

American quantized calibration in stochastic volatility

Giorgia Callegaro ^{*} Lucio Fiorin [†] Martino Grasselli [‡]

January 3, 2018

Abstract

We apply the recursive marginal quantization methodology to the pricing of vanillas and American style options, extending the results of Callegaro et al. (2015) to stochastic volatility models. The methodology is fast and accurate to the point that calibration can be efficiently performed. As a motivating example, we calibrate the Heston model on a book on Google stock that includes American style options.

1 Introduction

This paper constitutes the sequel of Callegaro et al. (2015). In that paper we calibrated a local volatility model using a quantization approach that had been introduced by Pagès and Sagna (2015). This technique has been further developed to be applied to pricing in financial markets in Callegaro et al. (2016), Fiorin et al. (2017) and McWalter et al. (2017). Here we extend the methodology to stochastic volatility models, in order to price vanillas and American style options. The challenge in our framework is represented by the presence of the volatility process, which increases the dimension of the problem. Starting from the Euler scheme of the stochastic volatility model, we propose a fast and accurate discretization of the asset price that improves the ones proposed in Callegaro et al. (2016) and Fiorin et al. (2017) (see Remark 3.1 below). We first focus on the pricing of vanillas and we show that pricing can be efficiently performed. Then, we exploit the idea in Bally et al. (2005), where it is shown that pricing of American options can also be performed through a backward procedure, like in a multinomial lattice. As a result, we provide the first calibration example of the Heston (1993) model using a book of real data that includes American style options. Of course, for this affine model one could also calibrate on vanillas using the Fourier methodology of Carr and Madan (1999). What is more, we emphasize that our methodology is very flexible insofar it applies to any stochastic volatility model (well beyond the family of affine models), including the SABR model and many others.

The paper is organized as follows: in Section 2 we give a quick application-oriented overview of the vector quantization methodology that can be also found in Callegaro et al. (2015). Section 3 extends the marginal quantization method to the class of stochastic volatility models and illustrates the idea of the algorithm. Further details can be found in Callegaro et al. (2016)

^{*}Department of Mathematics, University of Padova, via Trieste 63, 35121 Padova, Italy. Email: gcallega@math.unipd.it

[†]Department of Mathematics, University of Padova, via Trieste 63, 35121 Padova, Italy. Email: fiorin@math.unipd.it

[‡]Department of Mathematics, University of Padova (Italy) and DVRC Devinci Research Center, Pole Universitaire Léonard de Vinci, Paris La Défense (France) Acknowledgments: we are especially indebted to Abass Sagna for comments and suggestions.

with a slightly different method. Section 4 illustrates our numerical results on the Heston (1993) model, with particular emphasis on the calibration exercise on real data including American style options. Section 5 concludes.

2 What you need to know on quantization

In this section we provide a brief overview of optimal quadratic quantization and of recursive marginal quantization (henceforth RMQ). We focus here on vector quantization: we will only deal with discretization of random vectors, thus avoiding functional quantization (discretization of stochastic processes in the space of trajectories). This will be sufficient to discretize a stochastic process via RMQ. We refer to Graf and Luschgy (2000) and Pagès (2015) for vector quantization and to Pagès and Sagna (2015) for the first paper on RMQ ¹.

Optimal quadratic vector quantization answers to the following question:

how is it possible to optimally (in an L^2 -sense) approximate a continuous random vector X by a discrete one, \widehat{X} , taking a finite number of values?

The interest in such a discretization \widehat{X} is obvious, since integrals of the form $\mathbb{E}[h(X)]$, for sufficiently regular functions h , would be approximated by finite sums $\mathbb{E}[h(\widehat{X})]$.

Let us be more precise. We consider a real valued random variable X defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and having probability distribution \mathbb{P}_X . We suppose that it has finite second order moment.

