

## WEIGHTED MONTE CARLO: A NEW TECHNIQUE FOR CALIBRATING ASSET-PRICING MODELS

MARCO AVELLANEDA, ROBERT BUFF, CRAIG FRIEDMAN,  
NICOLAS GRANDECHAMP, LUKASZ KRUK and JOSHUA NEWMAN

*Courant Institute of Mathematical Sciences, New York University,  
251 Mercer Street, New York, NY 10012, USA*

Received 2 September 1999

Accepted 4 January 2000

A general approach for calibrating Monte Carlo models to the market prices of benchmark securities is presented. Starting from a given model for market dynamics (price diffusion, rate diffusion, etc.), the algorithm corrects price-misspecifications and finite-sample effects in the simulation by assigning “probability weights” to the simulated paths. The choice of weights is done by minimizing the Kullback–Leibler relative entropy distance of the posterior measure to the empirical measure. The resulting ensemble prices the given set of benchmark instruments exactly or in the sense of least-squares. We discuss pricing and hedging in the context of these weighted Monte Carlo models. A significant reduction of variance is demonstrated theoretically as well as numerically. Concrete applications to the calibration of stochastic volatility models and term-structure models with up to 40 benchmark instruments are presented. The construction of implied volatility surfaces and forward-rate curves and the pricing and hedging of exotic options are investigated through several examples.

### 1. Introduction

According to basic asset-pricing theory, security prices should be equal to the expectations of their discounted cashflows under a suitable probability measure. This “risk-neutral” measure represents the economic value of consuming one unit of account on a given future date and state of the economy. A risk-neutral probability implemented in the context of a specific market is often called a *pricing model*. It is natural to require that a pricing model reproduce correctly the prices of liquid instruments which are actively traded. This ensures that “off-market”, less liquid instruments are realistically priced by the model.<sup>a</sup>

Here, we consider pricing models based on Monte Carlo (MC) simulations of future market scenarios (“paths”).<sup>b</sup> Prices are computed by averaging discounted

<sup>a</sup>Throughout this paper, a pricing model refers to a model for pricing less liquid instruments relatively to more liquid ones (the benchmarks) in the context of a particular market. This type of financial model is used by most large investment banks to manage their positions.

<sup>b</sup>See Dupire (1998) for an up-to-date collection of papers by academics and practitioners on Monte Carlo methods in finance.

cashflows over the different paths. We shall be concerned with the *calibration* of such models, i.e., with specifying the statistics of the sample paths in such a way that the model matches the prices of benchmark instruments traded in the market.

Most calibration procedures rely on the existence of explicit formulas for the prices of the benchmark instruments. The unknown parameters of the underlying stochastic process are found by inverting such pricing formulas, either exactly or in the sense of least-squares. Unfortunately, in Monte Carlo simulations, this method may not be accurate enough due to sampling errors (the “finite sample effect”). Furthermore, closed-form solutions for prices may not always be available or easy to code. In the latter case, fitting the model to market prices implies searching the parameter space through direct simulation, which is a computationally expensive proposition.

This paper considers an alternative, non-parametric approach for calibrating Monte Carlo models and applies it to several practical situations. The main idea behind our method is to put the emphasis on determining directly the risk-neutral probabilities of the future states of the market, as opposed to finding the parameters of the differential equations used to generate the paths for the state-variables.

One way to motivate our algorithm is to observe that Monte Carlo simulations can be divided (somewhat arbitrarily) into two categories: those that are *uniformly weighted* and those that are *non-uniformly weighted*. To wit, consider a set of sample paths, denoted by  $\omega_1, \dots, \omega_\nu$ , generated according to some simulation procedure. By definition, a uniformly weighted simulation is such that all sample paths are assigned the same probability. Thus, a contingent claim that pays the holder  $h_i$  dollars if the path  $\omega_i$  occurs, has model value

$$\Pi_h = \frac{1}{\nu} \sum_{i=1}^{\nu} h_i. \quad (1)$$

A non-uniformly weighted simulation is one in which the probabilities are not necessarily equal. Suppose that we assign, respectively, probabilities  $p_1, \dots, p_\nu$  to the different paths. The value of the contingent claim according to the corresponding “non-uniformly weighted” simulation is

$$\Pi_h = \sum_{i=1}^{\nu} h_i p_i. \quad (2)$$

Our approach is based on non-uniformly weighted simulations. First, we simulate a large number of paths of a stochastic process followed by the state-variables (prices, rates, etc.) under a prior distribution. Second — and here we depart from the conventional Monte Carlo method — we assign a different probability to each path. Probabilities are determined in such a way that (i) the expected values of the discounted cashflows of benchmark instruments coincide (either exactly or within tolerance) with the market prices of these securities and (ii) they are as close as possible to uniform probabilities ( $p_i = 1/\nu$ ) corresponding to the simulated prior.

This method allows us to incorporate market information in two stages. The first step gives a prior probability measure that corresponds to our best guess for the risk-neutral measure given the information available. This guess may involve real statistics, such as estimates of rates of return, historical volatilities, correlations. It may also use parameters which are implied from market prices (implied volatilities, cost-of-carry, etc). In other words, the path simulation is used to construct a “backbone” or “prior” for the model which incorporates econometric or market-implied data. The second step has two purposes: it reconciles the econometric/prior information with the prices observed at any given time and also corrects finite-sample errors on the prices of the benchmark instruments which arise from the Monte Carlo simulation.

We denote the mid-market prices of the  $N$  benchmark instruments by  $C_1, \dots, C_N$  and represent the present values of the cashflows of the  $j$ th benchmark along the different paths by

$$g_{1j}, g_{2j}, \dots, g_{\nu j} \quad j = 1, \dots, N. \quad (3)$$

The price relations for the benchmark instruments can then be written in the form

$$\sum_{i=1}^{\nu} p_i g_{ij} = C_j, \quad j = 1, \dots, N, \quad (4)$$

where  $(p_1, \dots, p_{\nu})$  are the probabilities that we need to determine. Generically, this (linear) system of equations admits infinitely many solutions because the number of paths  $\nu$  is greater than the number of constraints.<sup>c</sup> The criterion that we propose for finding the calibrated probability measure is to minimize the *Kullback–Leibler relative entropy* of the non-uniformly sampled simulation with respect to the prior. Recall that if  $p_1, \dots, p_{\nu}$  and  $q_1, \dots, q_{\nu}$  are probability vectors on a probability space with  $\nu$  states, the relative entropy of  $p$  with respect to  $q$  is defined as

$$D(p|q) = \sum_{i=1}^{\nu} p_i \log \left( \frac{p_i}{q_i} \right). \quad (5)$$

In the case of a Monte Carlo simulation with  $q_i = 1/\nu \equiv u_i$  we have<sup>d</sup>

$$D(p|u) = \log \nu + \sum_{i=1}^{\nu} p_i \log p_i. \quad (6)$$

We minimize this function under the linear constraints implied by (4). To this effect, we implement a dual, or Lagrangian, formulation which transforms the problem into an unconstrained minimization over  $N$  variables. Minimization of the dual-objective function is made with L-BFGS [Byrd *et al.* (1994)], a gradient-based quasi-Newton optimization routine.

<sup>c</sup>It is also possible that the system of equations admits no solutions if the prior is inadequate or if the prices give rise to an arbitrage opportunity. We shall not dwell on this here.

<sup>d</sup>We shall denote the uniform probability vector by  $u$ , i.e.,  $u = (1/\nu, \dots, 1/\nu)$ .

The use of minimization of relative entropy as a tool for computing Arrow–Debreu probabilities was introduced by Buchen and Kelly (1996) and Gulko (1995, 1996) for single-period models. Other calibration methods based on minimizing a least-squares penalization function were proposed earlier by Rubinstein (1994) and Jackwerth and Rubinstein (1995). Samperi (1997), Avellaneda *et al.* (1997) and Avellaneda (1998) generalized the minimum-entropy method to intertemporal lattice models and diffusions. More recently, Laurent and Leisen (1999) considered the case of Markov chains. These studies suggest that this is a computationally feasible approach that works in several “classical” settings, such as generalizations of the Black–Scholes model with volatility skew or for one-factor interest rate models.

The use of minimum relative entropy for selecting Arrow–Debreu probabilities has also been justified on economic grounds. Samperi (1997) shows that there exists a one-to-one correspondence between the calibration of a model starting with a prior probability measure and using a “penalization function” on the space of probabilities and the calculation of state-prices via utility maximization. More precisely, the Arrow–Debreu prices coincide with the marginal utilities for consumption obtained by maximizing the expectation of the utility function  $U(x) = -\exp(-\alpha x)$  by investing in a portfolio of benchmark instruments. This correspondence is quite general. It implies, most notably, that other “distances” or “penalization functions” for Arrow–Debreu probabilities of the form

$$\tilde{D}(p|q) = \sum_{i=1}^{\nu} \psi \left( \frac{p_i}{q_i} \right) q_i, \quad \psi(x) \text{ convex} \quad (7)$$

can be used instead of relative entropy (which corresponds to the special case  $\psi(x) = x \log x$ ). For each such penalization function, there exists a corresponding concave utility  $U(x)$ , obtained via a Legendre transformation, such that the Arrow–Debreu probabilities are consistent with an agent maximizing his/her expected utility for terminal wealth by investing in a portfolio of benchmarks.<sup>e</sup>

The use of relative entropy also has consequences in terms of price sensitivity analysis and hedging. Avellaneda (1998) shows that the sensitivities of model values with respect to changes in the benchmark prices are equal to the linear regression coefficients of the payoff of contingent claims under consideration on the linear span of the cashflows of the benchmark instruments. In particular, the price sensitivities can be computed directly using a single Monte Carlo simulation, i.e., without having to disturb the  $N$  input prices and to repeat the calibration procedure each time.

