

Reading the smile: the message conveyed by methods which infer risk neutral densities

Eric Jondeau ^a, Michael Rockinger ^{b,*}

^a *Banque de France, Centre de recherche, 31 rue Croix des Petits Champs, 75049 Paris, France*

^b *HEC-School of Management, Department of Finance, 78351 Jouy-en-Josas, France*

Abstract

In this study we compare the quality and information content of risk neutral densities obtained by various methods. We consider a non-parametric method based on a mixture of log-normal densities, the semi-parametric ones based on an Hermite approximation or based on an Edgeworth expansion, the parametric approach of Malz which assumes a jump-diffusion for the underlying process, and Heston's approach assuming a stochastic volatility model. We apply those models on FF/DM exchange rate options for two dates. Models differ when important news hits the market (here anticipated elections). The non-parametric model provides a good fit to options prices but is unable to provide as much information about market participants expectations than the jump-diffusion model. © 2000 Elsevier Science Ltd. All rights reserved.

JEL classification: C52; G14; G15; F31; F33

Keywords: Risk-neutral density; Exchange rate options; Option pricing

1. Introduction

Much of the literature following the seminal work on option pricing by Black and Scholes (1973) and Merton (1973) assumed that the asset underlying an option follows a log-normal diffusion process. Empirical studies of option volatility, such as Rubinstein's (1994) presidential address, have shown that exchange rate options, out of or in the money, are associated with a different level of volatility than at the

* Corresponding author. Tel.: +33-139-67-72-59; fax: +33-139-67-70-85.

E-mail addresses: ejondeau@banque-france.fr (E. Jondeau), rockinger@gwsmt.hec.fr (M. Rockinger).

money options, a feature called the options smile. This finding is in contradiction with the assumption of a log-normal distribution for the underlying asset and shows that to correctly price options more general models are required.

Various methods have been suggested to extract out of options' prices the underlying risk neutral density (RND). This density is related to market participants' expectations of the future price process in a risk-neutral environment. As shown by Bahra (1996) and Campa et al. (1997), once such a density is obtained it is possible to compute moments as well as confidence intervals. As such, the RND plays an important role as a tool to evaluate the credibility of the central bank. RNDs are also important for an investor, for instance in risk management, who needs to quantify in terms of probability how a market may evolve in the future. RNDs can also be used to price exotic options.

The contribution of this study is the comparison of the advantages and drawbacks of various methods which extract risk neutral densities applied to FF/DM European type exchange rate options. We were able to obtain a time series of observations of OTC options covering 20 dates ranging between May 1996 and June 1997. For each day we dispose of a set of maturities up to one year. First, we discuss the implementation of the various methods in a cross sectional framework by focusing on just two dates: 17 May 1996, a day when the exchange rate markets were known to be calm, and on 25 April 1997, a few days after the French President Chirac announced dissolution of the National Assembly, which implied nation-wide elections. Second, we run all methods in a time series context which allows us to further retain a satisfying model for the exchange rate data at hand. The discussion of the message contained in a time series of confidence intervals obtained from RNDs illustrates the usefulness of this type of research. During the period under investigation we have another noticeable event in the summer of 1996 where we find a significant depreciation of the FF/DM due to the uncertainty about the ability of the French government to satisfy the Maastricht criteria (especially the deficit criteria).

We first provide a description of a large number of methods that allow construction of a RND. A first method based on approximating the RND with a mixture of densities, which could be called *non-structural*, is advocated by Bahra (1996) and Campa et al. (1997). Melick and Thomas (1997) indicate in addition how to price American options. In a study by Söderlind and Svensson (1997) it is shown how this mixture of densities method can be applied to various financial assets asking what can be learnt from the point of view of a policy-maker.

We also consider an approach based on the work of Jarrow and Rudd (1982) who developed a method for option pricing under the assumption that the underlying asset is not log-normally distributed. They show how the RND can be obtained as an Edgeworth expansion around a log-normal density. We consider this approach to be of *semi-nonparametric* nature. Their approach has been implemented by Corrado and Su (1996, 1997) who show that with this method options can be better priced.

In a similar spirit Madan and Milne (1994) describe the underlying RND with an Hermite polynomial approximation. Abken et al. (1996) provide an application and show how higher moments of the underlying asset are perceived to vary through time.

Bates (1996a,b) and Malz (1996a,b) go one step further and consider a *structural*

model by assuming that the underlying process follows a jump-diffusion, and in particular a Bernoulli version. Thus, they assume a full specification for the underlying price process. The RND obtained in their model depends on some parameters which can be estimated from options prices. Their work aims at extracting information concerning market participants expectations out of options prices.

Further structural models in which the price process of the underlying asset is fully specified are models of stochastic volatility. Hull and White (1987), Chesney and Scott (1989), Melino and Turnbull (1990), and Ball and Roma (1994), Chesney and Scott (1989), Melino and Turnbull (1990) and Ball and Roma (1994), assume that volatility follows a diffusion process. To make their models tractable they have to make simplifying assumptions concerning the correlation between volatility and the underlying asset's return. Heston (1993) by assuming a different process for volatility and by using a different numerical approach provides an almost closed form solution for option prices for a more general stochastic volatility environment.

The aim of most of those studies is to provide a pricing tool. Breeden and Litzenberger (1978) observe that the second derivatives of an options' price with respect to the strike price yields the RND. This observation makes it possible to derive from any option pricing model the underlying RND. In a similar vein, Gesser and Poncet (1997) derive an interesting term structure of volatility and compare the actual term structure with the ones generated by Hull and White and by Heston.

Several other approaches to obtain a RND have been proposed. Aït-Sahalia and Lo (1998) provide a non-parametric method based on time-series analysis and kernel estimates. Stutzer (1996) suggests a multistep procedure where the initial step also involves historical prices of the underlying asset. Rubinstein (1994) and Jackwerth and Rubinstein (1995) develop a method based on binomial trees. We restrict ourselves to models which do not involve trees and in which no history of the underlying asset is required.

Unlike some of the literature which has addressed the question how to price options under non-constant volatility (e.g. Derman and Kani, 1994; Dumas et al., 1998; Dupire, 1994; Shimko, 1993 as well as Stein and Stein, 1991) we address the question of the information content in options of various maturities.

In Section 2 we review various non-structural, semi-nonparametric, and structural methods. In Section 3 we introduce the data. Section 4 contains a cross-sectional comparison of the methods with a discussion of the parameters obtained for our structural models and a comparison of higher moments and confidence intervals. In Section 5 we turn to the time-series comparison. Section 6 concludes. Estimation issues are relegated to Appendix A.

2. Recovering RNDs

The following section outlines notation and the general paradigm within which we evaluate RNDs. Several of the methods described below could be adapted to instances in which the underlying asset is not an exchange rate. Such instances include Black's (1976) model for options on futures.

Let S_t be the price at t of a unit of foreign currency in local money.¹ A European call option written on S_t gives its owner the right to buy the underlying asset for the exercise (also called strike) price K at the expiration (or maturity) date τ . Since a rational investor will only exercise his right if he realizes a profit, the payoff for a call is $\max(S_\tau - K, 0)$.

A European put option written on S_t gives the owner the right to sell the underlying asset for the exercise price K at the expiration date τ . Exercise before τ is not possible. The payoff for a put is $\max(K - S_\tau, 0)$.

Under the assumption that the market is arbitrage free, Harrison and Kreps (1979) show that there exists a probability density for the underlying price process such that the call and put option price can be written as

$$C_t = e^{-rT} \int_{S_\tau=K}^{+\infty} (S_\tau - K) a(S_\tau, \tau; S_t, t | \theta) dS_\tau, \tag{1}$$

and

$$P_t = e^{-rT} \int_0^{S_\tau=K} (K - S_\tau) a(S_\tau, \tau; S_t, t | \theta) dS_\tau, \tag{2}$$

where θ is a vector of parameters describing the RND $a(\cdot)$, and where we defined the time to expiration as $T = \tau - t$.²

2.1. The benchmark case of log-normality: Garman–Kohlhagen

2.1.1. The model

Much of the early research on options has assumed a given price process for S_t , for instance that S_t follows a log-normal diffusion such as:

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \tag{3}$$

where μ , and σ represent respectively the instantaneous mean and volatility and where W_t is a Brownian motion with respect to some probability measure P .

Under such assumptions for the underlying asset, it can be shown that in a risk-neutral world the process S_t can be written as:

$$dS_t = (r - r^*) S_t dt + \sigma S_t dW_t^*, \tag{4}$$

where W_t^* is again a Brownian motion with respect to Q , an equivalent martingale measure and where r and r^* represent the domestic and foreign continuously compounded risk free interest rates. Under log-normality the RND associated with the

¹ For instance for the DM/FF options, S_t will represent the number of FF necessary to acquire one unit of DM

² A textbook level derivation can be found in Duffie (1988, p. 115).

future exchange rate can be obtained by the fact that $\ln(S_\tau)$ follows a normal with mean $\ln(S_t) + (r - r^* - \sigma^2/2)T$ and variance $\sigma^2 T$. This result follows from Ito's lemma.³ Thus, the RND is

$$a(S_\tau) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma S_\tau} \exp \left\{ -\frac{1}{2} \left(\frac{\ln(S_\tau) - \ln(S_t) - (r - r^* - \sigma^2/2)T}{\sigma\sqrt{T}} \right)^2 \right\}.$$

For this situation call and put options can be evaluated as truncated expectations. Garman and Kohlhagen (1983), following the methodology outlined by Black and Scholes (1973) and Merton (1973) obtain:

$$C(S_t, T, K, \sigma, r, r^*) = e^{-r^*T} S_t \Phi(d_1) - e^{-rT} K \Phi(d_2), \tag{5}$$

$$P(S_t, T, K, \sigma, r, r^*) = -e^{-r^*T} S_t [1 - \Phi(d_1)] - e^{-rT} K [1 - \Phi(d_2)], \tag{6}$$

$$d_1 = \frac{\ln(S_t/K) + (r - r^* + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \tag{7}$$

$$d_2 = \frac{\ln(S_t/K) + (r - r^* - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}. \tag{8}$$

As a consequence of non-arbitrage, under the risk-neutral probability the discounted expectation of the future price must be equal to the current price. This translates into the following *martingale restriction*:

$$S_t = e^{-(r-r^*)T} \int_0^{+\infty} S_\tau a(S_\tau) dS_\tau.$$

³ From (Eq. (4)) we obtain $d \ln(S_t) = (r - r^* - [1/2]\sigma^2)dt + \sigma dW_t^*$ and hence $\ln(S_\tau) = \ln(S_t) + (r - r^* - [1/2]\sigma^2)T + \sigma(W_\tau^* - W_t^*)$. Since $W_\tau^* - W_t^*$ is distributed as a normal random variable with mean 0 and variance T we can conclude. We recall that if $\ln(S) \sim \mathcal{N}(\mu, \sigma^2)$ then the density of S is $\phi((\ln(S) - \mu)/\sigma)/(\sigma S)$ and its distribution function is $\Phi((\ln(S) - \mu)/\sigma)$. In this work ϕ and Φ represent always the density and the cumulative density of the normal distribution.