A quantization grid of level N , $N \geq 1$, is a subset of \mathbb{R} , $\Gamma^N = \{x_1, \dots, x_N\}$, of size at most N having pairwise distinct components. A quantization function, or *quantizer*, is a Γ -valued (we drop N in the notation Γ^N since we will consider N as fixed) Borel function $q : \mathbb{R} \rightarrow \Gamma$ and N -quantizing X means projecting X on the grid Γ following the closest neighbor rule

$$q(X) = \text{Proj}_\Gamma(X) := \sum_{i=1}^N x_i \mathbb{1}_{C_i(\Gamma)}(X) \quad (1)$$

where $(C_i(\Gamma))_{1 \leq i \leq N}$ is a Borel partition of $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, also known as Voronoi partition, satisfying $C_i(\Gamma) \subset \{\xi \in \mathbb{R} : |\xi - x_i| = \min_{i \leq j \leq N} |\xi - x_j|\}$, $i = 1, \dots, N$. In what follows we use the notation \widehat{X}^Γ or \widehat{X} (when no ambiguity is possible with respect to the grid) to denote the Voronoi Γ -quantization of X : $\widehat{X}^\Gamma = \widehat{X} = q(X)$.

The L^2 -error induced by quantization is given by

$$e_N(X, \Gamma) := \|X - q(X)\|_2 = \left\| \min_{1 \leq i \leq N} |X - x_i| \right\|_2$$

where $\|X\|_2 := [\mathbb{E}(|X|^2)]^{1/2}$ is the usual L^2 -norm. Optimally quadratic quantizing X boils down to looking for a grid Γ with size at most N which minimizes the *distortion function* defined here below (see (Graf and Luschgy, 2000, Equation (3.4))).

Definition 2.1. *Let X be a real valued random vector belonging to $L^2(\mathbb{P})$. The L^2 -distortion function is a positive valued function defined on \mathbb{R}^N by*

$$D : (x_1, x_2, \dots, x_N) \longmapsto \mathbb{E} \left[\min_{1 \leq i \leq N} |X - x_i|^2 \right] = e_N(X, \Gamma)^2 \quad (2)$$

for $\Gamma = (x_1, x_2, \dots, x_N)$.

¹A further reference to be mentioned is the quantization website: <http://www.quantize.maths-fi.com>, where grids of the d -dimensional Gaussian distributions $\mathcal{N}(0; I_d)$, for $N = 1$ up to 10^4 and for $d = 1, \dots, 10$ are available. At the same link one also finds functional quantization grids (ie, trajectories) for the standard Brownian motion over $[0; 1]$ and for the Brownian bridge.

What about existence and uniqueness of an optimal grid? It is possible to show (see eg (Pagès, 2015, Prop. 1.1)) that if X belongs to $L^2(\mathbb{P})$, then the distortion function D attains (at least) one minimum Γ^* . The grid Γ^* and Proj_{Γ^*} are called *optimal quadratic quantizers*, respectively. In the case when $\text{card}(\text{supp}(\mu)) \geq N$, then Γ^* has pairwise distinct components. Moreover $\lim_{N \rightarrow +\infty} e_N(X) = 0$ and the convergence rate is given by the Zador theorem (see Graf and Luschgy (2000))

$$\min_{\Gamma, |\Gamma|=N} \|X - q(X)\|_2 = \min_{\Gamma, |\Gamma|=N} e_N(X, \Gamma) = Q_2(\mathbb{P}_X)N^{-1} + o(N^{-1})$$

where $Q_2(\mathbb{P}_X)$ is a nonnegative constant. Note that, when computing $\mathbb{E}[f(X)]$, where X is the value at maturity of the underlying, f a Lipschitz (payoff) function, and we denote by \hat{X}^Γ the quantization of X on the grid Γ , then, by Jensen's inequality, $|\mathbb{E}[f(X)] - \mathbb{E}[f(\hat{X}^\Gamma)]| \leq [f]_{\text{Lip}} e_N(X, \Gamma)$. In particular, the error of the quantized price decays at rate of $\frac{1}{N}$, in contrast to the Monte Carlo method, which has an error, ruled by the Central Limit Theorem, of order $\frac{1}{\sqrt{N}}$.