<sup>e</sup>While the particular choice of the mathematical distances  $\tilde{D}(p|q)$  remains to be justified, the different distances between probabilities which result are, roughly speaking, economically equivalent — except possibly for the particular choice of smooth, increasing, convex utility that might represent the agent’s preferences. The Kullback–Leibler distance is convenient because it leads to particularly simple mathematical computations, as we shall see hereafter. Another important feature of relative entropy is that it is invariant under changes of variables and therefore independent of the choice parameterization used to describe the system (Cover and Thomas (1991)). We refer the reader to Samperi (1999) for an in-depth discussion of this correspondence principle.

Thus, we hope that this method may provide an efficient approach for computing hedge ratios as well.

Practical considerations in terms of model implementation are studied in the last four sections. We show that calibration of Monte Carlo models to the prices of benchmark instruments results in a strong reduction of variance, or simulation noise. This is due to the fact that the model effectively only needs to estimate the *residual* cashflows (modulo the linear space spanned by the benchmarks). Therefore, instruments which are well-approximated by benchmarks have very small Monte Carlo variance. In particular, the interpolation of implied volatilities and prices of option between strikes and maturities is numerically efficient.

In practice, the success of any calibration method will depend on the characteristics of the market where it is applied. To evaluate the algorithm, we consider a few concrete examples. We study option-pricing models in the foreign-exchange and equity markets, using forward and liquidly traded options as benchmarks. The models that we use incorporate stochastic volatility and are calibrated to the observed volatility skew. We also discuss the calibration of fixed-income models, and apply the algorithm to the construction of forward-rate curves based on the prices of on-the-run US Treasury securities.

## 2. Relative Entropy Distance and the Support of the Risk-Neutral Measure

Relative entropy measures the deviation of the calibrated model from the prior. Intuitively, if the relative entropy is small, the model is “close” to the prior and thus is “more desirable” than a model that has a large distance from the prior. Let us make this statement more precise in the context of Monte Carlo simulations. The relative entropy distance,

$$D(p|u) = \log \nu + \sum_{i=1}^{\nu} p_i \log p_i, \quad (8)$$

takes values in the interval  $[0, \log \nu]$ . The value zero corresponds to  $p_i = 1/\nu$  (the prior) whereas a value of  $\log \nu$  is obtained when all the probability is concentrated on a single path. More generally, consider a probability distribution which is supported on a subset of paths of size  $\mu$  and is uniformly distributed on these paths. If we take  $\mu = \nu^\alpha$ , with  $0 < \alpha < 1$ , and substitute the corresponding probabilities in (8), we find that

$$D(p|u) = \log \nu + \log \left( \frac{1}{\nu^\alpha} \right) = (1 - \alpha) \log \nu. \quad (9)$$

Within this class of measures, the relative entropy distance counts the number of paths in the support on a logarithmic scale. If  $\frac{D(p|u)}{\log \nu} \ll 1$  the support of the calibrated measure is of size  $\nu$ , whereas  $\frac{D(p|u)}{\log \nu} \approx 1$  corresponds to a measure with a “thin support”. Thin supports are inefficient from a computational viewpoint. They

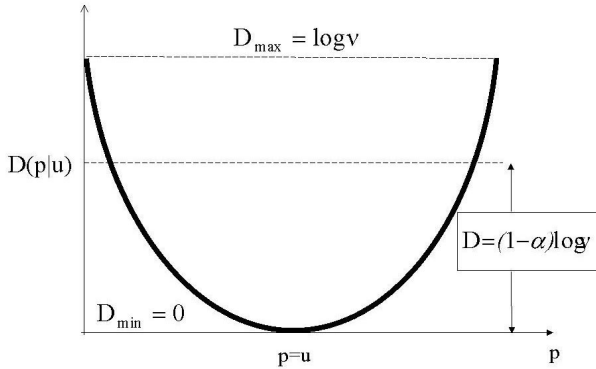


Fig. 1. Schematic graph of the relative entropy function. A probability with  $D(p|u) = (1-\alpha) \log \nu$  is supported essentially on a subset of paths of cardinality  $\nu^\alpha$ . Probabilities with small  $D(p|u)$  have large support whereas probabilities supported on a single path have the highest Kullback–Leibler distance,  $\log \nu$ .

imply that the calibration algorithm “discards” a large number of simulated paths. In this case, the *a priori* support of the distribution constructed by simulation will be very different from the the *a posteriori* support. This confirms the intuition which tells us that calibrations with small relative entropy are desirable.

This analysis can be applied to more general probability distributions. Let us write

$$p_i = \frac{1}{\nu^{\alpha_i}}, \quad i = 1, 2, \dots, \nu. \tag{10}$$

Let  $N_\alpha$  represent the number of paths with  $\alpha_i = \alpha$ , so that we have

$$\sum_\alpha N_\alpha = \nu, \quad \sum_\alpha \frac{N_\alpha}{\nu^\alpha} = 1. \tag{11}$$

Substituting (10) into (8), we find that

$$\begin{aligned} D(p|u) &= \log \nu \left( 1 + \sum_\alpha \frac{N_\alpha}{\nu^\alpha} \log \left( \frac{1}{\nu^\alpha} \right) \right) \\ &= \log \nu \left( 1 - \sum_\alpha \frac{N_\alpha}{\nu^\alpha} \alpha \right) \\ &= \log \nu (1 - \mathbf{E}^p(\alpha)), \end{aligned} \tag{12}$$

which shows that the relative entropy increases if the expected value of  $\alpha$  is large. Due to the constraints implied by (11), this is possible only if there is a wide range of exponents  $\alpha_i$ . Since probabilities are measured on a logarithmic scale, the measure will be concentrated on those paths which correspond to small values of  $\alpha$ . A wide mismatching of probabilities between the calibrated measure and the prior is undesirable because this means that certain state-contingent claims will have very different values under the prior and the posterior measures.

### 3. Calibration Algorithm

We describe the algorithm for calibrating Monte Carlo simulations under market price constraints. It is a simple adaptation of the classical dual program used for entropy optimization (see Cover and Thomas (1991)). The new idea proposed here is to apply the algorithm to the state-space which consists of a collection of sample paths generated by a Monte Carlo simulation of the prior.

To fix ideas, we shall consider a model in which paths are generated as solutions of the stochastic difference equations

$$X_{n+1} = X_n + \sigma(X_n, n) \cdot \xi_{n+1} \sqrt{\Delta T} + \mu(X_n, n) \Delta T, \quad n = 1, 2, \dots, M \quad (13)$$

where  $M\Delta T = T_{\max}$  is the horizon time. Here  $X_n \in \mathbf{R}^d$  is a vector of state variables and as a multi-dimensional process with values and  $\xi_n \in \mathbf{R}^{d'}$  is a vector of independent Gaussian shocks ( $d, d'$  are positive integers). The variance-covariance structure is represented by the  $\nu \times \nu'$  matrix  $\sigma(X, t)$  and the drift is the  $\nu$ -vector  $\mu(X, t)$ .<sup>f</sup>

Using a pseudo-random number generator, we construct a set of sample paths of (13) of size  $\nu$ , which we denote by

$$\omega^{(i)} = (X_1(\omega^{(i)}), \dots, X_M(\omega^{(i)})) \quad i = 1, 2, \dots, \nu. \quad (14)$$

We assume throughout this paper that the benchmark instruments are such that their *cashflows along each path  $\omega$  are completely determined by the path itself*. In the case of equities, where the components of the state-vector  $X$  generally represent stock prices, instruments satisfying this assumption include forwards, futures and standard European options. It is also possible to use barrier options or average-rate options. American-style derivatives do not satisfy this assumption because the early-exercise premium depends on the value of the option (and hence on the full pricing measure defined on the paths) as well as on the current state of the world. For fixed-income securities, benchmark instruments can include interest rate forwards, futures contracts, bonds, swaps, caps and European swaptions.<sup>g</sup> Under these circumstances, the price relations can be written in the form (4) where  $g_{ij}$  is the present value of the cashflows of the  $j$ th instrument along the  $i$ th path. The mathematical problem is to minimize the convex function of  $p$

$$D(p|u) = \log \nu + \sum_{i=1}^{\nu} p_i \log p_i \quad (15)$$

under linear constraints. This problem has been well studied (Cover and Thomas, 1991). Introducing Lagrange multipliers  $(\lambda_1, \dots, \lambda_N)$ , we can reformulate it as a

<sup>f</sup>This formulation extends trivially to the case of jump-diffusions or more general Markov processes and the MRE algorithm applies to these more general stochastic processes.