2.1.2. The link between deltas and strike prices

OTC options’ quotation is not done in terms of prices for a set of exercise prices but in terms of volatilities for options of various deltas. Given volatility, the spot exchange rate, the various interest rates, and time to maturity, there exists as we indicate below, a one-to-one relation between deltas and the strike price.

The delta of an option is defined as the derivative of the price with respect to the underlying value. Hence, for a call, and respectively for a put, we have

$$\delta^C = \frac{\partial C}{\partial S_t}(S_t, T, K, \sigma, r, r^*) = e^{-r^*T} \Phi(d_1),$$

$$\delta^P = \frac{\partial P}{\partial S_t}(S_t, T, K, \sigma, r, r^*) = e^{-r^*T} \Phi(-d_1),$$

where d_1 is defined in Eq. (7). Since δ is a strictly decreasing function of K , for each δ there corresponds a unique strike price which can be extracted numerically.

Since European calls and puts are related through the put-call parity, if we have the K for a call then $1 - \delta$ corresponds to a put with the same volatility and the same K . In other words, rather than working with calls and puts we focus only on calls. In practice only in the money call and put options are quoted. The non-existence of call and put options with a same strike implies that we cannot back out further information such as an implied spot exchange rate.

Once the strike price K is obtained it is possible to invert the pricing Eqs. (5)–(8) for each option and to obtain for each one a price in FF.⁴

2.2. A non-structural approach

Focusing on Eq. (1), we obtain by applying Leibniz’ rule, as in Breeden and Litzenberger (1978), that

$$\frac{\partial^2 C_t}{\partial K^2} = e^{-rT} a(K, \tau; S_t, t | \theta). \tag{9}$$

Thus, a simple computation of second derivatives gives us the actualized RND. This suggests a first method to extract a RND where the only (yet key) assumption to be made is that there exist enough strike prices to approximate numerically the density and where we need the assumption of arbitrage-free markets.⁵

However, numerical derivatives are known to be numerically unstable, and a more fruitful strategy is to assume that the RND, $a(\cdot)$, takes certain particular expressions.

⁴ If prices were quoted in numeraire, then, as the underlying asset changes, it would be necessary to continuously update the options price. Further, if options were quoted for a given set of exercise prices, as the spot rate moves it would be necessary to introduce new strike prices.

⁵ It should be mentioned that $a(\cdot)$ is the undiscounted RND on which we focus in this study whereas $e^{-rT}a(\cdot)$ represents an Arrow–Debreu state price. In the literature this state price gets sometimes referred to as the RND.

Sherrick et al. (1996) assume for a the Burr III distribution and Abadir and Rockinger (1998) fit densities derived from Kummer functions. In this work we do not pursue this road but follow Bahra (1996), Melick and Thomas (1997) and Söderlind and Svensson (1997) who describe a as a mixture of log-normal distributions.

Let $l(S_\tau; \mu_i, \sigma_i)$ ($L(S_\tau; \mu_i, \sigma_i)$) denote the log-normal density (and its associated cdf) with parameters μ_i and σ_i then

$$C_i = e^{-rT} \sum_{i=1}^M \alpha_i \int_{S_\tau=K}^{+\infty} (S_\tau - K) l(S_\tau; \mu_i, \sigma_i) dS_\tau$$

will describe the option price as a mixture of M log-normal distributions. The α_i are positive and sum up to 1. This formula can be evaluated easily since the formula for truncated expectations of log-normals,⁶

$$\int_{S_\tau}^{+\infty} (S_\tau - K) l(S_\tau, \theta) dS_\tau = (E[S_\tau | S_\tau > K] - K) Pr[S_\tau | S_\tau > K],$$

gives us a formula, equivalent, from the point of view numerical complexity, to the Garman–Kohlhagen formula,

$$C_i = e^{-rT} \sum_{i=1}^M \alpha_i \exp\left(\mu_i + \frac{1}{2}\sigma_i^2 T\right) \left(\left[1 - \Phi\left(\frac{\ln(K) - \mu_i - \sigma_i^2 T}{\sigma_i \sqrt{T}}\right) \right] - K \left[1 - \Phi\left(\frac{\ln(K) - \mu_i}{\sigma_i \sqrt{T}}\right) \right] \right)$$

In addition, the martingale constraint can be imposed with

$$S_0 \exp((r - r^*)T) = \sum_{i=1}^M \alpha_i \exp\left(\mu_i + \frac{1}{2}\sigma_i^2 T\right).$$

2.3. A semi-nonparametric approach involving Edgeworth expansions

In the following section we outline the method developed by Jarrow and Rudd (1982) for which a numerical application can be found in Corrado and Su (1996).⁷

⁶ Johnson et al. (1994), p. 241. indicate that if $S \sim N(\mu, \sigma^2)$ then $E[S | S > K] = \exp\left(\mu + \frac{1}{2}\sigma^2\right) \frac{1 - \Phi(U_0 - \sigma)}{1 - \Phi(U_0)}$, where $U_0 = \frac{\ln(K) - \mu}{\sigma}$.

⁷ Below we adapt their work to the pricing of European foreign exchange options.

The idea of Jarrow and Rudd (1982) is to capture deviations from log-normality by an Edgeworth expansion of the RND $a(S_\tau, \tau; S_t, t | \theta)$ in Eq. (9) around the log-normal density.⁸ The use of an Edgeworth expansion in this context is conceptually similar to Taylor expansions but applies to functions. In a conventional Taylor expansion, some function is approximated at a given point by a simpler polynomial. Here, the RND is approximated by an expansion around a lognormal density. A further difference is that expansions are usually made to obtain simplifications whereas here the approximation, by involving parameters which can be varied, allows us to generate more complicated functions.

In the next section we will present an alternative approach given by Madan and Milne (1994). There it is assumed that the RND can be obtained as a multiplicative perturbation of some given density. This multiplicative error allows for a certain control of higher moments. As shown further on, both methods can yield numerically similar results, conceptually, however, they are different.

First we will sketch how Edgeworth expansions can be obtained. Let A be the cumulative distribution function of a random variable X and its density. Define the characteristic function of X as $\xi(A, t) \equiv \int e^{itx} a(x) dx$. If moments of X exist up to order n then there exist cumulants $\kappa_j(A)$ implicitly defined by the expansion

$$\ln(\xi(A, t)) = \sum_{j=1}^{n-1} \kappa_j(A) \frac{(it)^j}{j!} + o(t^n - 1).$$

If a characteristic function is known, by taking an expansion of its logarithm around $t=0$, it is possible to obtain the cumulants. Between cumulants and moments up to the fourth order we have $\kappa_1(A) = E[X]$, $\kappa_2(A) = \text{Var}[X]$, $\kappa_3(A) = E[(X - E[X])^3]$, $\kappa_4(A) = E[(X - E[X])^4] - 3\text{Var}[X]$. Jarrow and Rudd show that an Edgeworth expansion of the fourth order for the true probability distribution A around the log-normal distribution L can be written, after imposing the condition that the first moment of the approximating density and the true probability are equal, ($\kappa_1(L) = \kappa_1(A)$):

$$a(s) = l(s) + \frac{\kappa_2(A) - \kappa_2(L)}{2!} \frac{d^2 l(s)}{ds^2} + \frac{(\kappa_3(A) - \kappa_3(L))}{3!} \frac{d^3 l(s)}{ds^3} + \frac{(\kappa_4(A) - \kappa_4(L)) + 3(\kappa_2(A) - \kappa_2(L))^2}{4!} \frac{d^4 l(s)}{ds^4} + \varepsilon(s),$$

where $\varepsilon(s)$ captures terms neglected in the expansion. The various terms in the expansion correspond to adjustments of the variance, skewness, and kurtosis. This expression is similar to a Taylor expansion, yet it is not the same since the coefficients of the terms in $d^j l/ds^j$ are parameters and not raised to a power.

⁸ Edgeworth expansions are frequently used in statistical theory to obtain distributions which deviate from the normal one.

Jarrow and Rudd show further that the price of an European call option struck at K can be written as

$$C(A) = C(L) + e^{-rT} \frac{\kappa_2(A) - \kappa_2(L)}{2!} l(K) - e^{-rT} \frac{\kappa_3(A) - \kappa_3(L)}{3!} \frac{dl(K)}{dS_\tau} \tag{10}$$

$$+ e^{-rT} \frac{(\kappa_4(A) - \kappa_4(L)) + 3(\kappa_2(A) - \kappa_2(L))^2}{4!} \frac{d^2(K)}{dS_\tau^2} + \varepsilon(K).$$

Since L stands for the log-normal distribution, it follows that $C(L)$ corresponds to the Garman–Kohlhagen formula and higher order cumulants can be obtained as functions of elementary components:

$$\kappa_1(L) = S_t \exp((r - r^*)T), \quad \kappa_2(L) = [\kappa_1(L)q]^2, \quad \kappa_3(L) = [\kappa_1(L)q]^3(3q + q^3),$$

$$\kappa_4(L) = [\kappa_1(L)q]^4(16q^2 + 15q^4 + 6q^6 + q^8),$$

where $q = (e^{\sigma^2 T} - 1)^{1/2}$ and where the first relation follows from risk neutral valuation.

