BANCA D'ITALIA

Temi di discussione

del Servizio Studi

The Probability Density Function of Interest Rates Implied in the Price of Options

by Fabio Fornari and Roberto Violi

Number 339 - October 1998

THE PROBABILITY DENSITY FUNCTION OF INTEREST RATES IMPLIED IN THE PRICE OF OPTIONS

by Fabio Fornari and Roberto Violi (*)

Abstract

The paper contributes to the stochastic volatility literature by developing simulation schemes for the conditional distributions of the price of long term bonds and their variability based on non-standard distributional assumptions and volatility concepts; it illustrates the potential value of the information contained in the prices of options on long and short term lira interest rate futures for the conduct of monetary policy in Italy, at times when significant regime shifts have occurred.

Risk-neutral probability density functions (PDFs) of interest rates are estimated for several episodes since 1992, using an extended version of the model developed by Söderlind and Svensson (1997), in which the true PDF of an asset price is approximated by a parametric mixture of gaussian distributions with displaced means and mean-reverting deterministic volatilities. This approach generalizes the familiar Black and Scholes option pricing model by letting the returns of the underlying asset follow a mean-reverting stochastic process governed by the existence of two regimes. The parameters of this functional form are obtained by minimising the squared deviations between predicted and true option prices. The ability of the model to fit observed option prices seems encouraging; a significant degree of skewness (so-called *risk-reversal*), large changes over time and fatness of the tails (which gives rise to the *volatility smile* effect) are the elements which characterize the estimated PDFs).

The analysis of the information conveyed by PDFs of future interest rates begins in the wake of the 1992 EMS crisis, when a large negative skewness emerged in the distribution of lira-denominated bond futures, suggesting that the market factored in the possibility of a sharp decline in prices, triggered by a monetary tightening aimed at defending the exchange rate peg and by the mounting inflationary risk induced by a large devaluation. We subsequently examine a recent sequence of monetary easings, which provide an interesting variety of market participants' reactions.

The official rate reduction of July 1996, which followed a sequence of monetary tightenings in 1994-5, appears to have been discounted by the market, since the estimated PDFs of bond and 3-month T-bill futures prices did not change and an almost gaussian shape was maintained. The subsequent reductions - October 1996 and January 1997 - showed instead a different kind of market reaction, especially for long rate PDFs which returned from negative skewness and fairly high kurtosis, observed before the policy move, to normality, with narrower interquartile range and less kurtosis after the interest rate reduction.

(*) Banca d'Italia, Research Department.

Index

1. Introduction[1]

The use of forward-looking information extracted from financial asset prices is now a well-established practice among central banks, helping to formulate and implement monetary policy. The information content of financial asset prices has until recently been limited to the knowledge of market participants' expectations about the future level of such prices, which often reflect up-to-date macroeconomic data not yet published or available to policy-makers. For example, it is customary to relate changes in the forward yield curve to shifts in market perception of future short-term rates, through the achievement of new prospects for future output growth and inflation. More recently, also, monetary authorities, in the attempt to divine market expectations, have introduced various forms of inflation targeting, thus increasing the need for assessing the credibility of the regime as well as the market perception about future monetary policy.

The ever expanding liquidity of derivatives markets, coupled with greater international financial integration, has provided reliable new sources of information, among which option prices play a prominent role. They overcome traditional procedures based on the evaluation of the future price of an asset by providing the whole (yet risk-neutral) probability distribution of that price, which is then used to measure the range and the frequency of the possible outcomes around the central value.

The main aim of this paper is to illustrate the potential value for the conduct of monetary policy of information extracted from the prices of options on lira-denominated interest rate futures traded on LIFFE; this information has been particularly relevant in Italy, where significant regime shifts have taken place. The importance of this type of information in assessing market expectations of future monetary policy and monetary conditions and in choosing the timing and the effectiveness of official rates changes is highlighted.

---

1 The views expressed in the paper are those of the authors and not necessarily those of the Banca d'Italia. We wish to thank the participants at the QMF 1997 Conference in Sydney for many stimulating comments and especially the referee and Ignazio Angeloni for many useful suggestions. The usual disclaimer applies.