Remark 2.2. *As soon as \mathbb{P}_X is absolutely continuous with respect to a log-concave density, then there exists exactly one optimal quantization grid of level N (see (Graf and Luschgy, 2000, Theorem. 5.1)).*

The last crucial question is how to obtain an optimal quantizer. It is known (see (Graf and Luschgy, 2000, Lemma 4.10) or (Pagès, 2015, Prop. 1.1)) that the distortion function D is differentiable at any N -tuple having pairwise distinct components $\Gamma = \{x_1, \dots, x_N\}$ and the gradient is given by

$$\nabla D(x_1, \dots, x_N) = 2 \left(\int_{C_i(\Gamma)} (x_i - \xi) d\mathbb{P}_X(\xi) \right)_{1 \leq i \leq N} = 2 \left(\mathbb{E}[\mathbb{1}_{X \in C_i(\Gamma)}(x_i - X)] \right)_{1 \leq i \leq N} \quad (3)$$

Based on the existence of the gradient of the distortion function, many stochastic algorithms, like the gradient descent and fixed point procedures, have been developed (see (Pagès, 2015, Section 3)). Moreover, when the gradient itself is differentiable, it is possible to apply the classical Newton-Raphson procedure as it turns out that the determination of optimal quantizers boils down to the solution of

$$\mathbb{E}[\mathbb{1}_{X \in C_i(\Gamma)} X] - x_i \mathbb{P}(X \in C_i(\Gamma)) = 0 \quad \forall i \in \{1, \dots, N\}$$

known as *Master Equation*. If we know the density of the random variable X , then it is possible to write all the components of the system in closed form, and its solution is trivial. Of course, this is no more the case when we consider X as the price of an asset at maturity, since the density of the process is typically unknown (except in trivial cases).

In order to exploit the information of the asset distribution, Pagès and Sagna (2015) introduced the recursive marginal quantization to discretize a stochastic process Y , in dimension one, using a Euler-Maruyama scheme. The essence of RMQ lies on the knowledge of the conditional law of $(Y_{t_k} | Y_{t_{k-1}})$, $k = 1, \dots, M$, which allows to recursively obtain (via a Newton-Raphson procedure) the quantizers for the marginals $\{(Y_{t_k})\}_{k=0, \dots, M}$ based on an explicit derivation of the gradient and of the Hessian of the distortion function. In the following section we provide some details on RMQ applied to our multi dimensional setting, namely for stochastic volatility models.

3 Recursive quantization of stochastic volatility models

Let us consider a continuous time Markov process $Y = (S, V)$ corresponding to the pair (price, volatility), whose evolution is specified by

$$\begin{pmatrix} dS_t \\ dV_t \end{pmatrix} = \begin{pmatrix} \mu_S(t, S_t, V_t) \\ \mu_V(t, V_t) \end{pmatrix} dt + \begin{pmatrix} \Sigma_{SS}(t, S_t, V_t) & \Sigma_{SV}(t, S_t, V_t) \\ \Sigma_{VS}(t, V_t) & \Sigma_{VV}(t, V_t) \end{pmatrix} dW_t, \quad \begin{pmatrix} S_0 = s_0 \\ V_0 = v_0 \end{pmatrix} \quad (4)$$

where W is a 2 dimensional Brownian motion and $\mu_S, \mu_V, \Sigma_{SS}, \Sigma_{SV}, \Sigma_{VS}, \Sigma_{VV}$ are such that the SDE (4) is well defined. In this section we only introduce our novel methodology, so that we do not focus on the technical conditions which would ensure existence of a strong solution to the very general SDE (4). Nevertheless, in the following Section 4 we will provide more details on existence of a strong solution in the case of the model in use, ie, in Heston model.

Fix a time horizon T and a time grid with constant step size Δ and length M , ie $t_k = \frac{k}{M}T$ for $k = 0, \dots, M$, so that the Euler scheme for the pair (S, V) reads