Jarrow and Rudd suggest identification of the second moment by imposing $\kappa_2(L) = \kappa_2(A)$. This argument is also justified on numerical grounds by Corrado and Su (1996) who notice that without this condition there will exist a problem of multicollinearity between the second and the fourth moment. Corrado and Su (1996) rather than estimating the remaining cumulants, ($\kappa_3(A)$ and $\kappa_4(A)$), estimate skewness and kurtosis (written respectively $\gamma_1(A)$ and $\gamma_2(A)$) through

$$\gamma_1(A) = \frac{\kappa_3(A)}{[\kappa_2(A)]^{3/2}}, \quad \gamma_2(A) = \frac{\kappa_4(A)}{[\kappa_2(A)]^2}.$$

Clearly, similar expressions hold for the distribution L . Given the assumption of equality of the second cumulants in the approximating and the true distribution it follows that

$$C(A) = C(L) - e^{-rT} (\gamma_1(A) - \gamma_1(L)) \frac{\kappa_2^{3/2}(L)}{3!} \frac{dl(K)}{dS_\tau} + \tag{11}$$

$$e^{-rT} (\gamma_2(A) - \gamma_2(L)) \frac{\kappa_2^2(L)}{4!} \frac{d^2l(K)}{dS_\tau^2}.$$

Using this expression it is easy to estimate with NLLS the implied volatility, (σ^2), skewness, ($\gamma_1(A)$), and kurtosis, ($\gamma_2(A)$).

The expression of the RND can be obtained after twice differentiating Eq. (11) with respect to K and then evaluating at S_τ :

$$a(S_\tau) = l(S_\tau) - (\gamma_1(A) - \gamma_1(L)) \frac{\kappa_2^{3/2}(L)}{6} \frac{\partial^3 l(S_\tau)}{\partial S_\tau^3} + (\gamma_2(A) - \gamma_2(L)) \frac{\kappa_2^2(L)}{24} \frac{\partial^4 l(S_\tau)}{\partial S_\tau^4},$$

where the partial derivatives can be computed iteratively using

$$\frac{\partial l}{\partial S_\tau} = - \left(1 + \frac{\ln(S_\tau) - m}{\sigma^2 T} \right) \frac{l(S_\tau)}{S_\tau},$$

$$\frac{\partial^2 l}{\partial S_\tau^2} = - \left(2 + \frac{\ln(S_\tau) - m}{\sigma^2 T} \right) \frac{1}{S_\tau} \frac{\partial l(S_\tau)}{\partial S_\tau} - \frac{1}{S_\tau^2 \sigma^2} l(S_\tau),$$

$$\frac{\partial^3 l}{\partial S_\tau^3} = - \left(3 + \frac{\ln(S_\tau) - m}{\sigma^2 T} \right) \frac{1}{S_\tau} \frac{\partial^2 l(S_\tau)}{\partial S_\tau^2} - \frac{2}{S_\tau^2 \sigma^2} \frac{\partial l(S_\tau)}{\partial S_\tau} + \frac{1}{S_\tau^3 \sigma^3} l(S_\tau),$$

$$\frac{\partial^4 l}{\partial S_\tau^4} = - \left(4 + \frac{\ln(S_\tau) - m}{\sigma^2 T} \right) \frac{1}{S_\tau} \frac{\partial^3 l(S_\tau)}{\partial S_\tau^3} - \frac{3}{S_\tau^2 \sigma^2} \frac{\partial^2 l(S_\tau)}{\partial S_\tau^2} + \frac{3}{S_\tau^3 \sigma^3} \frac{\partial l(S_\tau)}{\partial S_\tau} - \frac{1}{S_\tau^4 \sigma^4} l(S_\tau),$$

and where $m = \ln(S_t) + (r - r^* - \sigma^2/2)T$. These computations indicate that the RND in the Edgeworth case will be a polynomial whose coefficients directly command the skewness and kurtosis of the RND. We also notice that the RND involves rather complicated terms involving derivatives of the log-normal density.

2.4. A semi-parametric approach involving hermite polynomials

The theoretical foundations of this method are elaborated in Madan and Milne (1994) and applied in Abken et al. (1996). Other recent research using Hermite approximations within an option pricing context is Knight and Satchell (1997).

Their model operates as follows. First, they assume that the underlying asset follows a lognormal diffusion

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \tag{12}$$

where W_t is a Brownian motion with respect to some abstract reference density $\phi(\cdot)$ assumed to be Normal with mean zero and variance 1. This implies, when we move to a discretization, that

$$S_\tau = S_t \exp\left(\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}z\right) \tag{13}$$

where $z \sim \mathcal{N}(0,1)$.

The key idea of this approach is that the RND can be obtained through a multiplicative perturbation, (λ) , to the normal density so that $a(z) = \lambda(z)\phi(z)$. This can be alternatively viewed as a change in probability. Rather than assuming specific expressions for λ to go from one probability to another as one does under the martingale approach for option valuation, they assume a parametric structure for λ . The main thrust of their work aims at estimating $\lambda(z)$.

The key observation of their approach is that the reference measure being a normal one, the various components involved in the option pricing can be expressed as linear combinations of Hermite polynomials. Let $\{h_k\}_{k=1}^\infty$ be those polynomials. Such poly-

nomials are known to form an orthogonal basis with respect to the scalar product

$$\langle f, g \rangle = \int f(z)g(z)\phi(z)dz.^9$$

Since under the reference measure, $\phi(z)$, the dynamics of the underlying asset are perfectly defined, Madan and Milne show how it is possible to write any payoff, such as for instance the payoff of a call option as:

$$(z - K)^+ = \sum_{k=0}^{\infty} a_k h_k(z).$$

The a_k are well defined and their expression depends on μ, σ, T, τ .

On the other hand, it is also possible to write $\lambda(z)$ with respect to the basis as $\lambda(z) = \sum_{j=0}^{\infty} b_j h_j(z)$. Following Eq. (1) and given the orthogonality property of Hermite polynomials, the price of a call option can then be written as

$$C = \sum_{k=0}^{\infty} a_k \pi_k,$$

where the $\pi_k = e^{-rT} b_k$ are interpreted as the implicit price of polynomial risk h_k . Since the Hermite polynomial of order k will depend on a k th moment we will also refer to π_3 and π_4 as the price of skewness and kurtosis.

For practical purposes, the infinite sum can be truncated up to the fourth order. One can then either estimate $\pi_k, k=1, \dots, 4$ or follow Abken et al. (1996) and impose $\pi_0 = e^{rT}, \pi_1 = \pi_2 = 0$ and estimate μ, σ , and π_3, π_4 . In this case the RND simplifies to

$$\tilde{a}(z) = \phi(z) \left[1 + \frac{b_3}{\sqrt{6}} (z^3 - 3z) + \frac{b_4}{\sqrt{24}} (z^4 - 6z^2 + 3) \right], \tag{14}$$

where the b_i are parameters to be estimated. The parameters b_3 and b_4 correspond to the skewness and kurtosis if z follows a normal distribution. It is important to emphasize that unlike the Edgeworth case, since a further change of variable from z to S_τ has to be made, b_3 and b_4 will not correspond in general to the skewness

⁹ The Hermite polynomial of order k is defined by $H_k(x) = (-1)^k \frac{\partial^k \phi}{\partial x^k} \frac{1}{\phi(x)}$ where ϕ is the mean zero and unit variance normal density. After standardization of the polynomials H_k to unit norm, one obtains that the first four standardized Hermite polynomials are $h_0(x) = 1, h_1(x) = x, h_2(x) = (x^2 - 1)/\sqrt{2}, h_3(x) = (x^3 - 3x)/\sqrt{6}, h_4(x) = (x^4 - 6x^2 + 3)/\sqrt{24}$.

and kurtosis of the future exchange rate S_τ . It is also worth mentioning that the expression given by Eq. (14) is sometimes called a Gram–Charlier expansions which is the basis for other recent research (as in Knight and Satchell, 1997).

In the empirical part of this work we further pin down μ by imposing the martingale restriction and estimate only σ and the future value of the third and fourth price of risk. The actual risk neutral density $a(S_\tau)$ can then be inferred using the change of variable $z=[\ln(S_\tau)-\ln(S_t)-(r-r^*-\sigma^2/2)T]/\sigma\sqrt{T}$. Careful comparison of this RND with the one obtained in the previous section shows that, even though both involve a polynomial of the fourth degree, those polynomials are not equal even though they may yield similar shapes in numerical applications.

2.5. Risk neutral density for a process with jumps

In this section we assume that S_t is a log–normal jump–diffusion and hence the sum of a geometric Brownian motion and a Poisson jump process. The importance of jumps is emphasized by Jorion (1989) and Taylor (1994). Pricing formula for the jump–diffusion can be found in Merton (1976), Cox and Ross (1976) and Bates (1991, 1996a,b). Within this framework Malz (1996a,b) shows how information can be recovered from options when only very little information is available.

Under the assumption that the price process is the sum of a geometric Brownian motion and a jump component we can write that

$$dS_t = \mu S_t dt + \sigma S_t dW_t + k S_t dq_t$$

where q is a Poisson counter with average rate of jump occurrence λ and jump size k . In a very general set-up k could be assumed to be a random variable.

The price process under the risk neutral probability can be shown to be

$$dS_t = (r - r^* - \lambda E[k]) S_t dt + \sigma S_t dW_t^* + k S_t dq_t$$

Ball and Torous (1983, 1985) and Malz (1996a,b) assume for simplicity that over the horizon of the option there will be at most one jump of constant size. In this case, referred to as the Bernoulli version of the jump diffusion, the call and put prices become respectively:

$$(1 - \lambda T)C(S_t, T, K, \sigma, r, r^* + \lambda k) + (\lambda T)C(S_t(1 + k), T, K, \sigma, r, r^* + \lambda k),$$

$$(1 - \lambda T)P(S_t, T, K, \sigma, r, r^* + \lambda k) + (\lambda T)P(S_t(1 + k), T, K, \sigma, r, r^* + \lambda k).$$

In these formulae, $1 - \lambda T$ represents the probability of no jump before maturity. Bates and Malz point to the difficulty of disentangling λ and k numerically. For this reason we will only interpret the expected jump size λk .