<sup>g</sup>American-style securities, such as Bermudan swaptions or callable bonds do not satisfy this assumption.

mini-max program (the “dual” formulation of the constrained problem)

$$\min_{\lambda} \max_p \left\{ -D(p|u) + \sum_{j=1}^N \lambda_j \left( \sum_{i=1}^{\nu} p_i g_{ij} - C_j \right) \right\} \quad (16)$$

A straightforward argument shows that probability vector that realizes the supremum for each  $\lambda$  has the Boltzmann–Gibbs form

$$p_i = p(\omega^{(i)}) = \frac{1}{Z(\lambda)} \exp \left( \sum_{j=1}^N g_{ij} \lambda_j \right). \quad (17)$$

To determine the Lagrange multipliers, define the “partition function”, or normalization factor,

$$Z(\lambda) = \frac{1}{\nu} \sum_{i=1}^{\nu} \exp \left( \sum_{j=1}^N g_{ij} \lambda_j \right). \quad (18)$$

and consider the function

$$W(\lambda) = \log(Z(\lambda)) - \sum_{j=1}^N \lambda_j C_j = \log \left\{ \frac{1}{\nu} \sum_{i=1}^{\nu} \exp \left( \sum_{j=1}^N g_{ij} \lambda_j \right) \right\} - \sum_{j=1}^N \lambda_j C_j. \quad (19)$$

We shall denote by  $g_j(\omega)$  the present value of the cashflows of the  $j$ th instrument along the path  $\omega$ . (Thus,  $g_j(\omega_i) = g_{ij}$ ). At a critical point of  $W(\lambda)$ , we have

$$\begin{aligned} 0 &= \frac{1}{Z} \frac{\partial Z}{\partial \lambda_k} - C_k \\ &= \frac{1}{Z(\lambda)} \sum_{i=1}^{\nu} g_{ik} \exp \left( \sum_{j=1}^N g_{ij} \lambda_j \right) - C_k \\ &= \mathbf{E}^P \{ g_k(\omega) \} - C_k. \end{aligned} \quad (20)$$

Hence, if  $\lambda$  is a critical point of  $W(\lambda)$ , the probability vector defined by Eq. (17) is calibrated to the benchmark instruments.

Notice that the function  $W(\lambda)$  is convex: differentiating both sides of Eq. (19) with respect to  $\lambda$  yields

$$\frac{\partial^2 W(\lambda)}{\partial \lambda_j \partial \lambda_k} = \mathbf{Cov}^P \{ g_j(\omega) g_k(\omega) \}, \quad (21)$$

which is a non-negative definite matrix. In particular, the critical point, if it exists, should correspond to a *minimum* of  $W(\lambda)$ .

Based on this, we have the following algorithm:

- (i) Construct a set of sample paths using the difference equations (1) and a pseudo-random-number generator.
- (ii) Compute the cashflow matrix  $\{g_{ij}, i = 1, \dots, \nu, j = 1, 2, \dots, N, \}$ .



- (iii) Using a gradient-based optimization routine, minimize the function  $W(\lambda)$  in (19).<sup>h</sup>
- (iv) Compute the risk-neutral probabilities  $p_i, i = 1, 2, \dots, \nu$  for each path using Eq. (27) and the optimal values of  $\lambda_1, \dots, \lambda_N$ .

#### 4. Implementation Using Weighted Least-Squares: An Alternative to Exact Fitting

It may not always be desirable to match model values to the price data exactly due to bid-ask spreads, asynchronous data, and liquidity considerations. Alternatively, we can minimize the sum of the weighted least-squares residuals and the relative entropy. We define the sum of the weighted least-squares residuals as

$$\chi_w^2 = \frac{1}{2} \sum_{j=1}^N \frac{1}{w_j} (\mathbf{E}^p \{g_j(\omega)\} - C_j)^2, \quad (22)$$

where the  $w = (w_1, \dots, w_N)$  is a vector of positive weights.

The proposal is to minimize the quantity

$$\chi_w^2 + D(p|u) \quad (23)$$

over all probability vectors  $p = (p_1, \dots, p_\nu)$ . Notice that the limit  $w_i \ll 1$  corresponds to an exact matching of the constraints. The discrepancy between the model value and market price with a weight  $w_i$  is typically of order  $\frac{1}{\sqrt{w_i}}$ .

We indicate how to modify the previous algorithm to compute the probabilities  $(p_1, \dots, p_\nu)$  that minimize  $\chi_w^2 + D(p|u)$ . Using the inequality

$$ab \leq \frac{1}{2}a^2 + \frac{1}{2}b^2, \quad (24)$$

we find that, for all  $p$ ,

$$\chi_w^2 \geq - \sum_{j=1}^N \lambda_j (\mathbf{E}^p \{g_j(\omega)\} - C_j) - \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2. \quad (25)$$

It follows that

$$\begin{aligned} \inf_p [D(p|u) + \chi_w^2] &\geq \sup_\lambda \left\{ \inf_p \left[ D(p|u) - \sum_{j=1}^N \lambda_j (\mathbf{E}^p \{g_j(\omega)\} - C_j) \right] - \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2 \right\} \\ &= - \inf_\lambda \left[ \log(Z(\lambda)) - \sum_{j=1}^N \lambda_j C_j + \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2 \right] \\ &= - \inf_\lambda \left[ W(\lambda) + \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2 \right]. \end{aligned} \quad (26)$$

<sup>h</sup>In our implementation, we use the L-BFGS algorithm.

Here,  $W(\lambda) = \log(Z(\lambda) - \sum_j \lambda_j C_j)$  is the function used in the case of exact fitting. The inequality expressed in (26) is in fact an equality. To see this, observe that the function  $D(p|u) + \chi_w^2$  is convex in  $p$  and grows quadratically for  $p \gg 1$ . Therefore, a probability vector realizing the infimum exists and is characterized by the vanishing of the first variation in  $p$ . A straightforward calculation shows that if  $p$  is a minimum of this function, we have

$$p_i^* = \frac{1}{Z(\lambda^*)} \exp\left(\sum_{j=1}^N \lambda_j^* g_{ij}\right) \quad (27)$$

with

$$\lambda_j^* = -\frac{1}{w_j} (\mathbf{E}^{p^*} \{g_j(\omega)\} - C_j). \quad (28)$$

In particular, notice that this value of  $\lambda$  is such that (25) is an equality. Furthermore, the probability (27) is of the exponential type, so we have

$$\begin{aligned} D(p^*|u) + \chi_w^2 &= D(p^*|u) - \sum_{j=1}^N \lambda_j^* (\mathbf{E}^{p^*} \{g_j(\omega)\} - C_j) - \frac{1}{2} \sum_{j=1}^N w_j (\lambda_j^*)^2 \\ &= \inf_p \left[ D(p|u) - \sum_{j=1}^N \lambda_j^* (\mathbf{E}^p \{g_j(\omega)\} - C_j) \right] - \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^{*2} \\ &= -\log(Z(\lambda^*)) + \sum_{j=1}^N \lambda_j^* C_j - \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^{*2} \\ &\leq -\inf_{\lambda} \left( W(\lambda) + \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2 \right), \end{aligned} \quad (29)$$

so equality must hold. This calculation shows that the pair  $(\lambda^*, p^*)$  is a saddlepoint of the min-max problem and that there is no “duality gap” in (26) and (29).

The algorithm for finding the probabilities that minimize  $\chi_w^2$  under the entropy penalization consists of minimizing

$$\log(Z(\lambda)) - \sum_{j=1}^N \lambda_j (\mathbf{E}^{\lambda} \{g_j(\omega)\} - C_j) + \frac{1}{2} \sum_{j=1}^N w_j \lambda_j^2 \quad (30)$$

among candidate vectors  $\lambda$ . This algorithm represents a small modification of the one corresponding to the exact fitting of prices and can be implemented in the same way, using L-BFGS.<sup>i</sup>

<sup>i</sup>More generally, we can consider the minimization of  $W(\lambda) + \sum_{j=1}^N w_j \psi(\lambda_j)$ , where  $\psi$  is a convex function. A similar argument shows that this program corresponds to minimizing the quantity  $\sum \frac{1}{w_j} \psi^*(\mathbf{E}^{\lambda} \{g_j(\omega)\} - C_j)$ , where  $\psi^*$  is the Legendre dual of  $\psi$ . The case  $\psi(x) = |x|$  can be used to model proportional bid-ask spreads in the prices of the benchmark instruments, for example.

### 5. Price Sensitivities and Hedge Ratios

The MRE setting provides a simple method for computing portfolio price sensitivities, under the additional assumption that *the prior measure remains fixed* as we disturb the benchmark prices and recalibrate.<sup>j</sup> We show, under this assumption, that sensitivities can be related to regression coefficients of the target contingent claim on the cashflows of the benchmarks. For simplicity, we discuss only the case of exact fitting, but the analysis carries over to the case of least-squares residuals with minor modifications.

