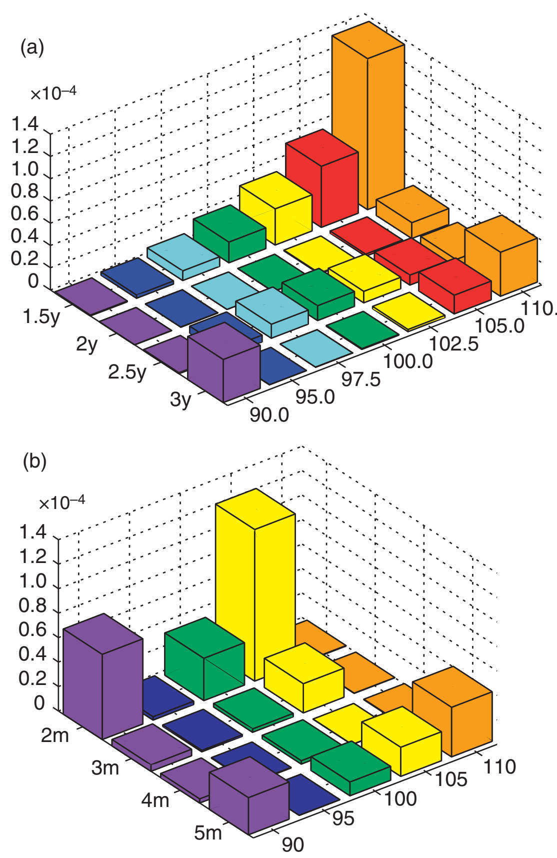

Long maturities in (a), short maturities in (b)

4 Example of fit of the implied volatility smile

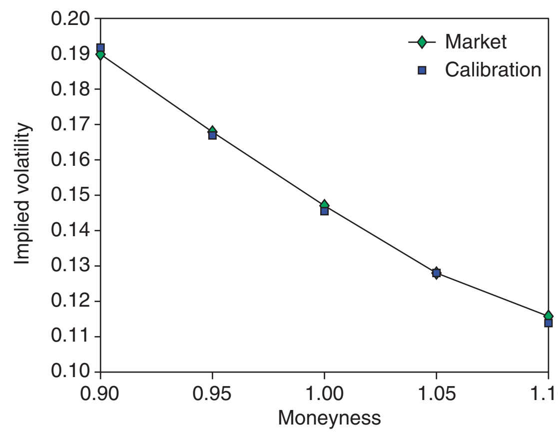

For short maturities (here the maturity is four months) we consider five strikes

D. Results on the pricing of up-and-out put options with strike $K = 100\%$ via quantization and Monte Carlo simulation

L	Benchmark price	Q price	Q error (%)	MC confidence interval
103.75	305.17096	296.70996	2.77	[295.43121,321.15975]
105	320.88575	313.47255	2.29	[311.23967,337.03611]
106.25	327.77641	320.99300	2.00	[317.02684,342.86616]
107.5	330.20446	323.76159	1.89	[319.36426,345.19677]
108.75	330.87364	324.63297	1.83	[319.90923,345.73776]
Comp. time		2.11746 s		2.71984 s

The barrier L is a percentage of the initial price. Q stands for quantization and MC stands for Monte Carlo (10^4 simulations). We consider a similar computational cost and in the last column we display the confidence interval for the corresponding MC

closed-form formulas. What is more, closed-form formulas do not perform well for short maturities, to the point that we are not able to present results of the calibration based on closed-form formulas in this case, while we note that the flexibility of the quantization approach allows us to overcome these difficulties. We therefore show the joint results of calibration via closed-form formulas and via quantization only with long maturities (from 1.5 years up to three years), while with short maturities (from two months up to five months) we only display the calibration results for the recursive marginal quantization. The calibrated parameters are displayed in table B.

With long maturities (respectively short maturities) the residual norm is given in table C, containing four maturities and seven strikes (respectively four maturities and five strikes). Overall, the quality of the fit is not excellent (see figures 3 and 4), but this is due to the particular model, which is very parsimonious (only three parameters). Nevertheless, we emphasise that despite the simplicity and the limits of the model, this represents the first successful calibration example based on quantization. Moreover, it is important to note that the procedure here illustrated is very robust since it can be easily applied to any local volatility (diffusive) model for which the Euler scheme is available.

Pricing of barrier options In order to test the goodness of this pure quantization method, we use the same data as in the previous subsection, focusing on short maturities. We fix the maturity $T = \frac{1}{3}$ and the strike $K = 100\%$ (ATM).

We compare the prices of an up-and-out put option obtained via quantization and with Monte Carlo simulation. The aim is to show that the quantization approach outperforms the Monte Carlo method in terms of computational cost. We first compute a Monte Carlo price with 10^7 simulations. The corresponding confidence interval is very sharp (about 0.3%), and we consider this price as our benchmark.