We also would like to mention, at this stage, that we will estimate this structural model for various dates and maturities. This will yield for each date and maturity a set of estimates. This may appear to contradict the assumption of constant parameters in the underlying process, on the other hand this issue is the same as with quoting

options in terms of volatilities. We will follow the literature and interpret the estimates as being those perceived to be valid at some point of time by market participants until the expiration of the option. It should be further noticed that the time series of parameters so obtained may correspond to a process of the underlying asset which has little to do with historically observed processes.

2.6. Risk neutral density for a model with stochastic volatility

An alternative to assuming jumps is to assume, as in Heston’s (1993) model, that volatility is stochastic. In the following, we recall the formulas used in Heston’s model.

The price dynamics are assumed to be given by

$$dS_t = \mu S_t dt + \sqrt{v_t} S_t dW_{1,t}, \quad dv_t = \kappa(\theta - v_t) dt + \gamma \sqrt{v_t} dW_{2,t}.$$

The parameters of Heston’s model are: θ the long-run volatility, κ the mean-reversion speed, and γ the volatility of the volatility diffusion. v_t is the instantaneous volatility. A priori v_t is not a parameter to estimate but the realization of a random variable. However, since it is unobservable, it is fairly natural to estimate it with the true parameters. Lastly, ρ denotes the correlation between the two Brownian motions $W_{1,t}$, and $W_{2,t}$.

Heston shows that the call option price is

$$C = e^{-r^*T} S_t P_1 - e^{-rT} K P_2, \quad P_j = \frac{1}{2} + \frac{1}{\pi} \int_0^{+\infty} \text{Re} \left(\frac{\exp[-ix \ln(K)] f_j(x)}{ix} \right) dx \quad \forall j=1,2,$$

where the integrand can be constructed with¹⁰

$$u_1 = 1/2, \quad u_2 = -1/2, \quad a = \kappa\theta, \quad b_1 = \kappa + \lambda - \rho\gamma, \quad b_2 = \kappa + \lambda,$$

$$d_j = [(\rho\gamma ix - b_j)^2 - \gamma^2(2u_j ix - x^2)]^{1/2},$$

$$g_j = \frac{b_j - \rho\gamma ix + d_j}{b_j - \rho\gamma ix - d_j},$$

$$D_j = \frac{b_j - \rho\gamma ix + d_j}{\gamma^2} \frac{1 - \exp(d_j T)}{1 - g_j \exp(d_j T)},$$

$$C_j = (r - r^*) ix T + \frac{a}{\gamma^2} \left[(b_j - \rho\gamma ix + d_j) T - 2 \ln \left(\frac{1 - g_j \exp(d_j T)}{1 - g_j} \right) \right],$$

$$f_j = \exp(C_j + D_j v_t + ix \ln(S_t)),$$

where λ stands for the price of the volatility risk. The parameters to be estimated

¹⁰ i is the complex number, solution to $i^2 = -1$.

are $a, b_1, b_2, \rho, \gamma, v_t$. Because λ is not identifiable we introduce $\kappa^* = \kappa\theta / (\kappa + \lambda)$ and $\theta^* = \kappa + \lambda$. Thus, only 5 parameters have to be estimated. Given Eq. (9), the RND can be easily inferred. Since the option pricing formula involves integrals, clearly, the computation of the RND will also involve integrals. For numerical purposes this evaluation will take a significant amount of time.

3. The data

The OTC data used were provided by a large French bank. Options are issued on a regular basis and reach maturity between a few days and one year. Anecdotal evidence suggests that market participants consider this market liquid. We were able to obtain data for 20 irregularly spaced dates.¹¹ The first one was 17 May 1996 while the last one was 27 June 1997.

As discussed in Section 2.1.2, this type of option is quoted in terms of δ . For all dates, we have at least information for options with δ taking the values 10, 15, 20, 30, 40, 50 (corresponding to the at the money option), 60, 70, 80, 85, 90. Between the first date and June 1996 we also have information for the 5 and 95 delta options. Since options in the extremes were rather illiquid, their quotation was given up at that time. In this study we used data for all possible δ .

For all dates, we were given bid and ask prices for in the money put and call options. Following the literature, we decided to work with the average between the bid and ask prices. Even though we obtained all results for options with 1, 2, 3, 6, 9, and 12 month to maturity, we decided to report the results for fewer maturities.¹²

The interest rates r and r^* are the domestic (French) and foreign (German) euro-currency interest rates chosen to match the expiration of the options. We transformed these rates into their continuously compound equivalents. The spot exchange rate is easily available.

By using a numerical procedure and the methodology outlined in Section 2.1.2 we extracted for each option of a given maturity the corresponding strike price. The difference between the actual data and the delta obtained for the optimal K was in all cases smaller than 0.07% of the initial delta!

4. Cross-sectional comparison

In this section we are going to present and interpret estimation results for two dates only. In the next section we will compare the methods within a time series context.

To get a feel for the data at hand, we trace the volatility of an option as a function

¹¹ Even strenuous efforts did not allow us to obtain more dates.

¹² The full set of estimates for the two dates can be found in a working paper version of this study.

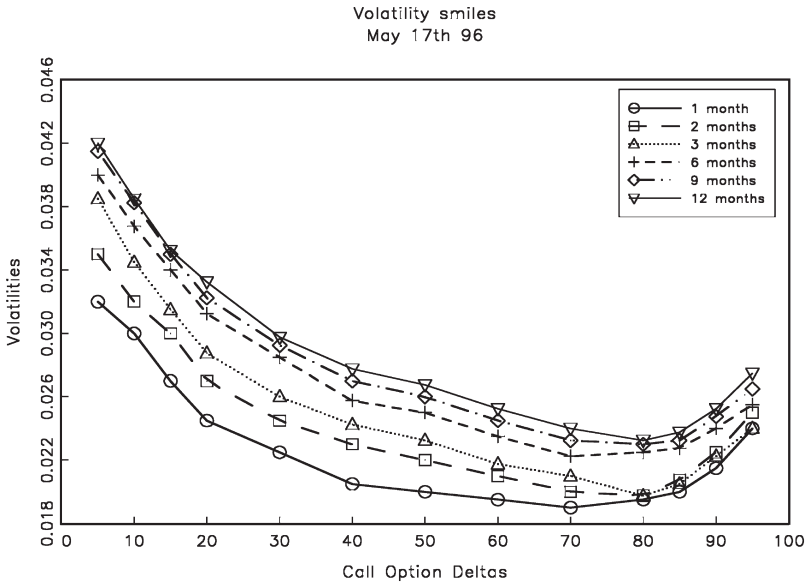


Fig. 1. Volatility smiles 17 May 1996.

of the delta and maturity in Figs. 1 and 2. If log-normality held, then we should observe one straight line independent of maturity. For a given maturity, the deviation from the straight line is called the volatility smile. The shift across maturities is the term structure of volatilities. Here options with low δ (high strike prices) are highly

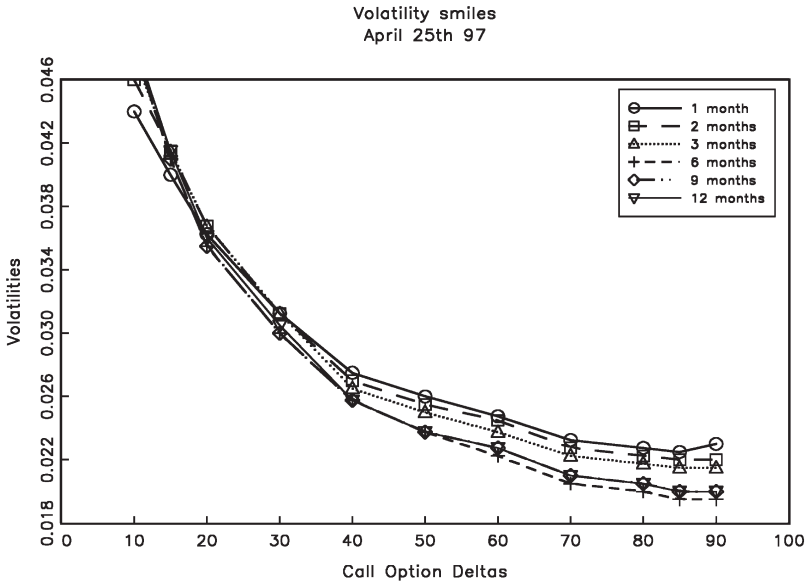


Fig. 2. Volatility smiles 25 April 1997.

valued, meaning that the market expects an increase in the exchange rate (a FF depreciation).

The smiles indicate that more complicated models than the Garman–Kohlhagen model should be considered for the data at hand. For future comparisons we nonetheless estimated this model, using the NLLS procedure outlined in Appendix A. This yielded for each data point and maturity a single volatility estimate. These volatilities are then used to construct a set of benchmark RNDs which will be presented later on for comparison purposes.

We also estimated the parameters for the other non-structural models. For the mixture of log-normals the values of the parameter estimates have no obvious explanation but they could be used to infer the various moments of the mixture density. For the Edgeworth expansion the parameters correspond to the volatility, skewness, and kurtosis of the underlying density. We decided, however, to compare the moments of all models simultaneously at a later stage. Before discussing moments we wish to present the parameter estimates for the structural models of Malz and Heston which do have an economic meaning.

4.1. Parameter estimates for structural models

4.1.1. The jump-diffusion case

We first turn to the parameter estimates for the jump-diffusion model of Malz presented in Table 1. Turning to the first date we notice that σ increases from 0.0172 to 0.0205. This means that investors expect a greater uncertainty about price movements in the longer run. The probability that a jump occurs before maturity, (λT) , varies from 0.0399 to 0.0699 suggesting that for the calm date investors do not believe in a great likelihood of a jump occurrence.