Let  $F(\omega^{(i)})$ ,  $i = 1, \dots, \nu$  represent the discounted cashflows of a portfolio or contingent claim. To compute the price sensitivities of the model value of the portfolio we utilize the “chain rule”, differentiating first with respect to the Lagrange multipliers. More precisely,

$$\frac{\partial \mathbf{E}^p(F(\omega))}{\partial C_k} = \sum_{j=1}^{N_{\text{sec}}} \frac{\partial \mathbf{E}^p(F(\omega))}{\partial \lambda_j} \frac{\partial \lambda_j}{\partial C_k}. \tag{31}$$

We note, using Eq. (17) for the probability  $p_i$ , that

$$\frac{\partial \mathbf{E}^p(F(\omega))}{\partial \lambda_j} = \mathbf{Cov}^p\{F(\omega), g_j(\omega)\}. \tag{32}$$

Moreover, we have, on account of Eq. (21),

$$\frac{\partial}{\partial \lambda_j} (\mathbf{E}^p(g_k(\omega))) = \frac{\partial}{\partial \lambda_j} \left( \frac{\partial \log(Z(\lambda))}{\partial \lambda_k} \right) = \mathbf{Cov}^p\{g_j(\omega), g_k(\omega)\}. \tag{33}$$

In particular,

$$\frac{\partial C_k}{\partial \lambda_j} = \mathbf{Cov}^p\{g_j(\omega), g_k(\omega)\}. \tag{34}$$

Substitution of the expressions in (33) and (34) into Eq. (31) gives

$$\nabla_C \mathbf{E}^p\{F(\omega)\} = \mathbf{Cov}^p\{F(\omega), g(\omega)\} \cdot [\mathbf{Cov}^p\{g(\omega), g(\omega)\}]^{-1}, \tag{35}$$

with the obvious matrix notation.<sup>k</sup> This implies, in turn, that the sensitivities of the portfolio value with respect to the input prices are the linear regression coefficients of  $F(\omega)$  with respect to  $g_j(\omega)$ . Namely, if we solve

$$\min_{\beta} \left\{ \mathbf{Var}^p \left[ F(\omega) - \beta_0 - \sum_{j=1}^N \beta_j g_j(\omega) \right] \right\}, \tag{36}$$

<sup>j</sup>This assumes, implicitly, that the prior probability represents information that “varies slowly” with respect to the observed market prices. For example, the assumption is consistent with interpreting the prior as a historical probability.

<sup>k</sup>The invertibility of the covariance matrix presupposes that the cashflow vectors of the benchmark instruments,  $g_j(\omega)$   $j = 1, \dots, N$ , are linearly independent. This assumption is discussed, for example, in Avellaneda (1998).

we obtain, from (35),

$$\beta_k = \frac{\partial \mathbf{E}^P(F(\omega))}{\partial C_k} \quad k = 1, \dots, N \quad (37)$$

and<sup>1</sup>

$$\beta_0 = \mathbf{E}^P(F(\omega)) - \sum_{j=1}^N \beta_j \mathbf{E}^P(g_j(\omega)). \quad (38)$$

We conclude that a sensitivity analysis with respect to variations of the input prices can be done without the need to perform additional Monte Carlo runs or disturb the input prices one by one. Instead, the MRE framework allows us to compute prices and hedges with a single Monte Carlo simulation, which is much less costly.<sup>m</sup>

Notice that the characterization of the hedge ratios as regression coefficients shows that they are “stable” in the sense that they vary continuously with input prices. In practice, the significance of this hedging technique depends on details of the implementation procedure, such as the number of paths used in the simulation. The main issue is whether the support of the probability measure induced by the prior — the basic scenarios of the simulation — is sufficiently “rich in scenarios”, for example.

## 6. Variance Reduction

The calibration of Monte Carlo simulations can significantly reduce pricing errors. Claims that are “well-replicated” by the benchmarks — in the sense that the variance in (36) is small compared to the variance of  $F(\omega)$ <sup>n</sup> — will benefit from a significant noise reduction compared to standard MC evaluation.

In fact, given any vector  $\zeta = (\zeta_1, \dots, \zeta_N)$ , we have

$$\mathbf{E}^P(F(\omega)) = \mathbf{E}^P \left\{ F(\omega) - \sum_{j=1}^N \zeta_j g_j(\omega) \right\} + \sum_{j=1}^N \zeta_j C_j. \quad (39)$$

Since the second term on the right-hand is constant, the variance of the Monte Carlo method for pricing the cashflow  $F$  is the same as the one associated with  $F - \zeta \cdot g$ . This statement is true for any value of the vector  $\zeta$  and so, in particular, for the regression coefficients  $(\beta_1, \dots, \beta_N)$ . Since, by definition,  $F - \beta \cdot g$  has the

<sup>1</sup>This result can be interpreted as follows. Assume that an agent hedges the initial portfolio by shorting  $\beta_j$  units of the  $j$ th benchmark instrument for  $j = 1, \dots, N$ . In this case, the model value of the net holdings (initial portfolio + hedge) is  $\beta_0$ . It represents the expected cost of the dynamic replication of the residual.

<sup>m</sup>In contrast, a perturbation analysis that uses centered differences to approximate the partial derivatives with respect to input instruments requires  $2N + 1$  Monte Carlo simulations.

<sup>n</sup>The ratio of the variances is the statistic  $1 - R^2$  in the risk-neutral measure.

least possible true variance among all choices of  $\zeta$ , the cashflow  $\beta \cdot \Gamma$  is an “optimal control variante” for the simulation. Our method implicitly uses such control variates.

To measure this variance reduction experimentally in a simple framework, we considered the problem of calibrating a Monte Carlo simulation to the prices of European stock options, assuming a lognormal price with constant volatility.

We considered European options on a stock with a spot price of 100 with no dividends. The interest rate was taken to be zero. Taking a “maximum horizon” for the model of 120 days, we used as benchmarks all European options with maturities of 30, 60 and 90 days and strikes of 90, 100 and 110, as well as forward contracts with maturities of 30, 60 and 90 days. We assumed that the prices of the benchmarks were given by the Black–Scholes formula with a volatility of 25%. The prior was taken to be a geometric Brownian motion with drift zero and volatility 25%.<sup>o</sup>

The test consisted of pricing various options (target options) with strike/maturity distributed along a regular grid (maturities from 20 days to 120 days at 1-day intervals; all integer strikes lying between two standard deviations from the mean of the distribution). For each option, we compared the variances resulting from pricing with the simulated lognormal process with and without calibrating to the “benchmarks”. As a matter of general principle, when pricing an option contract, we also include the forward contract corresponding to the option’s expiration date in the set of calibration instruments.<sup>p</sup>

All Monte Carlo runs were made with 2000 paths. Each run took roughly half a second of CPU time on a SunOS 5.6. This includes the time required to search for the optimal lambdas. We verified the correctness of the scheme by checking that all model prices fell within three (theoretical) standard deviations of the true price, both with and without the min-entropy adjustment.

We found that there was significant variance reduction in all cases, with the exception of options having strikes far from the money and maturities which did not match the benchmark maturities. As expected, the best results were observed for those options with strikes and maturities near to those of the benchmark options. In particular, options with the “benchmark maturities” (30, 60 and 90 days) yielded some of the best results for most strikes which were not too far away from the money. We also obtained some of the best results from options with strikes at or close to the forward values. The following table gives the factor by which the entropy method improved the variance for selected strikes and maturity dates. Note that the table below includes the benchmark instruments which yield an infinite improvement since the entropy method always prices them correctly (indicated by INF on Fig. 2). Benchmark strikes and maturities are shown in boldface.

<sup>o</sup>We assumed that all benchmark options were correctly priced with the prior to separate the issues of calibration and variance reduction.

<sup>p</sup>Doing so guarantees that the mean of the distribution of the asset price at the expiration date is fitted exactly.

Maturity (Days)	Strike								
	80	85	90	95	100	105	110	115	120
20	N/A	N/A	1.03	2.22	8.24	2.78	1.54	N/A	N/A
<b>30</b>	N/A	N/A	INF	13.66	INF	19.11	INF	3.09	N/A
45	N/A	1.25	2.38	5.21	14.83	6.63	3.75	2.25	N/A
<b>60</b>	N/A	5.63	INF	49.83	INF	73.98	INF	10.83	3.03
75	1.54	3.04	6.25	10.61	25.71	14.08	9.12	5.45	3.05
<b>90</b>	2.47	8.79	INF	92.40	INF	150.36	INF	22.96	5.77
120	1.96	2.77	4.09	6.08	13.57	8.34	5.77	4.15	3.03

Fig. 2. Variance improvement ratio.