$$\begin{pmatrix} \tilde{S}_{k+1} \\ \tilde{V}_{k+1} \end{pmatrix} = \begin{pmatrix} \tilde{S}_k \\ \tilde{V}_k \end{pmatrix} + \begin{pmatrix} \mu_S(t_k, \tilde{S}_k, \tilde{V}_k) \\ \mu_V(t_k, \tilde{V}_k) \end{pmatrix} \Delta + \begin{pmatrix} \Sigma_{SS}(t_k, \tilde{S}_k, \tilde{V}_k) & \Sigma_{SV}(t_k, \tilde{S}_k, \tilde{V}_k) \\ \Sigma_{VS}(t_k, \tilde{V}_k) & \Sigma_{VV}(t_k, \tilde{V}_k) \end{pmatrix} \sqrt{\Delta} \tilde{W}_k, \quad \begin{pmatrix} \tilde{S}_0 = s_0 \\ \tilde{V}_0 = v_0 \end{pmatrix} \quad (5)$$

where $\tilde{W}_k := \frac{W_{t_{k+1}} - W_{t_k}}{\sqrt{\Delta}}$ is a bivariate standard Gaussian random variable having mean $(0, 0)$ and variance the identity matrix 2×2 . Similarly to the one dimensional case (see Callegaro et al. (2015)), we have that

$$\left(\begin{pmatrix} \tilde{S}_{k+1} \\ \tilde{V}_{k+1} \end{pmatrix} \middle| \begin{pmatrix} \tilde{S}_k \\ \tilde{V}_k \end{pmatrix} = \begin{pmatrix} s \\ v \end{pmatrix} \right) \stackrel{\text{Law}}{=} \mathcal{N}(\mu(t_k, s, v), \Sigma(t_k, s, v)) \quad (6)$$

where $\mu(t_k, s, v) = \begin{pmatrix} s + \mu_S(t_k, s, v)\Delta \\ v + \mu_V(t_k, v)\Delta \end{pmatrix}$ and

$$\Sigma(t_k, s, v) = \Delta \begin{pmatrix} (\Sigma_{SS}^2 + \Sigma_{SV}^2)(t_k, s, v) & (\Sigma_{SS}\Sigma_{VS} + \Sigma_{SV}\Sigma_{VV})(t_k, s, v) \\ (\Sigma_{SS}\Sigma_{VS} + \Sigma_{SV}\Sigma_{VV})(t_k, s, v) & (\Sigma_{VV}^2 + \Sigma_{VS}^2)(t_k, v) \end{pmatrix}$$

As the dynamics of the volatility factor can be written independently of the price process, one can use Callegaro et al. (2016) in order to quantize the volatility. The difficult point is the construction of the distortion function relative to \tilde{S}_{k+1} , that we denote by D_{k+1} , which depends also on the volatility process V . Let us fix a grid $\mathbf{x}^{k+1} = \{x_1^{k+1}, \dots, x_N^{k+1}\}$. The distortion function for \tilde{S}_{k+1} reads

$$D_{k+1}(\mathbf{x}^{k+1}) = \sum_{i=1}^N \int_{C_i(\mathbf{x}^{k+1})} (s_{k+1} - x_i^{k+1})^2 \mathbb{P}(\tilde{S}_{k+1} \in ds_{k+1}) \quad (7)$$

Having quantized the pair $(\tilde{S}_\ell, \tilde{V}_\ell)$, $\ell = 0, \dots, k$, we can write the distribution of \tilde{S}_{k+1} , using (6). This gives the recursive formula to compute the quantizers

$$\mathbb{P}(\tilde{S}_{k+1} \in ds_{k+1}) = ds_{k+1} \int_{\mathbb{R}} \int_{\mathbb{R}} \phi_{\mu_S(t_k, s_k, v_k), (\Sigma_{SS}^2 + \Sigma_{SV}^2)^{\frac{1}{2}}(t_k, s_k, v_k)}(s_{k+1}) \mathbb{P}(\tilde{S}_k \in ds_k, \tilde{V}_k \in dv_k) \quad (8)$$

where ϕ is the density function of a scalar Normal distribution with mean $\mu_S(t_k, s_k, v_k)$ and variance $(\Sigma_{SS}^2 + \Sigma_{SV}^2)(t_k, s_k, v_k)$. It is then possible to compute the gradient and the Hessian matrix of the distortion D_{k+1} , in order to find (numerically) its minima. Note that, once we have quantized $(\tilde{S}_\ell, \tilde{V}_\ell)$, $\ell = 0, \dots, k$, the integrals in (8) become a finite sum, thus leading to extremely fast computations. In the sequel, we will apply this methodology to the celebrated Heston (1993) stochastic volatility model.