Turning to the expected jump size, (λk) , we notice that this measure decreases

Table 1
Estimates of the Bernoulli version of the jump-diffusion model^a

	1 month	3 months	6 months	12 months
17.05.96				
σ	0.0172	0.0178	0.0193	0.0205
λT	0.0399	0.0621	0.0655	0.0699
λk	0.0104	0.0095	0.0075	0.0058
25.04.97				
σ	0.0186	0.0176	0.0160	0.0165
λT	0.0717	0.0608	0.0600	0.0574
λk	0.0230	0.0128	0.0089	0.0063

^a This presents the parameter estimates for the Bernoulli version of a jump-diffusion. σ is the diffusion volatility. The jump will occur with probability λ within one year. Its size is k . λT represents the probability of a jump to occur before the maturity of the option. λk is the annualized impact of a possible jump. Parameters are always estimated using non-linear least squares as further explained in the appendix. All options are European. For 17.05.96 (25.04.97) we have options for 13 (11) deltas. The first date corresponds to a calm market whereas the second one to an agitated market.

from 0.0104 to 0.0058. This means that what is considered to be a jump in the short run becomes normal in the long run. To sum up, investors expect that a jump will occur with a higher probability in the long run but then only large variations will be considered jumps.¹³

Turning to the second date, when the market was more agitated, we notice that σ decreases across maturities. Further, for the one month to maturity, σ is higher for the new date than for the first date (0.0186 against 0.0172). In the long run, instead σ is smaller for the new date. Those results suggest that there was higher non-directional uncertainty for the short run after Chirac's announcement of a snap election: markets were expected by investors to either move up or down. In the long run, however, since then fundamental uncertainty given by σ is now smaller than for the first date, investors appear to anticipate the creation of a single currency area. Clearly, for a single currency area one expects σ to vanish completely.

The jump probability λT decreases from 0.0717 to 0.0574 showing that investors attach also a higher probability to a depreciation of the Franc in the short run. When turning to the impact of a jump on prices, given by λk , we notice its sharp increase relative to the first date and this for all maturities. The sign, which is always positive for this component, suggests that, if anything, the FF was expected to depreciate against the mark. To sum up, Chirac's announcement led to important market turbulences. On 25 April 1997 in an environment of agitated foreign exchange markets, investors expected that a jump of rather large magnitude would occur in the short run.

4.1.2. Stochastic volatility

After estimating this model for each maturity, given the great instability of the parameters across maturities, we decided to also report in Table 2, the estimates for the stochastic volatility model where for a given date we used all maturities simultaneously.

We notice for the first date that the long-run volatility ($\sqrt{\theta^*}$) increases from 0.0264 to 0.0349 whereas for the second date it decreases from 0.0720 to 0.0038. This variable captures a similar message than the diffusion volatility namely that on a normal, calm day there should be an upward sloping term structure of volatilities, and a decreasing one (or at least a less steep one) on a day with agitated markets.

The parameter ρ captures the skewness of the distribution, i.e. the probability of an asymmetric event. Its impact on the RND has to be read in combination with γ , the volatility of volatility. We notice for both dates that γ decreases whereas ρ increases with maturities. Those findings appear similar to the ones for λk of the jump diffusion model, namely, that in the long run an event has to be very large in order to be considered as a shock (i.e. to be generating skewness through ρ). In other words, in the long run most of the events are considered normal.¹⁴

Some of the parameter estimates display rather large variability. For this reason we also estimate the model with all maturities simultaneously. We first notice that

¹³ We are grateful here to Allan Malz for helping us getting the interpretations straight.

¹⁴ The situation of normality would correspond to a situation with $\gamma=0$.

Table 2
Parameter estimates of the stochastic volatility model^a

	1 month	3 months	6 months	12 months	Combined
17.05.96					
κ^*	3.2556	3.3781	3.4815	2.2940	4.0300
$\sqrt{\theta^*}$	0.0264	0.0362	0.0386	0.0349	0.0316
γ	0.1562	0.1423	0.1596	0.1064	0.1500
ρ	0.4497	0.5727	0.5434	0.5968	0.5430
$\sqrt{v_t}$	0.0221	0.0190	0.0020	0.0167	0.0224
25.04.97					
κ^*	3.2514	3.3687	3.4267	3.8023	3.2820
$\sqrt{\theta^*}$	0.0720	0.0432	0.0184	0.0038	0.0283
γ	0.3068	0.1837	0.1332	0.1430	0.1570
ρ	0.5176	0.6269	0.6226	0.6537	0.6170
$\sqrt{v_t}$	0.0185	0.0208	0.0367	0.0587	0.0300

^a This table presents the results for Heston's stochastic volatility model described by $dS_t = \mu S_t dt + \sqrt{v_t} dW_{1,t}$, $dv_t = \kappa(\theta - v_t)dt + \gamma \sqrt{v_t} dW_{2,t}$ where $W_{1,t}$, $W_{2,t}$ are two Brownian motions with possible correlation ρ . γ is the volatility of volatility. $\sqrt{v_t}$ is a measure of instantaneous volatility. κ and $\sqrt{\theta}$ represent the intensity of mean reversion and long run volatility. If λ is the risk premium then $\kappa^* \equiv \kappa\theta/(\kappa+\lambda)$ and $\theta^* \equiv \kappa+\lambda$. We estimated parameters in two stages, first running κ^* on a grid between 2 and 5 and then running an estimation with κ^* free using as starting value the optimal one from the first stage. The last column combines all maturities for a given date.

for the first date the measure of current volatility, $\sqrt{v_t}$, (0.0224), is smaller than for the second date, (0.030). This shows that the joint estimation is able to capture the increased market uncertainty due to political risk on the second maturity. The parameter ρ which captures the slope of the smile has also increased. The parameter $\sqrt{\theta^*}$ corresponds to the long-run volatility. This parameter takes the value 0.0316 for the first date and 0.0283 for the second one. This decrease in value confirms what we obtained with the jump-diffusion namely that investors are more confident on the second date that in the long run market volatility will be small because of a possible unique European currency. The parameter κ^* , capturing the speed by which volatility is mean-reverting, decreases from 4.03 down to 3.282.¹⁵ This means that for the more agitated date investors expect that it will take longer before the market reverts to normal. This observation is further corroborated by γ , the volatility of volatility. This parameter increases slightly from the first to the second maturity.

4.2. Moments for the various models

To further compare the different models we check the statistical properties of the various RNDs, displayed in Table 3. First, we verify that the first moment of the RNDs is equal to the forward rate. Second, we check how the constraints imposed

¹⁵ We notice here the large difference in the parameter estimates between the model with all maturities combined and the others. This illustrates our difficulties to pin down the mean-reversion parameter.

Table 3
Moments of the risk neutral density^a

	Forward	Volatility	Skewness	Kurtosis
17.05.96				
<i>1-month</i>				
Log-normal	3.3898	0.0202	0.0179	0.0006
Log-normal mixture	3.3898	0.0225	0.9096	4.3917
Hermite approximation	3.3898	0.0224	0.7127	3.2319
Edgeworth expansion	3.3898	0.0224	0.6898	3.2137
Jump-diffusion	3.3898	0.0219	1.2932	3.5955
Stochastic-volatility	3.3899	0.0215	1.1647	3.4252
<i>3-months</i>				
Log-normal	3.3933	0.0227	0.0348	0.0022
Log-normal mixture	3.3933	0.0253	1.3548	4.1869
Hermite approximation	3.3935	0.0248	1.1410	2.7441
Edgeworth expansion	3.3933	0.0251	1.0211	2.9609
Jump-diffusion	3.3933	0.0249	1.3715	3.0700
Stochastic-volatility	3.3932	0.0244	1.3375	3.6976
<i>12-months</i>				
Log-normal	3.4131	0.0267	0.0813	0.0118
Log-normal mixture	3.4130	0.0292	1.3369	3.6170
Hermite approximation	3.4132	0.0289	1.1495	2.5381
Edgeworth expansion	3.4131	0.0291	1.0215	2.7068
Jump-diffusion	3.4131	0.0289	1.2982	2.6747
Stochastic-volatility	3.4132	0.0284	1.4897	4.3789
25.04.97				
<i>1-month</i>				
Log-normal	3.3740	0.0257	0.0228	0.0009
Log-normal mixture	3.3740	0.0300	1.8572	6.1805
Hermite approximation	3.3741	0.0291	1.4135	3.4687
Edgeworth expansion	3.3740	0.0294	1.3080	3.6579
Jump-diffusion	3.3740	0.0291	1.6362	3.5315
Stochastic-volatility	3.3738	0.0284	1.4149	4.7203
<i>3-months</i>				
Log-normal	3.3758	0.0248	0.0382	0.0026
Log-normal mixture	3.3758	0.0307	2.3917	9.4481
Hermite approximation	3.3749	0.0307	1.3105	4.6650
Edgeworth expansion	3.3758	0.0299	1.3830	4.8034
Jump-diffusion	3.3758	0.0296	2.0354	5.2717
Stochastic-volatility	3.3757	0.0274	1.7039	4.9168
<i>12-months</i>				
Log-normal	3.3820	0.0241	0.0741	0.0098
Log-normal mixture	3.3801	0.0290	2.5141	9.2247
Hermite approximation	3.3787	0.0315	1.3475	5.2982
Edgeworth expansion	3.3820	0.0303	1.3888	5.6886
Jump-diffusion	3.3820	0.0297	2.2992	6.4565
Stochastic-volatility	3.3810	0.0271	1.8602	5.6927

^a Table 3 displays a comparison of various moments for the RNDs. For 17.05.96 (25.04.97) the actual forward prices for the 1, 3, and 12 month options are 3.3989, 3.3933, and 3.4131 (3.3740, 3.3758, and 3.3820).

by log-normality on the third and fourth moments can bias the variance estimates. Last, it is tempting to compare the estimates of the skewness and the kurtosis obtained under the different RNDs.