Maturity (Days)	Strike								
	80	85	90	95	100	105	110	115	120
20	N/A	N/A	0.56	0.37	0.32	0.37	0.51	N/A	N/A
<b>30</b>	N/A	N/A	0	0.16	0	0.15	0	0.38	N/A
45	N/A	0.47	0.35	0.28	0.26	0.28	0.34	0.47	N/A
<b>60</b>	N/A	0.25	0	0.10	0	0.09	0	0.21	0.45
75	0.50	0.32	0.23	0.21	0.21	0.22	0.24	0.30	0.42
<b>90</b>	0.38	0.19	0	0.07	0	0.07	0	0.15	0.32
120	0.40	0.33	0.30	0.29	0.29	0.29	0.31	0.35	0.40

Fig. 3. Standard errors from the entropy method (in percentage of implied volatility).

Maturity (Days)	Strike								
	80	85	90	95	100	105	110	115	120
20	N/A	N/A	0.25	0.57	0.85	0.65	0.44	N/A	N/A
<b>30</b>	N/A	N/A	1	0.93	1	0.95	1	0.75	N/A
45	N/A	0.40	0.63	0.80	0.91	0.83	0.72	0.55	N/A
<b>60</b>	N/A	0.83	1	0.97	1	0.98	1	0.91	0.68
75	0.41	0.68	0.83	0.89	0.95	0.92	0.88	0.81	0.67
<b>90</b>	0.62	0.89	1	0.99	1	0.99	1	0.95	0.80
120	0.51	0.63	0.73	0.79	0.90	0.85	0.80	0.73	0.64

Fig. 4.  $R^2$  statistic from the entropy method.

The variance reduction from the entropy method translated into some excellent data for the computed standard errors. The figure below contains this information. The data is given in terms of Black–Scholes implied volatility and is obtained by taking the standard error of the option price and dividing it by the Black–Scholes value of vega.

Finally, we examined the  $R^2$  statistic given by the entropy method. We found that  $R^2$  was greatest for values with benchmark maturity dates and strikes whose values were close to that of the forward price. The results are given in Fig. 4. Our interpretation is that the options with benchmark dates or near-the-money strikes have only a small component of their cashflows which is orthogonal to the benchmark instruments. One would expect both greater variance reduction and dependence on the values of the benchmark instruments in these cases. The variance reduction data given in Fig. 2 confirms this interpretation.

## 7. Example: Fitting a Volatility Skew

We apply the algorithm to calibrate a model using forwards and the prices of European options with different strikes and maturities. This example is taken from the interbank foreign exchange market. It is well-known that options with different strikes/maturities trade with different implied volatilities. The goal is to construct a pricing model that incorporates this effect. Notice that this problem has been addressed by many authors in the context of the so-called “volatility surface” (Dupire (1994), Derman and Kani (1994), Rubinstein (1994), Chriss (1996); see Avellaneda *et al.* (1997) for references up to 1997 on this problem). The method presented here is completely different since we do not interpolate option prices or use a parameterization of the local volatility function  $\sigma(S, t)$ .

We considered a dataset consisting of 25 contemporaneous USD/DEM option prices obtained from a major dealer in the interbank market on August 25, 1995. The maturities are 30, 60, 90, 180 and 270 days. Strikes (quoted in DEM) correspond to 50-, 20- and 25-delta puts and calls. Aside from these options, we introduced five additional “zero-strike options” which correspond to the present value of a dollar in DEM for delivery at the different expiration dates (see Fig. 5) — the forward prices implied by the interest rates and the spot price. Including these forward prices in the set of benchmark instruments ensures that the model is calibrated to the forward rates and hence there is no net bias in the forward prices.

As a prior, we considered the system of stochastic differential equations:

$$\begin{aligned} \frac{dS_t}{S_t} &= \sigma_t dZ_t + \mu dt \\ \frac{d\sigma_t}{\sigma_t} &= \kappa dW_t + \nu_t dt, \end{aligned} \tag{40}$$

where  $Z_t$  and  $W_t$  are Brownian motions such that  $\mathbf{E}(dZ_t dW_t) = \rho dt$ . In Eq. (40),  $S_t$  represents the value of one US Dollar in DM. The instantaneous volatility is denoted by  $\sigma_t$ . The additional parameters are:  $\mu$ , the cost-of-carry (interest rate differential),  $\kappa$ , the volatility of volatility, and  $\nu_t$  is the drift of the volatility. Therefore, we are calibrating a two-factor stochastic volatility model. We assume the following numerical values for the parameters that define the prior dynamics:

- (1)  $S_0 =$  mid-market USD/DEM spot exchange rate = 1.4887

exp (days)	type strike	price	exp (days)	type strike	price
30	call 1.5421	0.007	180	call 1.6025	0.0141
30	call 1.531	0.0093	180	call 1.5779	0.0191
30	call 1.4872	0.0234	180	call 1.4823	0.0505
30	put 1.4479	0.0092	180	put 1.3902	0.0216
30	put 1.4371	0.0069	180	put 1.3682	0.0162
60	call 1.5621	0.0094	270	call 1.6297	0.0173
60	call 1.5469	0.0126	270	call 1.5988	0.0226
60	call 1.4866	0.0319	270	call 1.4793	0.0598
60	put 1.4312	0.0128	270	put 1.371	0.0254
60	put 1.4178	0.01	270	put 1.3455	0.019
90	call 1.5764	0.0112	30	fwd 0	1.486695
90	call 1.558	0.0149	60	fwd 0	1.484692
90	call 1.4856	0.0378	90	fwd 0	1.482692
90	put 1.4197	0.0153	180	fwd 0	1.476708
90	put 1.4038	0.0114	270	fwd 0	1.470749

Fig. 5. Data used for fitting the implied volatilities of options. The implied volatilities, displayed on the left-hand side of the graph, range from 13% to 14.5%. From Avellaneda and Paras (1996).

- (2) US rate = 5.91%
- (3) DM rate = 4.27%
- (4)  $\mu = -1.64\%$  (for convenience, we take  $\mu = \text{DM rate} - \text{US rate}$  in the prior, i.e., we adjust the model to the standard risk-neutral drift).
- (5)  $\sigma_0 = \text{Initial value of the prior volatility of USD/DEM} = 14\%$ . (This is essentially the average of the observed implied volatilities.)
- (6)  $\kappa = 50\%$
- (7)  $\rho = -50\%$

We simulated 5000 paths of Eq. (40), consisting of 2500 paths and their antithetics. The gradient tolerance in the BFGS routine was set to  $10^{-7}$  and we used equal weights  $w_i = 10^{-5}$  for the least-squares approximation. We found that the difference between model prices and market prices was typically between the orders of  $10^{-4}$  and  $10^{-5}$  DM, representing relative errors of 1% in the deep-out-of-the-money short-term options and much less 0.1% for at-the-money options (see Fig. 6).

The algorithm initiated with  $\lambda_i = 0, i = 1, \dots, 30$  converges after approximately 20 iterations of the BFGS routine. The entire calibration procedure takes about four seconds on a desktop PC with a Pentium II 330 Mhz processor. In practice, computation times are much faster because the values of lambdas from the previous runs can be stored and used as better initial guesses.

The relative entropy of the calibrated risk-neutral measure was found to be  $D(p|u) = 7.39 \times 10^{-3}$ . We can interpret this result in terms of the parameter  $\alpha$  of



exp (days)	type strike	error	lambda	exp (days)	type strike	error	lambda		
30	c	1.5421	-0.000019	-22.31451	180	c	1.6025	-0.000015	-21.81038
30	c	1.531	0.000045	43.71395	180	c	1.5779	0.000032	24.07218
30	c	1.4872	-0.000042	-27.9719	180	c	1.4823	-0.000006	-10.61853
30	p	1.4479	0.000004	-5.30083	180	p	1.3902	0.000018	19.90037
30	p	1.4371	0.000016	12.82587	180	p	1.3682	-0.000016	-15.64954
60	c	1.5621	-0.000025	-26.09551	270	c	1.6297	-0.000001	6.951896
60	c	1.5469	0.000033	34.28321	270	c	1.5988	-0.000007	-3.66328
60	c	1.4866	-0.000003	-7.354312	270	c	1.4793	0.000009	-3.421416
60	p	1.4312	-0.000029	-31.52781	270	p	1.371	0.000019	10.23301
60	p	1.4178	0.000041	32.31679	270	p	1.3455	-0.000012	-13.96789
90	c	1.5764	-0.000019	-18.9968	30	f	0	0.00002	8.864074
90	c	1.558	0.000034	29.95544	60	f	0	0.000018	-1.028268
90	c	1.4856	0.000004	-18.24796	90	f	0	0.000027	7.661008
90	p	1.4197	0.000029	38.27566	180	f	0	0.000017	5.521524
90	p	1.4038	-0.00005	-36.20459	270	f	0	0.000013	0.71636

Fig. 6. Fitting errors and lambdas for the 30 instruments using 5000 paths. The relative entropy is  $D \approx 0.07$ .

$p$	
Mean	0.0002
Standard Error	1.12E-06
Median	0.000188
mode	0.000162
Standard Deviation	7.95E-05
Sample Variance	6.32E-09
Kurtosis	1.046163
Skewness	0.896427
Range	0.00059
Minimum	0.000039
Maximum	0.000629

Fig. 7. Descriptive statistics for the vector of calibrated probabilities  $(p_1, \dots, p_{5000})$ .