Remark 3.1. *Note that this result is different from the one in Callegaro et al. (2016), where we considered the Euler scheme separately for S and V . Indeed, using here the Euler scheme for the pair (Price, Volatility), it is possible to write the density of the price process more accurately, by taking into account properly the correlation between the price and the volatility. This allows us to price efficiently vanillas and American style options and then to calibrate the model.*

Remark 3.2. *Our technique can be applied to any stochastic volatility model for which an Euler discretization scheme is available, as also shown in Callegaro et al. (2016). Moreover, since recursive marginal quantization has been applied on top of a Milstein discretization scheme in McWalter et al. (2017) to discretize a geometric Brownian motion and a constant elasticity of variance process, we deem that an interesting research direction could be the extension of our methodology using a higher order discretization scheme.*

4 Numerical results for the Heston model

4.1 Pricing of vanillas

The Heston (1993) model assumes the following risk-neutral dynamics for the pair (S, V)

$$\begin{aligned} \frac{dS_t}{S_t} &= rdt + \sqrt{V_t}(\rho dW_t^1 + \sqrt{1 - \rho^2}dW_t^2) \\ dV_t &= \kappa(\theta - V_t)dt + \xi\sqrt{V_t}dW_t^1 \end{aligned}$$

where W^1 and W^2 are two independent standard Brownian motions and where r is the risk free interest rate, θ is the long run average price variance, κ is the rate at which the variance V reverts to θ , ρ is the correlation between the asset and the instantaneous variance and ξ is the vol of vol parameter, which determines the volatility of the variance process. We assume that r, κ, θ and ξ are strictly positive. In this case, the above system of SDEs admits a strong solution (see eg (Andersen and Piterbarg, 2007, Section 2)). The components of the Euler scheme here read

$$\mu(t_k, s, v) = \begin{pmatrix} s + rs\Delta \\ v + \kappa(\theta - v)\Delta \end{pmatrix} \quad \Sigma(t_k, s, v) = \Delta \begin{pmatrix} s^2v & \rho\xi sv \\ \rho\xi sv & \xi^2v \end{pmatrix}$$

We first compare the pricing of European calls obtained with our quantization methods with the ones provided by the Fourier based methodology as in Carr and Madan (1999), that we take as our benchmark. In Table 1 we display the prices together with the errors (relative error and absolute difference of implied volatilities) for European call options with maturity $T = 1$ year. The strike K is in percentage of the initial price $S_0 = 100$. ‘‘Q’’ stands for quantization, where we considered $N = 20$ points for the underlying, $N = 10$ points for the volatility process and $M = 12$ time steps. This choice for N, M represents a good compromise between precision and computational time.

	Bench. price	Q price	Rel error (%)	IV difference ($\times 10^{-3}$)
$K = 80$	24.91	24.99	0.32	3.87
$K = 85$	20.75	20.81	0.29	2.38
$K = 90$	16.75	16.79	0.24	1.33
$K = 95$	12.96	13.00	0.31	1.16
$K = 100$	9.43	9.45	0.21	0.92
$K = 105$	6.24	6.25	0.16	0.75
$K = 110$	3.50	3.52	0.57	0.72
$K = 115$	1.45	1.46	0.69	0.84
$K = 120$	0.40	0.40	0.75	0.89

Table 1: Results on the pricing of European call options via quantization and via the FFT approach of Carr and Madan (1999). The strike K is in percentage of the initial price $S_0 = 100$. The maturity is $T = 1$ year and $r = 0.04$. The parameters are those of the calibration exercise in Table 3. “Q” stands for quantization, where we considered grids of 20 points for the underlying, 10 points for the volatility process and 12 time steps. For the quantization, the computational times is less than 2 seconds for the whole book.

Results were obtained with *MATLAB*, on a laptop with a 2.7 GHz CPU and 8 Gb of memory. The computational time to get all the prices is less than two seconds. The results in Table 1 show the efficiency of our methodology, which is fast and accurate.