Some models (log-normality, Edgeworth expansion and jump diffusion) impose the constraint that the first non-central moment equals the forward rate. For other models, the better the adjustment, the closer the first moment is to the forward rate. We notice that for the first date, all the models give a first moment equal to the forward rate. For the second date however, the Hermite approach gives a small gap for the 3-months maturity (3.3749 instead of 3.3758) and similar for the 12 months maturity.

As far as volatilities are concerned, we see the bias implied by the log-normality assumption: the volatility induced by the log-normal model appears systematically smaller than what one obtains with the other approaches. Otherwise, we observe substantial homogeneity in the volatilities given by the other models.

The estimates of skewness and kurtosis are much more divergent, since at this level the specifics of the different models matter. The log-normal model is less interesting from this point of view, since on theoretical grounds it does not allow for asymmetry or fat tails. First, we observe that skewness as well as kurtosis are generally far from what is obtained under log-normality: for the first date for instance, skewness is between 0.68 and 1.48 and excess kurtosis is between 2.74 and 4.39. The skewness obtained from semi-nonparametric models is systematically lower than the skewness obtained with other models, although this difference is small. We do, however, notice pronounced differences, between models, of kurtosis; the log-normal mixture model and the stochastic volatility model generally give very large excess kurtosis (especially for the second date).

The graphs of the RND corroborate our earlier findings. All RNDs differ significantly from the benchmark. Further, we notice that the RNDs for the Hermite and Edgeworth expansion are very close. These two approaches have the unfortunate drawback of yielding negative densities. The reason for this is that only a limited range of skewness–kurtosis pairs are compatible with positive approximations.¹⁶ Going back to Table 3, we see that for those approximations skewness and kurtosis are always smallest: the reason is that those methods have difficulties accommodating higher moments beyond a certain range. Those models seem unable to capture the high skewness of exchange rate data.

When we inspect Figs. 3–6 we realize that the model with stochastic volatility distinguishes itself by a curvature which is less pronounced than the other models. This means that this type of model has difficulties in capturing the strong skewness which appears in the data. When going back to Table 3 we notice that the model with stochastic volatility always has smaller values of skewness but at times the largest kurtosis. This suggests that the stochastic volatility model is unable to capture the asymmetry in the data and suggests as a substitute for skewness a higher kurtosis.

¹⁶ See also Barton and Dennis (1952).

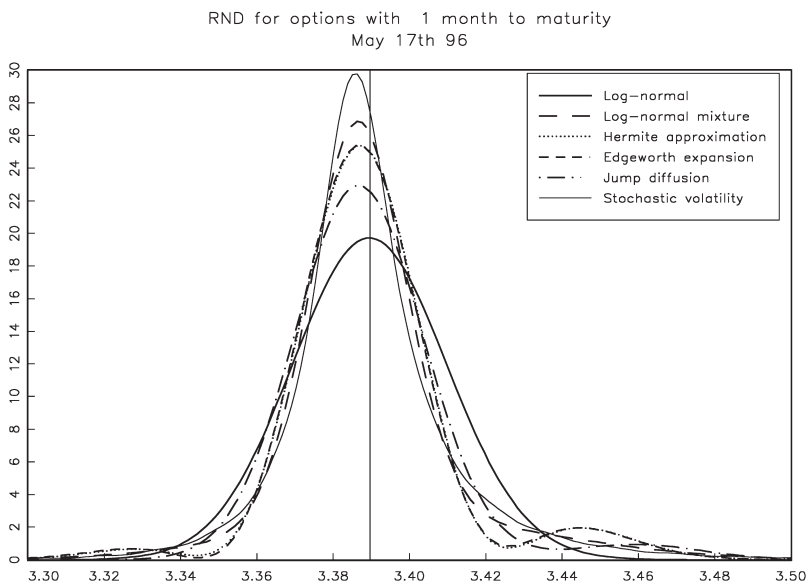


Fig. 3. RND for options with 1 month to maturity 17 May 1996.

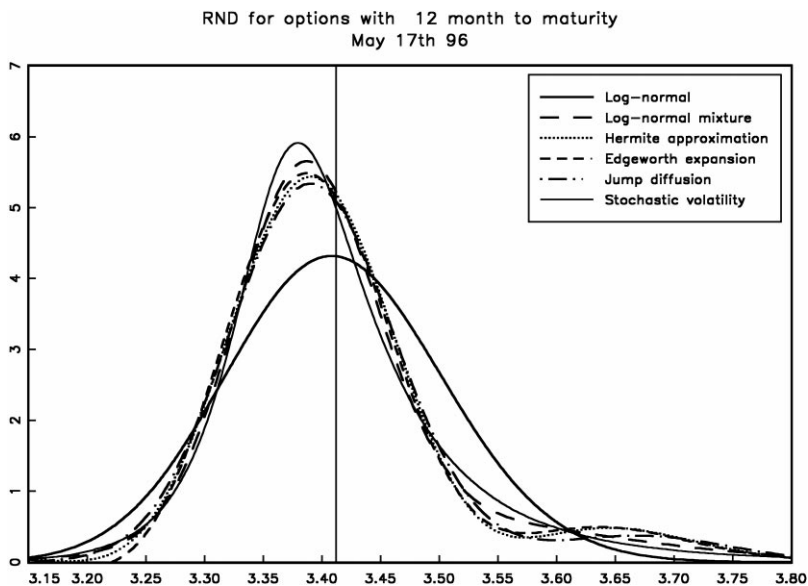
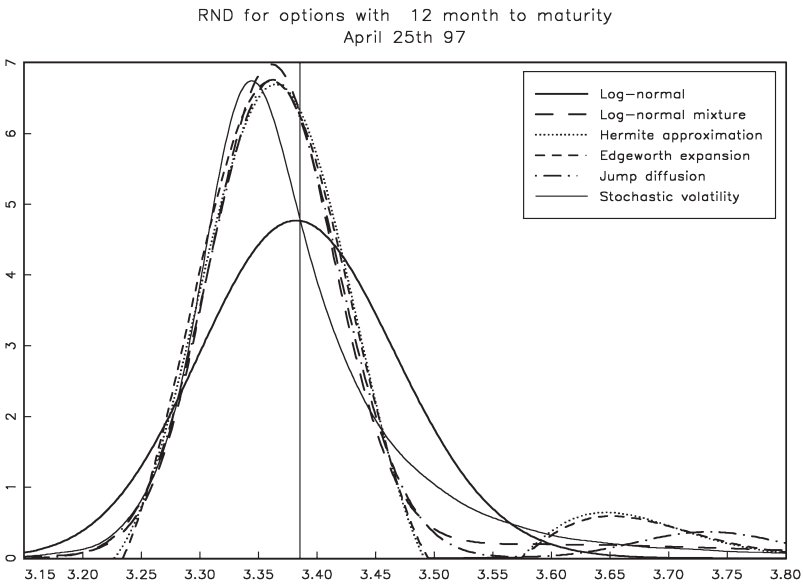



Fig. 4. RND for options with 12 months to maturity 17 May 1996.



In a situation in which fears are directional (such as for a devaluation) this feature seems to be somewhat beside the point.

To summarize, we notice a great deal of homogeneity for the different models as far as the first and the second moments are concerned. What really differentiates the models is their ability to capture the third and fourth moments.

4.3. *The use of RNDs*

An important point to check in the comparison of the various methods is whether they give similar confidence intervals. This point is of particular interest for policy-makers, since the bandwidth of confidence intervals can be seen as an indicator of credibility of monetary policy. As is well known, it is not possible to extract directly forecasts from option prices, since the underlying distributions are based on the assumption of risk neutrality of market participants. It might be argued that this type of analysis is misleading since one assumes risk neutrality. However, Rubinstein (1994, p. 804) using a numerical example states: "...despite warnings to the contrary we can justifiably suppose a rough similarity between the risk-neutral probabilities implied in option prices and subjective beliefs." For this reason we follow Campa et al. (1997) and construct RNDs which are based on the forward rate. In this case, confidence intervals are not interpreted in levels, because it is misleading to read a floor and a ceiling of an interval in FF/DM. One can analyze the relative intervals and the relative bandwidths expressed as a percentage of the forward rate. Thus, we estimate, for each maturity and each method, two confidence intervals: the bands of minimum width such that market participants put a 90% (and a 95%) probability on the fact that the FF/DM will be inside the band at the end of the period. As the RNDs are centered on the forward rate, we define the bandwidth as half the difference between the floor and ceiling expressed as a percentage of the forward rate.

Table 4 reports the estimates of the floor, the ceiling, and the bandwidth. Several points are worth noting: first, we clearly observe the asymmetry of the RNDs for all methods and all maturities since the forward to floor ratio is always smaller than the ceiling to forward ratio. For instance for the bandwidth containing 90% of the distribution on 17 May 1996 for the 1-month maturity, the former is about 0.85% whereas the latter is about 1.4%. For more distant maturities, the gap is even larger.

In the same way, we notice that the asymmetry increases for the second date, since the ceiling to forward ratio is at least twice the forward to floor ratio. This result clearly shows that the uncertainty on April 1997 was unfavorable to the FF.

Second, the excess kurtosis can be measured to a certain extent from the bandwidth. As clearly appears to be the case, for a given probability, the bandwidth of the log-normal model is always narrower than the ones of the other approaches. This means that, for a given bandwidth, the more sophisticated methods (which allow for fat tails) will give a higher probability outside the bandwidth than the log-normal model.