On the quantization side, we use 48-dimensional quantizers, which turn out to be a good tradeoff between precision and computational cost. Then we look for the number of paths required by the Monte Carlo that shares the same computation time required by quantization (about 2 seconds). It turns out that we need 10^4 simulations, as illustrated in table D. Note that the quantization price falls within the confidence interval of the Monte Carlo, which is quite large. What is more, Monte Carlo is much less precise than quantization when we fix the computational time for each case. Here, by precision we mean the distance between the benchmark and the quantization price, while for the Monte Carlo we mean the maximal distance between the benchmark price and the endpoints of the confidence interval.

We also perform another exercise, namely, we look for the number of simulations required by Monte Carlo in order to match the precision of the quantization method. Table E shows that Monte Carlo requires 8×10^4 simulations, which increases the computational cost of the quantization by approximately a factor of 10.

In conclusion, the quantization method is a very good alternative to Monte Carlo.

E. Results on the pricing of up-and-out put options with strike $K = 100\%$ via quantization and Monte Carlo simulation

L	Benchmark price	Q price	Q error (%)	MC confidence interval
103.75	305.17096	296.70996	2.77	[299.38595,308.44100]
105	320.88575	313.47255	2.29	[314.20251,323.31032]
106.25	327.77641	320.99300	2.00	[321.12872,330.24552]
107.5	330.20446	323.76159	1.89	[323.91769,333.03919]
108.75	330.87364	324.63297	1.83	[324.54593,333.66737]
Comp. time		2.11746 s		19.12999 s

The barrier L is a percentage of the initial price. Q stands for quantization and MC stands for the Monte Carlo (8×10^4 simulations) that shares the same precision as the Q¬method. In the last row we display the associated computation times

Conclusion

We have applied recursive marginal quantization to the local volatility model QNV to provide an alternative way to compute prices, without the numerical problems due to the real/complex nature of the roots. The procedure gives a fast way to price vanilla as well as barrier options, compared with Monte Carlo simulation. A successful calibration of the QNV model on real data shows the flexibility and the robustness of the quantization method, which can be considered a model-independent approach. Extensions of this work could include less parsimonious local volatility models, since the speed of the algorithm does not depend on the number of parameters, and pricing of structured contracts, in the spirit of Bardou, Bouthemy & Pagès (2009), who investigated the energy market. **R**

Giorgia Callegaro is an assistant professor in the Department of Mathematics at the University of Padova in Italy. Lucio Fiorin is a PhD student in the same department and Martino Grasselli is an associate professor at the department and the head of the finance lab at the Pôle Universitaire Léonard de Vinci in Paris and cofounder of Quanta Finanza Srl in Italy. **The authors are especially indebted to Abass Sagna for comments and suggestions. They also thank Daniele Marazzina and Anna Battauz for fruitful discussions.** **Email: gcallega@math.unipd.it, fiorin@math.unipd.it, grassell@math.unipd.it**

REFERENCES

Andersen L, 2011
Option pricing with quadratic volatility: a revisit
Finance and Stochastics 15(2), pages 191–219

Bardou O, S Bouthemy and G Pagès, 2009
Optimal quantization for the pricing of swing options
Applied Mathematical Finance 16, pages 183–217

Blacher G, 2001
A new approach for designing and calibrating stochastic volatility models for optimal delta-vega hedging of exotics
In *ICBI Global Derivatives, Conference Presentation*
ICBI Global Derivatives, Juan-Les-Pins

Carr P, T Fischer and J Ruf, 2013
Why are quadratic normal volatility models analytically tractable?
SIAM Journal of Financial Mathematics 4, pages 185–202

Glasserman P, 2003
Monte Carlo Methods in Financial Engineering
Springer Series, Applications of Mathematics

Graf S and H Luschgy, 2000
Foundations of quantization for probability distributions
Springer, New York

Ingersoll J, 1997
Valuing foreign exchange rate derivatives with a bounded exchange process
Review of Derivatives Research 1, pages 159–181

Kou S, 2003
On pricing of discrete barrier options
Statistica Sinica 13(4), pages 955–964

Lipton A, 2002
The vol smile problem
Risk February, pages 61–65

Lipton A, A Gal and A Lasis, 2014
Pricing of vanilla and first-generation exotic options in the local stochastic volatility framework: survey and new results
Quantitative Finance 14(11), pages 1899–1922

Lipton A and W McGhee, 2002
Universal barriers
Risk May, pages 81–85

Pagès G, 2014
Introduction to optimal vector quantization and its applications for numerics
Preprint, available at: http://hal.archives-ouvertes.fr/INSMI/hal-01034196.

Pagès G, H Pham and J Printems, 2003
Optimal quantization methods and applications to numerical problem in finance
In *Handbook of Computational and Numerical Methods in Finance*, ed. ST Rachev
Birkhauser, Boston

Pagès G and J Printems, 2005
Functional quantization for numerics with an application to option pricing
Monte Carlo Methods and Applications 11(4), pages 407–446

Pagès G and A Sagna, 2014
Recursive marginal quantization of an Euler scheme with applications to local volatility models
Preprint, available at: http://arxiv.org/abs/1304.2531

Sagna A, 2010
Pricing of barrier options by marginal functional quantization
Working Paper, available at: http://arxiv.org/abs/arXiv:1012.1037

Zühlsdorff C, 2002
Extended Libor market models with affine and quadratic volatility
Working Paper, University of Bonn, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=228250