4.2 Pricing of American options

Many numerical methods computing prices of such options in the Heston model have been developed. They typically approximate the solution to the partial differential equation under early exercise constraints by ad hoc finite difference discretization grids (see eg Ikonen and Toivanen (2008)) or by exploiting recombining two dimensional lattices (see eg Beliaeva and Nawalkha (2010) or Vellekoop and Nieuwenhuis (2009)) or by means of Fourier based methods, such as the Wiener-Hopf factorization (see eg Boyarchenko and Levendorskiy (2013)) or the Cosine method (see eg Fang and Oosterlee (2011)).

An interesting feature of our quantization approach is that from the Euler-Maruyama discretization scheme one can directly deduce some information on the transition probability of the asset price, so that the pricing of American options can be performed through a backward procedure on a multinomial lattice, as illustrated in Bally et al. (2005). Let us assume that the set of possible exercise times of an American put option of maturity T is finite, meaning that we are approximating the American option value with the price of a Bermudan option, where the exercise times are $t_k = \frac{kT}{M}$, with $k = 1, \dots, M$. The parameters for the Heston model are the same as in the case of European options of the previous subsection. We show in Table 2 the comparison of the quantization price with a benchmark for the Heston model provided by the method of Vellekoop and Nieuwenhuis (2009) (this latter method is implemented, for example, in PREMIA, for more details see <https://www.rocq.inria.fr/mathfi/Premia/>). For the benchmark price we consider 20 points for the discretization of V , 200 points for the discretization of S and 12 time steps.

Let us also point out that, in our numerical examples, the early exercise premium is always positive at any node of the discretization tree. In general, indeed, any approximation of an American option price might not be arbitrage free, in that the approximated American price may be lower than the price of the corresponding European option. Therefore, testing positivity of the early exercise premium is the most important check to be done.

The quantization approach leads to accurate prices and takes 1.9676 seconds to compute all the put prices. This opens the door to the possibility of calibrating a book including both European and American options, as we are going to do in the next subsection.

	Bench. price	Q price	Rel error (%)
$K = 80$	1.81	1.78	1.75
$K = 85$	2.48	2.50	0.69
$K = 90$	3.33	3.27	1.65
$K = 95$	4.42	4.47	1.13
$K = 100$	5.80	5.75	0.84
$K = 105$	7.61	7.62	0.10
$K = 110$	10.13	10.17	0.39
$K = 115$	14.79	14.87	0.53
$K = 120$	19.76	19.85	0.46

Table 2: Results on the pricing of American Put options via quantization and the method of Vellekoop and Nieuwenhuis (2009). The strike K is in percentage of the initial price $S_0 = 100$. The maturity is $T = 1$ year, $r = 0.04$. Q stands for quantization. The parameters of the Heston model are those of the calibration exercise in Table 3. The quantization grids are those used for the pricing in Table 1.

4.3 Calibration of American options to real data

In order to show the effectiveness of our pricing technique, we calibrate the Heston model to a book of American option prices on the GOOG Google stock as of date April, 26th 2017. The book includes 4 maturities (from 3 months to 14 months) with 12 strikes, ranging from 85% to 115% of the spot price, for a total of 96 options. Table 3 shows that the parameters calibrated with the quantization are in line with the ones found by the FFT methodology of Carr and Madan (1999). Res Norm indicates the average square error on implied volatilities (IV), defined as

$$\text{Res Norm} := \frac{1}{\#\text{strikes} \#\text{maturities}} \sum_{\ell=1}^{\#\text{strikes}} \sum_{k=1}^{\#\text{maturities}} (IV_{\ell,k}^{\text{market}} - IV_{\ell,k}^{\text{model}})^2$$

The implied volatility in the European case is the vol parameter one plugs into the Black Scholes formula to obtain the market price. The situation is different in the case of American options, as there is no industry-standard model for pricing these options, even under Black-Scholes assumptions. The available pricing models include binomial and trinomial trees of various sorts, the Barone-Adesi and Whaley approximation and its variants, the Longstaff and Schwartz algorithm, and many others. Here we follow the Bloomberg convention that adopts a binomial tree pricer. For an implementation of the method, see for example <http://westclintech.com/SQL-Server-Options-Functions/SQL-Server-Binomial-American-Implied-Volatility-function>.