The comparison of the various methods is also interesting. The log-normal model shows no asymmetry since the forward to floor ratio and the ceiling to forward ratio are almost the same. Other approaches are much more homogeneous, except perhaps

Table 4
95 and 90 percent confidence intervals^a

	95% boundaries			90% boundaries		
	fwd/floor	ceil/fwd	bandwith	fwd/floor	ceil/fwd	bandwith
17.05.96						
<i>1 month</i>						
Log-normal	1.1950	1.1850	1.1830	0.9719	1.0030	0.9828
Log-normal mixture	1.1950	1.7310	1.4559	0.8606	1.2214	1.0374
Hermite approximation	1.0833	1.8402	1.4559	0.8236	1.4398	1.1284
Edgeworth expansion	1.0461	1.8038	1.4195	0.8236	1.4034	1.1102
Jump-diffusion	1.0461	1.8766	1.4559	0.8977	1.0394	0.9646
Stochastic-volatility	1.0833	1.6218	1.3467	0.8606	1.2214	1.0374
<i>3 months</i>						
Log-normal	2.3179	2.2798	2.2726	1.9386	1.9161	1.9089
Log-normal mixture	1.9764	3.5524	2.7452	1.6372	2.5706	2.0907
Hermite approximation	1.8631	3.5524	2.6907	1.5997	2.8979	2.2362
Edgeworth expansion	1.8631	3.5524	2.6907	1.5997	2.8979	2.2362
Jump-diffusion	1.9764	3.7706	2.8543	1.6748	2.7888	2.2180
Stochastic-volatility	1.9764	3.2615	2.5998	1.5997	2.3888	1.9817
<i>12 months</i>						
Log-normal	5.5041	5.4113	5.3141	4.5866	4.5075	4.4465
Log-normal mixture	4.6658	8.0502	6.2540	3.8797	5.9174	4.8261
Hermite approximation	4.5076	8.1225	6.2179	3.8407	6.5319	5.1153
Edgeworth expansion	4.3106	8.0502	6.0913	3.7239	6.4596	5.0249
Jump-diffusion	4.6658	8.4117	6.4348	3.9578	6.4235	5.1153
Stochastic-volatility	4.5076	7.5803	5.9467	3.7239	5.4836	4.5369
25.04.97						
<i>1 month</i>						
Log-normal	1.5095	1.4887	1.4879	1.2716	1.2573	1.2564
Log-normal mixture	1.2377	2.6129	1.9177	1.0347	1.8194	1.4218
Hermite approximation	1.1361	2.5137	1.8185	1.0010	2.1169	1.5540
Edgeworth expansion	1.1361	2.5468	1.8351	1.0010	2.1169	1.5540
Jump-diffusion	1.2377	2.6460	1.9342	1.0685	2.1500	1.6036
Stochastic-volatility	1.4074	2.2823	1.8351	1.1023	1.6210	1.3556
<i>3 months</i>						
Log-normal	2.5296	2.5227	2.4949	2.1144	2.0931	2.0819
Log-normal mixture	2.0111	4.8690	3.4202	1.7026	3.0845	2.3793
Hermite approximation	1.9768	4.7037	3.3211	1.7368	3.9767	2.8419
Edgeworth expansion	1.8052	4.6046	3.1889	1.6001	3.9106	2.7428
Jump-diffusion	2.0111	5.0672	3.5194	1.7368	4.0098	2.8584
Stochastic-volatility	2.1144	4.1750	3.1228	1.7026	2.8862	2.2801
<i>12 months</i>						
Log-normal	4.9931	4.9421	4.8489	4.1634	4.1175	4.0572
Log-normal mixture	3.8779	10.087	6.9105	3.3115	5.9646	4.5850
Hermite approximation	3.9491	9.7910	6.7950	3.4525	8.2736	5.8054
Edgeworth expansion	3.3115	9.4611	6.3332	3.0306	8.0427	5.4921
Jump-diffusion	3.9135	10.6156	7.1908	3.3819	8.2736	5.7725
Stochastic-volatility	3.9847	8.1417	5.9869	3.3115	5.5358	4.3706

^a This table displays 90 and 95 percent confidence intervals for the 1, 3, and 12 month options. Actual forward prices are as in Table 3.

for Heston's model. Indeed this model seems less asymmetrical than the other ones. More precisely, in many cases, the ceiling is nearer the forward rate. This result can be explained by the already mentioned fact that the stochastic volatility model is unable to generate a hump (as the Malz approach is) and, thus, it has to compensate the lack of flexibility with a less rapidly decreasing density (see also Figs. 5 and 6). Accordingly we note for instance for the 1-month maturity on 25 April 1997 an important gap between confidence intervals evaluated by Heston's model and by the other approaches: the bandwidth containing 95% of the distribution is 1.40–2.28 for Heston's model and about 1.03–2.61 for the other models.

5. Time series comparison

In this section we compare the performance of the various models and show how they can be used to read information contained in the data.

5.1. *Relative performance*

As a preliminary remark, we have to mention that we decided, in the time-series context, to drop the model with stochastic volatility. The reason for this is the obvious difficulty of that model in capturing the large skewness which appears to reside in the data at hand.¹⁷

Tables 5 and 6 show the absolute relative errors for the various dates and models. We notice in Table 5 that for the short maturity, for most of the cases, the mixture of lognormals is the best model. For the short maturity we notice further that the jump diffusion model also does quite well. Table 6 shows that for the longer maturity option Malz's model is the best except for one date. For practical purposes, this suggests that one should use for short-run options the mixture of log-normals model and for long-run options the jump-diffusion model.

5.2. *The message contained in confidence intervals*

As an illustration we display in Fig. 7 the evolution of the 90% confidence interval over the 20 dates for which we have information. We have chosen as a model the mixture of lognormals since it appears to be a good method for the short run.

In the summer of 1996 we observe a strong widening of the interval. Anecdotal evidence suggests that this is related to the political uncertainty in France. First, at the beginning of the summer there was a cabinet reshuffling and more importantly the financial markets had doubts about the ability of France to satisfy the public deficit criterion of the Maastricht treaty. The depreciation of the FF, therefore, was accompanied by the widening and an upward shift of the confidence interval. After a reassuring budget announcement, we see both an appreciation for the FF and a

¹⁷ We did not experiment with this model on other data for which it might well be optimal.

Table 5
AREs for 1 month to maturity options^a

Date	Bench	Mix	HE	ED	JD
17.05.96	157.168	16.834	11.925*	13.922	36.260
31.05.96	169.834	17.695*	21.909	23.334	37.281
14.06.96	164.191	17.284*	17.376	18.656	36.895
28.06.96	186.198	38.823*	72.850	71.874	40.947
5.07.96	177.835	26.298*	75.016	73.909	37.268
26.07.96	176.727	17.410	14.344	13.098*	39.891
23.08.96	203.223	10.548	28.258	28.692	9.283*
6.09.96	198.543	10.189	25.635	26.200	9.540*
4.10.96	161.028	13.723*	18.941	20.000	31.611
31.10.96	170.431	16.270*	28.076	20.476	22.907
8.11.96	170.431	16.277*	28.076	29.202	22.923
4.12.96	197.480	24.942	30.579	30.486	16.212*
27.12.96	168.213	15.368	12.680	13.113	8.278
30.01.97	168.217	15.372	12.679	13.114	8.276*
28.2.97	228.196	29.606	35.285	35.237	26.843*
3.04.97	251.562	22.192*	56.416	55.709	33.403
25.04.97	229.733	7.734*	18.155	17.831	8.955
2.06.97	237.059	16.900*	27.901	27.432	21.176
28.06.97	283.858	15.231*	53.130	51.707	22.162

^a This table presents the absolute relative errors (ARE) for the various models for the 1 month to maturity options. The * marks the model with the smallest error for a given day. The mnemonics Bench, Mix, HE, ED, JD stand respectively for the benchmark, mixture of lognormals, Hermite approximation, Edgeworth expansion, and jump-diffusion model.

narrowing of the interval. At the end of 1996 we see a new widening of the interval, but without a depreciation of the FF. This can be explained by heterogeneity of beliefs. If a small number of investors believe that markets may increase strongly and a large number of others believe that markets will move downwards each week then we expect that the confidence interval to widen but the forward rate to remain unchanged.

Later on, the interval regularly narrowed up to April 1997. At this time President Chirac announced a sudden election. Once again, the widening of the interval is associated with an upward shift: the forward exchange rate is about 3.38, and the market participants attach a 10% probability to the event of an exchange rate higher than 3.34. After the election and the victory of the left-wing coalition, the interval tends to decline significantly, but the upward shift clearly remains. This is associated with the new government's reassuring statement about EMU and its general economic policy.

6. Conclusion

In this paper we implement several methods that extract risk-neutral densities. The methods range from the non-structural (given by a mixture of lognormals) to the

Table 6
AREs for the 12 month to maturity options^a

Date	Bench	Mix	HE	ED	JD
17.05.96	176.747	32.930	40.432	39.437	23.238*
31.05.96	186.032	41.181	52.648	51.178	32.722*
14.06.96	191.256	33.007*	64.557	64.634	34.589
28.06.96	185.231	24.795	47.709	46.682	18.800*
5.07.96	200.735	56.325	81.774	81.302	47.909*
26.07.96	204.238	63.751	42.407	42.400	19.539*
23.08.96	215.753	24.484	28.366	28.349	15.502*
6.09.96	215.862	32.244	36.051	36.254	24.515*
4.10.96	168.531	37.457	30.615	32.494	19.109*
31.10.96	185.067	41.579	38.035	39.915	21.904*
8.11.96	177.486	41.652	36.035	37.010	17.371*
4.12.96	222.621	34.087	28.767	28.179	11.270*
27.12.96	232.423	38.671	42.990	41.428	14.013*
30.01.97	232.434	38.743	42.982	41.437	14.016*
28.2.97	251.202	35.277	58.668	56.635	23.769*
3.04.97	255.440	47.828	63.055	60.506	15.203*
25.04.97	248.672	41.034	38.826	37.143	16.694*
2.06.97	256.277	35.634	45.607	43.402	19.789*
28.06.97	291.310	63.565	81.166	77.084	34.425*

^a This is similar to Table 5 but for the 12 month to maturity options.