Sec. 2. We find a value of  $\alpha = 1 - \frac{D(p|u)}{\log \nu} = 0.99913$ , which would correspond to an “effective number of paths”  $\nu^\alpha \approx 4963$  according to the heuristics of Sec. 2. This represents an excellent fit in terms of the support of the calibrated measure. In Fig. 7, we present descriptive statistics for the calibrated probabilities, in Fig. 8, we plot the probabilities, which appear to be randomly distributed with a small mean




Fig. 8. Snapshot of the 5000 probabilities obtained by the method. The values oscillate about  $1/5000 = 0.0002$ .




Fig. 9. Histogram of the logarithms of the calibrated probabilities multiplied by 5000. The distribution is unimodal and strongly peaked about  $\log(1/5000) = -8.517$ . Observe that there are some outliers corresponding to small probabilities.

about the uniform value  $1/5000 = 0.0002$ . Figure 9 displays a histogram of the logarithms of the probabilities. These results indicate a relatively small scattering about the mean and a large set of paths which support the posterior measure.

The values of the lambdas were between  $-0.84$  and  $0.79$ , which correspond to moderate variations of the probabilities about their mean. In Fig. 9, we present a

histogram of the 5000 probabilities obtained. The distribution of probabilities — or, equivalently, the distribution of state-price deflators — is unimodal and strongly peaked about its mean  $1/5000 = 0.0002$ . This confirms that the risk-neutral measure is supported on the full set of paths generated using the prior.

To gain insight into the calibrated model, we generated an *implied volatility surface*, by repricing a set of options on a fine grid in strike/maturity space. We determined the highest and lowest strikes in the input option, which are, respectively, 1.67 DEM and 1.34 DEM, corresponding to the 20-delta options with the longest maturity (270 days). We then considered strike increments going from the maximum to the minimum value according to the rule

$$k_{\max} = 1.67 = k_0, \quad k_i = \frac{k_{i-1}}{1.01}, \quad i = 1, 2, \dots, 20 \tag{41}$$

and maturities

$$t_{\min} = 30 = t_0, \quad t_i = t_{i-1} + 10, \quad i = 1, 2, \dots, 24. \tag{42}$$

We repriced these 480 options with the weighted Monte Carlo and generated a surface by interpolating the implied volatilities linearly between strikes and dates. The interpolation was done graphically using the Excel 5.0 graphics package. The results are exhibited in Figs. 10 and 11.

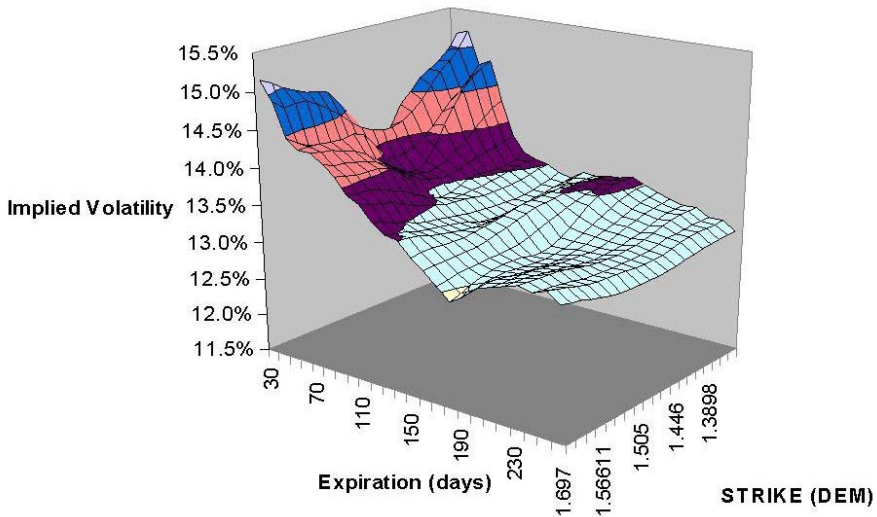


Fig. 10. Implied volatility surface for the USD/DEM options market based on the data in Fig. 5. The prior volatility is 14%.

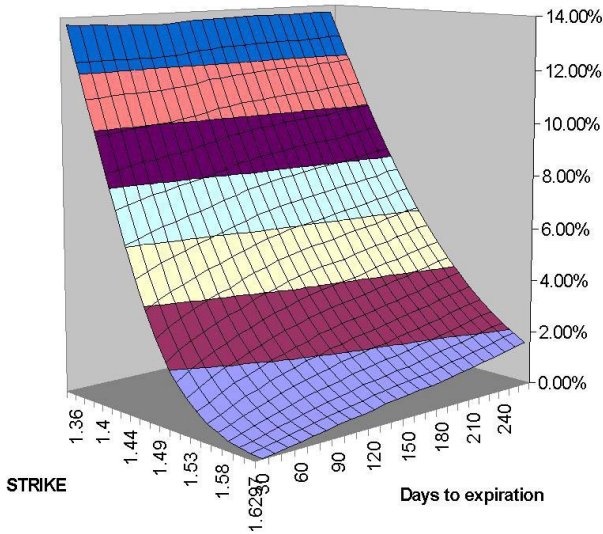


Fig. 11. Call prices associated with the implied volatility surface of Fig. 10.

Finally we tested the model by pricing a barrier option with the following characteristics:

- (1) Option type: USD put, DM call
- (2) Notional amount: 1 USD
- (3) Maturity: 180 days
- (4) Strike: 1.48 DM
- (5) Knockout barrier: 1.38 DM

and comparing the results with the price and delta hedge obtained using the Black–Scholes lognormal model.

Closed-form solutions for the price of barrier options in the lognormal setting were computed by Reiner and Rubinstein (1991). To compare with our model, we used a Black–Scholes constant volatility of 14% — which corresponds to the implied volatilities of short-maturity options — as well as with a constant volatility of 13%. The latter corresponds to the implied volatility of the ATM forward call expiring in 180 days. The differences in the Black–Scholes prices and deltas by changing the volatilities from 13% to 14% are quite small. In Fig. 12, the results of applying the closed-form solution with 14% volatility are compared to the Monte Carlo simulation model described above. The values obtained with both models agree to three significant digits.

To evaluate the quality of the hedges, we computed the “net delta” of the exotic option given by the model and compared it with the Black–Scholes delta. For this, we assumed that the input options have the same deltas under BS than under our model. Of course, this is a simplification, since the deltas of the individual options may

exp.	type	strike	hedge	exp.	type	strike	hedge
30	c	1.5421	-3.95%	180	c	1.6025	-7.35%
30	c	1.531	4.00%	180	c *	1.5779	14.08%
30	c	1.4872	-1.08%	180	c *	1.4823	13.49%
30	p	1.4479	4.19%	180	p *	1.3902	-111.23%
30	p	1.4371	-6.71%	180	p *	1.3682	91.36%
60	c	1.5621	4.41%	270	c	1.6297	-2.91%
60	c	1.5469	-5.89%	270	c	1.5988	4.47%
60	c	1.4866	2.30%	270	c	1.4793	-1.41%
60	p *	1.4312	-17.21%	270	p	1.371	-2.47%
60	p *	1.4178	17.66%	270	p	1.3455	2.41%
90	c	1.5764	-3.50%	30	f		0.48%
90	c	1.558	1.61%	60	f		-0.14%
90	c	1.4856	-4.53%	90	f		5.22%
90	p *	1.4197	-15.59%	180	f		-20.59%
90	p *	1.4038	21.96%	270	f		-0.07%

expiration	options	fwd	net delta
30	0.0%	0.5%	0.4%
60	1.3%	-0.1%	1.2%
90	-3.1%	5.2%	2.2%
180	18.3%	-20.6%	-2.3%
<b>ALL</b>	<b>16.6%</b>	<b>-15.0%</b>	<b>1.5%</b>

<b>MC Value</b>	= 0.358%
<b>BS Value</b>	= 0.364%
<b>BS Delta</b>	= 1%

Fig. 12. Price and hedge report for the reverse-knockout barrier option. Notice that the calibrated model computes exposures to all the options and forwards entered as reference instruments. Instruments that have an exposure (in notional terms) of more than 10% of the notional amount of the exotic option are labelled with asterisks. We observe significant exposure (1) at 60 and 90 days near the knockout barrier and (2) at the expiration date in ATM and low-strike options. Hedges at the barrier involve a spread in contracts with neighboring strikes, as expected.

be affected globally by the differences in implied volatilities. However, experience shows that such approximation is reasonable, in the sense that we do not expect the volatility skew in this market to distort significantly the deltas of the plain-vanilla options. The total amount of forward dollars needed to hedge the knockout is obtained by converting each option hedge into an equivalent forward position and summing over all contracts with the same maturity. To this amount, we also add the corresponding sensitivity to the forward contract (modeled here as a call with strike 0).