The numerical results presented in this subsection show that our procedure is also robust from the calibration point of view.

	Quantization
ρ	-0.9250
κ	0.1269
ξ	0.4058
θ	0.1922
V_0	0.0319
Res Norm	5.2255×10^{-5}

Table 3: Heston (1993) model calibrated on a book of American Call and Put options on the GOOG Google stock as of date April, 26th 2017. The book includes 4 maturities (from 3 months to 14 months) with 12 strikes, ranging from 85% to 115% of the spot price $S_0 = 871.71$, for both Call and Put options, for a total of 96 options. Res Norm indicates the average square error on implied volatilities. In analogy with the calibration on vanilla options for the Heston model, see eg Da Fonseca and Grasselli (2011), note that also here the calibrated parameters do not satisfy the Feller's condition $2\kappa\theta \geq \xi^2$.

Figure 1 displays the implied volatility squared errors for the calibration on American Call options (resp. left hand side) and Put options (resp. right hand side). Overall, the error is in line with the performance of the Heston (1993) model, namely with a resnorm around 10^{-5} .

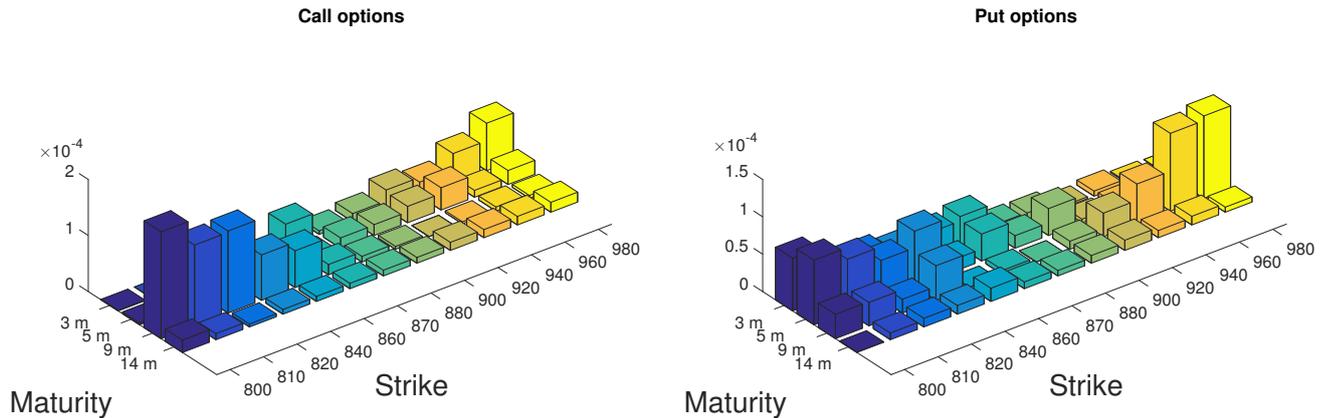


Figure 1: Implied volatility squared errors for the calibration on a book of American Call and Put options on the GOOG Google stock as of date April, 26th 2017 using the quantization method. American Call options on the left, American Put options on the right.

Finally, Figure 2 shows the quantization grids for the asset price in the Heston (1993) model for the parameters' set found in the calibration.