Section 2 of the paper surveys the available procedures for extracting the probability distribution (henceforth PDF) of financial asset prices. It starts with the traditional framework of Black and Scholes (1973), moves to a theoretical section presenting a simulation scheme for the conditional distribution of asset prices and their variability based on non-standard distributional assumptions as well as non-standard volatility concepts and ends with an extended version of an option pricing model developed by Söderlind and Svensson (1997) under the hypothesis that the true PDF is a parametric mixture of gaussian distributions with displaced means and mean-reverting, deterministic, volatilities. The last two parts of this Section are closely related, though they might appear barely complementary. One deals with the theoretical argument according to which — as is traditionally assumed in finance — prices evolve according to a specific continuous-time law of motion; this law is obviously thought to reproduce the price paths defined by the true — yet unknown — generating process. In this setup, which is further complicated by the adoption of non-standard assumptions, closed form solutions for the terminal distribution of prices are not available. Yet, it is our opinion that this is the appropriate way to tackle the problem. The second one — which is devoted to developing empirical procedures to extract information from option prices — obtains risk-neutral distributions through available prices of options. It is not our intention to compare the results of the two procedures; in fact, this is an applied paper and its main aim is to extract information from market data. Nonetheless an effort is made to inform the reader that accurate pricing strategies are available for options written on assets with very general forms of stochastic volatility, which can also be used to deliver the terminal distributions of the price of the assets. Sections 3 and 4 apply the methodologies of Section 2 to simulate the probability distribution of the BTP futures prices and their variability and to estimate the probability distribution from options written on long and short lira futures; the latter exercise is geared to investigating and interpreting some relevant monetary episodes of the nineties. Section 5 draws the main conclusions.

2. Option pricing and the identification of the probability distributions: survey of the literature and recent developments

In the nineties asset pricing has gradually moved from the attempt to generalise traditional univariate models so to create a multi-factor setting towards a *reverse engineering* approach, in which financial asset prices are used to estimate unknown parameters of the pricing function — for example volatility — relevant for understanding the behaviour of market participants.

The prices of financial assets have been traditionally used to derive the expectations of the future path of interest and exchange rates distributions, though such expectations were limited to the mean and the variance. Recent developments have been stimulated by emerging empirical evidence that the variance of asset price changes itself changes over time, so that knowing the expected values of the first two moments is insufficient to characterise the uncertainty perceived by *rational* market participants. The most recent models of this new stream of research use option prices to derive the whole probability distribution of the price of the underlying asset (see, for example, Neuhaus, 1995; Bahra, 1996; Malz, 1996; Melick and Thomas, 1997; Butler and Davis, 1998; Aït-Sahalia, Wang and Yared, 1998).

Within this framework, a theoretical model for the determination of option prices is required. However, since the true model used by market participants can, at best, only be approximated, one has to test more than a single specification, so as to limit the risk of getting a biased estimator of the volatility function.

2.1 *The Black and Scholes framework*

The Black and Scholes (henceforth BS) formula has been known to misprice *out-of-the-money* and *in-the-money* options since its original application in Black and Scholes (1972). What the subsequent research has noticed is that the volatility used by economic agents, obtained by inverting that formula, is not constant, but varies both over time and with the *moneyness* (i.e. the ratio of the spot price, S, to the strike price, K). Despite this bias, the BS model is still widely used by market participants, though in a peculiar fashion: in fact, it serves as a one-to-one mapping between volatility and the price of options. Thus, options are marketed according to their specific volatility (i.e. volatility referred to the time to expiration and the moneyness of the option) which, together with the other observable characteristics, is plugged into the BS formula to obtain the price.

The use of the BS formula to derive the full probability function of the price of the underlying asset was initially proposed by Breeden and Litzenberger (1978). Their result is based on the ability of portfolios of options to replicate the value of a pure Arrow-Debreu security which pays, at maturity, one unit of money conditionally on the realisation of a predetermined event (or state of nature). If one assumes writing N options with exercise prices separated by a unit of money (1, 2,...,N), the difference between the prices of a couple of options with contiguous strike prices, $C(i,\tau) - C(i+1,\tau)$, will always be unity, quite independent of the particular value chosen for i, provided that the terminal price of the asset is greater than or equal to i-1 and after choosing an appropriate unit of measure for the prices. Extending this reasoning, the difference

(1) $$[C(i,\tau)-C(i+1,\tau)]-[C(i+1,\tau)-C(i+2,\tau)]$$

will pay i, independently of the event that occurs, when the terminal asset price is i. If the difference between exercise prices is fixed and equal to ΔM, the value (P) of a portfolio made up of *call* options will be such that, by generalising (1):

(2) $$P(M,\tau;\Delta M)/\Delta M = [C(M+\Delta M,\tau)-C(M,\tau)]-[C(M,\tau)-C(M-\Delta M,t)]/(\Delta M)^2.$$

When one takes the limit of the value of the portfolio in (2), letting ΔM tend to nil, such value yields the density function of the price of the underlying asset at exercise price X, as the second partial derivative of the price of the option with respect to the exercise price:

(3) $P(M,\tau;\Delta M)/\Delta M = \partial^2 C(X,\tau)/\partial X^2$.