Adding the “forward deltas” associated with the different maturities, we find a total of 1.57% which is near the Black–Scholes value of 1.80. The results obtained in this example appear reasonable, despite the fact that the computation was done using Monte Carlo simulation with only 5000 paths. We believe that this is due in part to the reduction of variance which results from the calibration process.

### 8. Example 2: Fitting the Smile for America Online Options on May 1999

We calibrated another stochastic volatility model to the mid-market prices of 35 America Online call prices recorded on May 10, 1999 at the close of the trading day.

To ensure sufficient liquidity of the benchmarks, we took only options with a traded volume above 100 for shorter maturities and above 50 for those with maturity periods longer than half a year. It can be seen from Fig. 13 that the implied volatilities of the benchmark calls vary in the range 78.33–88.52%. These extreme volatilities correspond to deeply in- or out-of-the-money short-term options. We also included forward prices for the stock at the different delivery dates, namely 12, 40, 68, 159 and 257 days. Thus, we calibrated the simulation to 40 benchmark prices.

We simulated  $\nu = 10000$  Monte Carlo paths from this distribution on the time-horizon of 258 days with one time-step per day. We used the following parameters:  $S_0 = 128.375$  (the America Online closing price on May 10, 1999),  $\sigma_0 = 0.86$ ,  $\rho = -0.5$ ,  $r = 5\%$  and  $\kappa = 0.5$ . Then, we applied the MRE method described in Sec. 3.2 to calibrate the uniform distribution on the obtained sample paths to the set of 35 European call prices on America Online given in Fig. 13.

The program matched all the given 40 benchmark prices with the predefined accuracy to four decimal places using 181 iterations starting from  $\lambda = 0$ . The obtained entropy of the calibrated measure on the path space was 0.66 with the maximum possible entropy being  $\log 10000 = 9.21$ , indicating that the prior distribution is not far in the entropy distance from the calibrated one.

Maturity	Strike	Price	IVOL	Maturity	Strike	Price	IVOL
12	120	12.125	78.33	68	150	12.25	87.99
12	130	7.125	83.73	68	170	7.625	87.60
12	135	5.125	83.29	68	175	6.625	86.80
12	140	3.625	83.40	68	180	6	87.54
12	145	2.625	85.16	68	200	3.75	87.84
40	115	21.75	85.34	159	120	32.625	83.90
40	120	18.875	85.27	159	125	30.25	83.19
40	125	16.25	84.96	159	150	21.5	83.43
40	135	12.25	86.86	159	160	19	84.19
40	140	10.625	87.80	257	100	48.375	81.21
40	160	5.5	87.65	257	110	43.375	80.80
68	100	35.625	88.52	257	120	38.75	80.07
68	110	29	86.98	257	130	35.125	80.76
68	115	25.875	85.55	257	150	28.125	79.73
68	120	23.25	85.61	257	160	25.375	79.77
68	125	20.875	85.78	257	170	23	79.98
68	135	17.125	87.80	257	200	16.625	78.93
68	145	13.625	87.54				

Fig. 13. America Online call (mid-market) prices on May 10, 1999.

Figure 14 displays the implied volatility surface associated with the calibrated model. This surface was obtained by pricing a dense grid of plain-vanilla options with the calibrated Monte Carlo. Figure 15 represents the surface of corresponding call option prices. Notice that the shapes of the implied volatility surfaces in




Fig. 14. Implied volatility surface for AOL option closing prices, calibrated to the prices of Fig. 13. The additional parameters are spot price=128.312,  $\sigma_0 = 86\%$ ,  $r = 5\%$ ,  $\kappa = 50\%$  and  $\rho = -50\%$ . The relative entropy distance to the prior is  $D = 0.66$ .




Fig. 15. Surface of call prices corresponding to the implied volatility surface of Fig. 14.

Maturity	Coupon	Price
07/15/99	—	98.955
10/14/99	—	97.823
03/30/00	—	95.725
03/31/01	4.875	99.875
02/15/04	4.75	98.812
11/15/08	4.75	97.219
02/15/29	5.25	96.250

Fig. 16. Seven benchmark US-Treasury bills and bonds. Prices are as quoted on 04/15/99 and not adjusted for accrued interest. The alignment of prices reflects the direction of the sensitivities listed below in Fig. 20: left alignment indicates positive sensitivity, right alignment indicates negative sensitivity.

both examples are quite different. Of course, the call price surfaces appear to be more similar: from well-known no-arbitrage relations, they are both convex and decreasing in the strike direction and monote-increasing with expiration.

### 9. Example 3: Constructing a US Treasury Yield Curve

We considered the following calibration problem: given the current prices of on-the-run treasury securities, construct a smooth forward rate curve consistent with these prices.

Figure 16 shows the on-the-run instruments and the prices observed in the morning of Thursday, April 15, 1999.

A stochastic short-rate model was used to discount future cashflows. As a prior, we considered the modified Vasicek model

$$dr = \alpha(m(t) - r)dt + \sigma dW. \quad (43)$$

Here,  $m(t)$  is the possibly time-dependent level of mean reversion, and the constant  $\alpha$  controls the rate of mean reversion. We experimented with two types of mean reversion levels: constant levels and time-dependent levels, where the latter were taken to be  $m(t) =$  the piecewise-constant (bootstrapped) instantaneous forward-rate curve.<sup>9</sup>

Figure 17 shows several prior instantiations of the coefficients of (43).

We calibrated the modified Vasicek model with 15000 Monte Carlo paths and 24 time steps per year. Figure 18 shows the calibrated forward-rate curve and zero-coupon-bond yield curve for scenarios I and II, with constant level of mean

<sup>9</sup>These are arbitrary modeling choices. For example, we could start with an econometrically calibrated forward rate curve or with a level of mean reversion that corresponds to an estimate of forward rates for long maturities. Bootstrapping is standard method for building a forward rate curve: it works by assembling forward rates instrument by instrument, earlier maturities first. Rates are constant between maturities and jump at maturities.



Scenario	$\alpha$	$m(t)$	$\sigma$
I	0.25	0.0426035	<b>0.01</b>
II	0.25	0.0426035	<b>0.05</b>
III	0.25	bootstrap	<b>0.01</b>
IV	0.25	bootstrap	<b>0.05</b>

Fig. 17. Various prior instantiations of the coefficients of (43). 0.0426035 is the rate of the first leg of the piecewise constant bootstrapped forward rate curve.

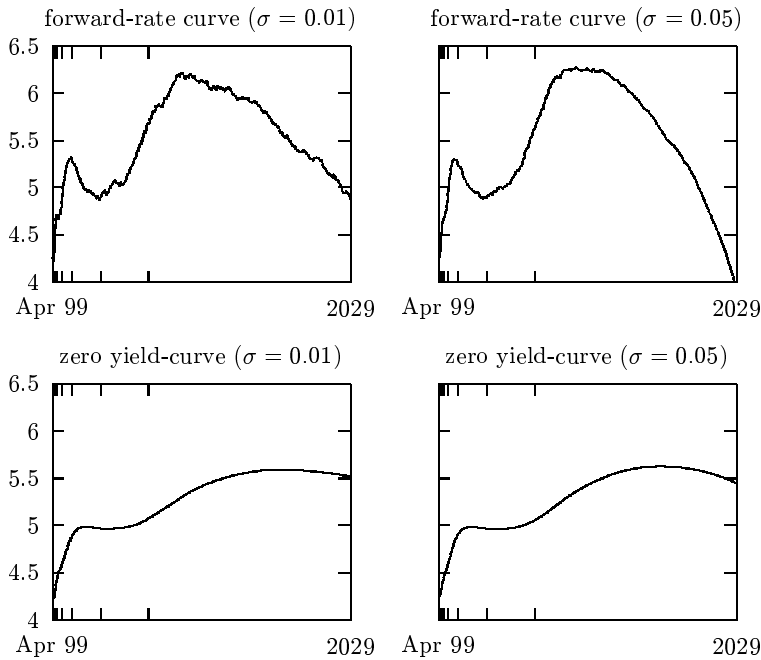


Fig. 18. Forward-rate curves and zero-coupon-bond yield curves for scenario I ( $\sigma = 0.01$ ) on the left side, and scenario II ( $\sigma = 0.05$ ) on the right side.

reversion. Figure 19 shows the calibrated forward-rate curve and zero-coupon-bond yield curve for scenarios III and IV. In these scenarios, the level of mean reversion  $m(t)$  is set to the piecewise constant bootstrapped forward rate curve consistent with the data in Fig. 16.