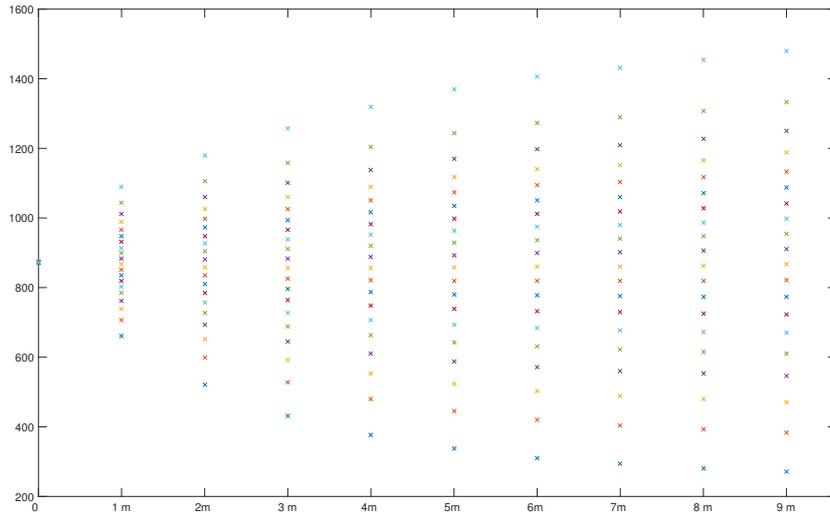


Figure 2: Quantization grids for the asset price in the Heston (1993) model for the parameter set found in the calibration exercise on the GOOG Google stock as of date April, 26th 2017, with $N = 20$ points at every time step and maturity 9 months.

5 Conclusion

In this paper we introduced a new efficient methodology to price options in a stochastic volatility environment through a quantization approach, which only assumes that a (Euler-Maruyama or more sophisticated) discretization scheme for the pair (price, volatility) is available. We have applied the methodology to the Heston (1993) model, for which an efficient benchmark is represented by the Fourier approach. However, we emphasize that our approach is flexible enough to consider also the pricing of American-style options, for which the Fourier technology is almost useless. As a motivating example, we have produced a calibration exercise on real data of the Heston model using a book of options that includes American style contracts. This can be useful, for example, in view of reconstructing the volatility of the index starting from the implied volatility of its constituents, for which only American style options are typically available.

References

- Andersen, L. and Piterbarg, V. (2007). Moment explosions in stochastic volatility models. *Finance and Stochastics*, 11(1):29–50.
- Bally, V., Pagès, G., and Printems, J. (2005). A quantization tree method for pricing and hedging multidimensional American options. *Mathematical Finance*, 15(1):119 – 168.
- Beliaeva, N. and Nawalkha, S. (2010). A simple approach to pricing American options under the Heston stochastic volatility model. *The Journal of Derivatives*, 17(4):25–43.
- Boyarchenko, S. and Levendorskiy, S. (2013). American options in the Heston model with stochastic interest rate and its generalizations. *Applied Mathematical Finance*, 20(1):26–49.

- Callegaro, G., Fiorin, L., and Grasselli, M. (2015). Quantized calibration in local volatility. *Risk Magazine*, 28(4):62 – 67.
- Callegaro, G., Fiorin, L., and Grasselli, M. (2016). Pricing via recursive quantization in stochastic volatility models. *Quantitative Finance*, Issue 6:855–872.
- Carr, P. and Madan, D. B. (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2:61–73.
- Da Fonseca, J. and Grasselli, M. (2011). Riding on the Smiles. *Quantitative Finance*, 11:1609–1632.
- Fang, F. and Oosterlee, C. W. (2011). A fourier-based valuation method for bermudan and barrier options under heston’s model. *SIAM Journal on Financial Mathematics*, 2(1):439–463.
- Fiorin, L., Sagna, A., and Pagès, G. (2017). Product Markovian quantization of an \mathbb{R}^d -valued Euler scheme of a diffusion process with applications to finance. *Preprint*.
- Graf, S. and Luschgy, H. (2000). *Foundations of quantization for probability distributions*. Springer, New York.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–343.
- Ikonen, S. and Toivanen, J. (2008). Efficient numerical methods for pricing American options under stochastic volatility. *Numerical Methods for Partial Differential Equations*, 24(1):104–126.
- McWalter, T., Rudd, R., Kienitz, J., and Platen, E. (2017). Recursive Marginal Quantization of Higher-Order Schemes. *Working Paper, available on SSRN*.
- Pagès, G. (2015). Introduction to vector quantization and its applications for numerics. *ESAIM: proceedings and surveys*, 48:29–79.
- Pagès, G. and Sagna, A. (2015). Recursive marginal quantization of the Euler scheme of a diffusion process. *Applied Mathematical Finance*, 22(5):463–498.
- Vellekoop, M. and Nieuwenhuis, H. (2009). A tree-based method to price American options in the Heston model. *Journal of Computational Finance*, 13(1):1 – 21.