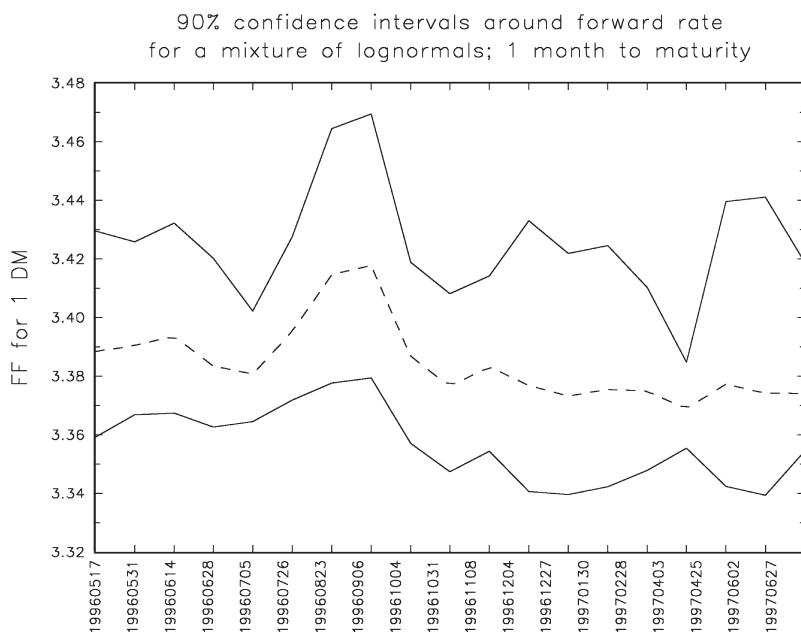


Fig. 7. 90% confidence intervals around forward rate for a mixture of log-normals; 1 month to maturity.

fully structural (a jump-diffusion and a stochastic volatility model). We also implement methods based on Hermite and Edgeworth expansions.

First, we compare these various methods for two dates. The first date is a rather calm one while the second date corresponds to an agitated market. We find that all models yield RNDs which differ significantly from the lognormal benchmark. Concerning stability and speed of estimation, we find that the mixture of lognormals and the stochastic volatility model require fixing some parameters on a grid and then estimating the remaining ones. This obviously results in a rather slow procedure. The other methods in contrast converge quickly and yield rather stable results.

We find further that models differ in their ability to capture the large skewness existing in the foreign exchange data at hand. In particular, the polynomial approximations and the stochastic volatility model have difficulties at this level.

Second, we compare the various methods on time-series data using as criterion the absolute relative error. We see that the mixture of lognormals model performs well on short-maturity options and that the jump diffusion model outperforms all models for longer maturities. The construction of confidence intervals reveals interesting patterns and shows their usefulness for policy makers and for investors who need to know what other market participants anticipate about a market's future.

Acknowledgements

The second author acknowledges help from the HEC Foundation and the European Community TMR Grant: "Financial Market Efficiency and Economic Efficiency". We are grateful to Bernard Dumas and Allan Malz for their valuable comments. The usual disclaimer applies. The Banque de France does not necessarily endorse the view expressed in this paper.

Appendix A

Here we describe how we implemented the non-linear least squares (NLLS) estimation. We first introduce some notation, then we discuss the traditional NLLS estimation. We then go on to explain how we estimated parameters in more difficult situations.

For a given date, we consider N options characterized by subscript i . The i th option has strike price K_i and maturity T . The market price, written C_{iT}^M , is given. Last, let $C_{iT}^X(\theta)$ be the theoretical price for the i th strike price and maturity T where θ is a parameter vector describing the RND associated with model X .

NLLS consists in finding the solution to the program.¹⁸

¹⁸ This type of program can be easily implemented within Gauss using the `optnum` module.

$$\min_{\theta \in \Theta} \sum_{i=1, \dots, N} (C_{iT}^M - C_{iT}^X(\theta))^2$$

where Θ is the domain to which θ can belong.

For some of the models, the parameter estimation turned out to be difficult. In particular, if parameters need to be obtained in a systematic way such as in the time-series framework, it becomes necessary to make sure that the algorithm does not diverge. In most cases, what did the trick was, first, to restrict parameters in certain intervals (such a restriction can be obtained by using a logistic transform) and, second, to force certain parameters to take values on a grid whereas the other parameters were obtained without restrictions. When a parameter was on a grid we eventually ran an unconstrained estimation using as starting values the estimates obtained over the grid that had a minimum error.

We encountered difficulties in the following cases: For the mixture of lognormals case we noticed that we often obtained parameter estimates in which all the weight was put on one density and yielding a degenerate density (with zero variance) for the density with no weight. Further experiments with this method revealed the existence of multiple minima. To mitigate this problem we decided to take the weight over a grid starting close to 0 and ending close to 1 and to estimate for each of the weights optimal parameters. We also decided to constrain, by using a logistic transform, the means of the various densities in a range deemed to be reasonable.

We encountered similar difficulties in estimating the stochastic volatility model. For this case, we forced the λ parameter on a grid to take values between two bounds chosen sufficiently wide apart to cover a reasonable range of values. All other methods tended to be fast and stable.

References

- Abadir, K.M., Rockinger, M., 1998. Density embedding functions, University of York, mimeo.
- Abken, P., Madan, D.B., Ramamurtie, S., 1996. Estimation of risk-neutral and statistical densities by Hermite polynomial approximation: With an application to Eurodollar Futures Options, Federal Reserve Bank of Atlanta, mimeo.
- Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. NBER Working Paper 5351.
- Bahra, B., 1996. Probability distributions of future asset prices implied by option prices. Bank of England Quarterly Bulletin August 1996, 299–311.
- Ball, C., Roma, A., 1994. Stochastic volatility option pricing. *Journal of Financial and Quantitative Analysis* 29, 589–607.
- Ball, C., Torous, W., 1983. A simplified jump process for common stock returns. *Journal of Financial and Quantitative Analysis* 18, 53–65.
- Ball, C., Torous, W., 1985. On jumps in common stock prices and their impact on call option pricing. *Journal of Finance* 50, 155–173.
- Barton, D.E., Dennis, K.E.R., 1952. The conditions under which Gram–Charlier and Edgeworth curves are positive definite and unimodal. *Biometrika* 39, 425–427.
- Bates, D.S., 1991. The Crash of '87: Was it expected? The evidence from options markets. *Journal of Finance* 46, 1009–1044.

- Bates, D.S., 1996a. Dollar jump fears, 1984:1992, distributional abnormalities implicit in currency futures options. *Journal of International Money and Finance* 15, 65–93.
- Bates, D.S., 1996b. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. *Journal of Financial Studies* 9, 69–107.
- Black, F., 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3, 167–179.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Campa, J.C., Chang, P.H.K., Reider, R.L., 1997. ERM bandwidths for EMU and after: evidence from foreign exchange options. *Economic Policy*, 55–87.
- Chesney, M., Scott, L., 1989. Pricing European currency options: a comparison of the modified Black-Scholes and a random variance model. *Journal of Financial and Quantitative Analysis* 24, 267–284.
- Corrado, C.J., Su, T., 1996. S&P500 index option tests of Jarrow and Rudd's approximate option valuation formula. *Journal of Futures Markets* 6, 611–629.
- Corrado, C.J., Su, T., 1997. Implied volatility skews and stock index skewness and kurtosis implied by S&P500 index option prices. *Journal of Derivatives* 4 (4), 8–19.
- Cox, J.C., Ross, S.A., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Derman, E., Kani, I., 1994. Riding on a smile. *Risk* 7 (2), 32–39.
- Duffie, D., 1988. *Security Markets: Stochastic Models*. Academic Press, Boston.
- Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: empirical tests. *Journal of Finance* 53, 2059–2106.
- Dupire, B., 1994. Pricing with a smile. *Risk* 7 (1), 18–20.
- Garman, M., Kohlhagen, S., 1983. Foreign currency option values. *Journal of International Money and Finance* 2, 231–238.
- Gesser, V., Poncet, P., 1997. Volatility patterns: theory and some evidence from the Dollar-Mark option market. *Journal of Derivatives* 5, 46–61.
- Harrison, J.M., Kreps, D., 1979. Martingales and arbitrage in multiperiod security markets. *Journal of Economic Theory* 20, 381–408.
- Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42, 281–300.
- Jackwerth, J.C., Rubinstein, M., 1995. Implied probability distributions: empirical analysis. Institute of Business and Economic Research, Finance Working Paper 250, University of California at Berkeley.
- Jarrow, R., Rudd, A., 1982. Approximate valuation for arbitrary stochastic processes. *Journal of Financial Economics*, 347–369.
- Johnson, N.L., Kotz, S., Balakrishnan, N., 1994. *Continuous Univariate Distribution*, Vol. 1, 2nd ed. John Wiley and Sons, New York.
- Jorion, P., 1989. On jump processes in the foreign exchange and stock markets, *Review of Financial Studies*, 431–445.
- Knight, J., Satchell, S., 1997. Pricing derivatives written on assets with arbitrary skewness and kurtosis, Trinity College, mimeo.
- Madan, D.B., Milne, F., 1994. Contingent claims valued and hedged by pricing and investing in a basis. *Mathematical Finance* 4, 223–245.
- Malz, A.M., 1996a. Options-based estimates of the probability distribution of exchange rates and currency excess returns, mimeo. Federal Reserve Bank of New York.
- Malz, A.M., 1996b. Using option prices to estimate realignment probabilities in the European monetary system: the case of Sterling Mark. *Journal of International Money and Finance* 15, 717–748.
- Melick, W.R., Thomas, C.P., 1997. Recovering an asset's implied PDF from options prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91–116.
- Melino, A., Turnbull, S., 1990. Pricing foreign currency options with stochastic volatility. *Journal of Econometrics* 45, 239–265.

- Merton, R., 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- Merton, R., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 125–144.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 771–818.
- Sherrick, B.J., Garcia, P., Tirapattur, V., 1996. Recovering probabilistic information from option markets: tests of distributional assumptions. *Journal of Futures Markets* 16, 545–560.
- Shimko, D., 1993. Bounds of probability. *Risk* 6, 33–47.
- Söderlind, P., Svensson, L.E.O., 1997. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics* 40, 383–429.
- Stein, E., Stein, J., 1991. Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies* 4, 727–752.
- Stutzer, M., 1996. A simple nonparametric approach to derivative security valuation. *Journal of Finance* 51, 1633–1652.
- Taylor, S.J., 1994. Modeling stochastic volatility. *Mathematical Finance* 4, 183–204.