In accordance with the work of Samperi (1997) and others, the optimal Lagrange multipliers  $\lambda_j^*$  can be interpreted in terms of an optimal investment portfolio. Specifically, consider an expected CARA utility function defined on the space of static portfolios  $(\theta_1, \dots, \theta_N)$  as follows:

$$U(\theta) = -\frac{1}{\nu} \sum_{i=1}^{\nu} e^{-\sum_{j=1}^N \theta_j (g_{ij} - C_j)}. \tag{44}$$

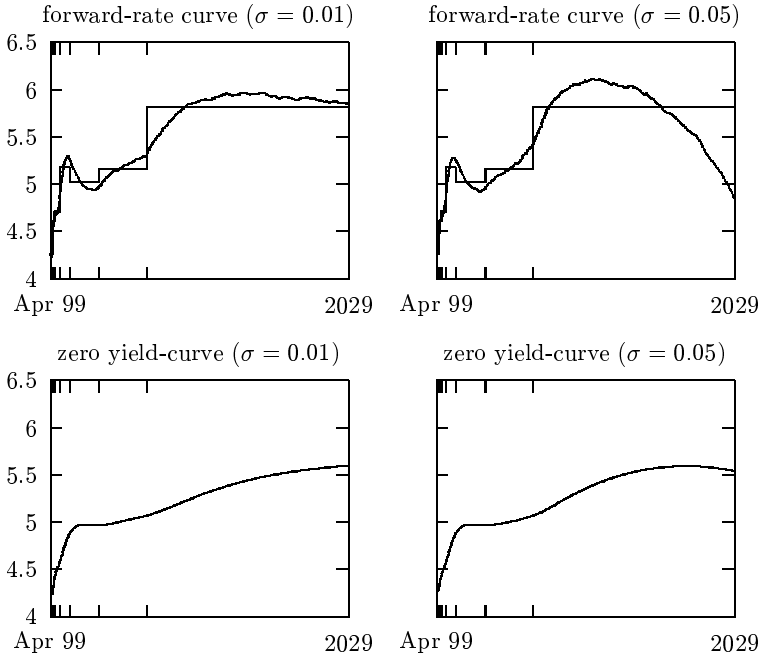


Fig. 19. Forward-rate curves and zero-coupon-bond yield curves for scenario III ( $\sigma = 0.01$ ) on the left side, and scenario IV ( $\sigma = 0.05$ ) on the right side. Both scenarios revert to the piecewise constant bootstrapped prior  $m(t)$  superimposed in the top row.

Since  $\lambda^*$  minimizes  $\log(Z(\lambda)) - \lambda \cdot C$ , it follows from the analysis of Sec. 3 that the vector of Lagrange multipliers and the optimal portfolio are in simple correspondence: we have

$$\lambda_j^* = -\theta_j, \quad j = 1, \dots, N. \tag{45}$$

Sensitivities are the opposites of the optimal portfolio weights for (44). A negative lambda corresponds to a “cheap” instrument (hence  $\theta > 0$ ) and a positive lambda to a “rich” instrument (hence  $\theta < 0$ ). The Lagrange multipliers, or sensitivities,  $\lambda_1^*, \dots, \lambda_7^*$  for all four scenarios are summarized in Fig. 20. Making  $m(t)$  time-dependent does not change the sensitivities significantly (scenario I versus III and II versus IV). The Lagrange multipliers corresponding to short-term instruments, however, are very high if the prior volatility is low ( $\sigma = 0.01$  in scenarios I and III).

This last application of the MRE algorithm has also been implemented by one of the authors (R. Buff) as a prototype software operating remotely via the Internet. The software, which uses periodically updated Treasury securities prices and/or bond prices entered by the users, is accessible on the Courant Finance Server (<http://www.courantfinance.cims.nyu.edu>).

Maturity	Sensitivities for scenario...			
	I	II	III	IV
07/15/99	<b>32.507</b>	0.582	<b>29.421</b>	0.569
10/14/99	<b>-17.638</b>	-0.242	<b>-15.307</b>	-0.298
03/30/00	<b>5.359</b>	0.114	<b>4.607</b>	0.123
03/31/01	<b>-1.667</b>	-0.056	<b>-1.692</b>	-0.055
02/15/04	0.085	0.009	0.211	0.010
11/15/08	0.340	0.013	0.013	0.005
02/15/29	-0.356	-0.021	-0.050	-0.011

Fig. 20. The sensitivities for the seven benchmark instruments in Fig. 16, in each of the four prior scenarios. Sensitivities with an absolute value greater than 1 are typeset in boldface.

## 10. Conclusions

We have presented a simple approach for calibrating Monte Carlo simulations to the price of benchmark instruments. This approach is based on minimizing the Kullback–Leibler relative entropy between the posterior measure and a prior measure. In this context, the prior corresponds to the uniform measure over simulated paths (hence to the “classical” Monte Carlo simulation). This approach is known to be equivalent to finding the Arrow–Debreu prices which are consistent with an investor which maximizes an expected utility of exponential type. The advantage of the minimum-entropy algorithm is that (i) it is non-parametric (and thus not market- or model-specific) and (ii) it allows the modeller to incorporate econometric information and *a priori* information on the market dynamics, effectively separating the specification of the dynamics from the issue of price-fitting.

We showed that the algorithm can be implemented as an exact fit to prices or in the sense of least-squares. The notion of entropy distance can be interpreted as a measure of the logarithm of the *effective number of paths* which are active in the posterior measure. Large entropy distances correspond therefore to “thin” supports and thus to an orthogonality (in the measure-theoretic sense) between the prior and posterior measures.

The sensitivities produced by the model can be computed via regression, without the need to simulate the market dynamics and to recalibrate each time that we disturb the price of a benchmark instrument. Another interesting feature of the weighted Monte Carlo algorithm is the reduction of variance which results from the exact pricing of benchmark instruments. In fact, the simulation effectively evaluates the “residual risk” obtained after projecting the payoff of interest onto the space of portfolios spanned by the benchmark instruments. Numerical experiments indicate that the reduction of variance can be significant.

We discussed concrete implementations of the algorithm for the case of foreign-exchange and equity options, calibrated the underlying dynamics to

two-factor stochastic volatility models. We have exhibited numerical evidence that shows that such an algorithm can be implemented in practice on desktop computers.

## Acknowledgments

We are grateful to Graciela Chichilnisky, Freddy Delbaen, Raphael Douady, Darrell Duffie, Nicole El Karoui, David Faucon, Olivier Floris, Helyette Geman, Jonathan Goodman, Robert Kohn, Jean-Pierre Laurent, Jean-Michel Lasry, Jerome Lebuchoux, Marek Musiela, Jens Nonnemacher and Frank Zhang for their enlightening comments and suggestions. We acknowledge the hospitality and generous support of Banque Paribas, Ecole Polytechnique, Ecole Normale Supérieure de Cachan, ETH-Zurich and the Center for Financial Studies (Frankfurt). This work was partially supported by the US National Science Foundation (DMS-9973226).

## References

- [1] M. Avellaneda, “Minimum-entropy calibration of asset-pricing models”, *Int. J. Theor. App. Finance* **1**(4) (1998) 447.
- [2] M. Avellaneda, C. Friedman, R. Holmes and D. Samperi, “Calibrating volatility surfaces via relative-entropy minimization”, *App. Math. Finance* **4**(1) (1997) 37–64.
- [3] M. Avellaneda and A. Paras, “Managing the volatility risk of portfolios of derivative securities: The Lagrangian uncertain volatility model”, *App. Math. Finance* **3** (1996) 21–52.
- [4] P. W. Buchen and M. Kelly, “The maximum entropy distribution of an asset inferred from option prices”, *J. Financial and Quantitative Analysis* **31**(1) (1996) 143–159.
- [5] N. Chriss, “Transatlantic trees”, *RISK* **9** (1996) 7.
- [6] E. Derman and I. Kani, “Riding on a smile”, *RISK* **7** (1994) 2.
- [7] B. Dupire, “Pricing with a smile”, *RISK* **7** (1994) 1.
- [8] B. Dupire, *Monte Carlo Methodologies and Applications for Pricing and Risk Management*, RISK Publications, London (1998).
- [9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, New York (1991).
- [10] L. Gulko, “The Entropy Theory of Option Pricing”, Yale University Working Paper (1995).
- [11] “The Entropy Theory of Bond Pricing”, Working Paper, Yale University (1996).
- [12] J. C. Jackwerth and M. Rubinstein, “Recovering Probability Distributions from Contemporaneous Security Prices”, Working Paper, Berkeley University, Haas School of Business (1995).
- [13] J. P. Laurent and D. Leisen, “Building a Consistent Pricing Model from Observed Option Prices”, Working Paper, the Hoover Institute, Stanford University (1999).
- [14] D. Samperi, “Inverse Problems, Model Selection and Entropy in Derivative Security Pricing”, PhD thesis, New York University (1997).
- [15] M. Rubinstein, “Implied binomial trees”, *J. Finance* **69**(3) (1994) 771–818.
- [16] M. Rubinstein, and E. Reiner, “Breaking down the barriers”, *RISK* **4**(8) (1991) 28–35.

- [17] C. Zhu, R. H. Boyd, P. Lu and J. Nocedal, *L-BFGS-B: FORTRAN Subroutines for Large-Scale Bound-Constrained Optimization*, Northwestern University, Department of Electrical Engineering (1994).
- [18] Y. Zhu and M. Avellaneda, "A risk-neutral stochastic volatility model", *Int. J. Theor. App. Finance* **1**(2) (1998) 289.