The first partial derivative of the call price with respect to the exercise price — which in the case of BS[2] reduces to $-e^{-r\tau}N(d_2)$ — gives the cumulative distribution function of the underlying up to X. More generally, in continuous time, the value of a call option under risk-neutrality is:

(3a) $C(X,T) = 1/(1+r_\tau)\int_{X,\infty}(S_T-X)f(S_T)dS_T$,

where $f(S_T)$ represents the risk-neutral probability density of the value of the underlying asset at the expiration date T, S_T.

Taking the partial derivative in (3a) with respect to the strike price X, yields:

(3b) $\partial C(S,T)/\partial X = -1/(1+r_\tau)[1-F(X)]$,

where F(X) is the risk-neutral cumulative distribution evaluated in X. Hence, the second derivative

(3c) $\partial^2 C(S,T)/\partial X^2 = 1/(1+r_\tau)f(X)$

provides the probability density, $f(X)=F'(X)$. Relations (3a)-(3c) between option values and strike prices allow us to estimate the probability distribution of the underlying asset.

2.2 *Valuing options on assets with stochastic volatility*

The main valuation schemes of financial asset prices, for both traditional cash and

---

[2] In the following expression r is the riskless rate of interest, τ the time to maturity, N(·) the cumulative normal density.

derivative instruments — including the BS formula — rest on the assumption that the price of the underlying asset follows, in continuous time, an integrated or *mean-reverting* stochastic process; this implies that the increments of the prices are generated by the following stochastic differential equation with fixed variance (where time subscripts have been omitted, S is the asset price, W a Brownian motion and σ, θ, μ, real fixed parameters)

(4) $$dS = (\mu - \theta S)dt + \sigma dW$$

The empirical evidence has strongly rejected the assumption of stability of the variance parameter, σ; it is a well established empirical regularity that volatility is subject to an unpredictable sequence of calm and turbulence, i.e. to heteroskedasticity.

With heteroskedastic financial market returns, the evaluation of derivative assets becomes an extremely difficult task. In these circumstances, the traditional homoskedastic approach suffers severe distorsions (see Hull and White, 1987); on the other hand, searching for closed-form solutions has so far required highly restrictive assumptions.[3]

In the Hull and White model, for example, the variability of the logarithmic rates of change of the price of the underlying asset is generated endogenously —as are the price changes — and follows a geometric stochastic process; their model comprises the following two equations:

(5) $$dS_t = \phi S_t dt + \sigma_t S_t dZ_{1,t}$$

(6) $$d\sigma_t^2 = \mu\sigma_t^2 dt + \xi\sigma_t^2 dZ_{2,t},$$

where S_t is the price of the underlying asset, σ_t^2 the variance of its (log-) changes, ϕ, μ, ξ, real parameters.

To derive a closed-form solution for the price of a call option based on model (5)-

[3] See, again, Hull and White (1987) and, more recently, Heston (1993), Duan (1995).

(6), Hull and White (1987) have to assume that the correlation between the increments of the two sources of risk, Z_1 and Z_2, is nil and that the intercept in the variance equation (6) is nil ($\mu = 0$). Though the latter hypothesis is not very restrictive and finds a theoretical justification in Nelson (1990), the former is consistent only with the assumption of market incompleteness, since a single financial asset (S) cannot possibly provide a hedge against two sources of uncertainty, Z_1 and Z_2, perfectly uncorrelated and hence not diversifiable.

The search for analytic solutions for the prices of derivative instruments has recently been associated with Monte Carlo simulations and with semiparametric evaluation schemes. The former tool is necessary whenever one attempts to generalise the Hull and White approach in (5) - (6); this happens since analytic solutions are no longer available and approximating schemes based on conditionally heteroskedastic autoregressive models represent a simple tool to recover the parameters of the stochastic differential equation which generates both the conditional mean and variance.[4] For example, Nelson (1990) and Fornari and Mele (1997a) have shown that the Hull and White system of difference equations can be approximated by an AR(1)-GARCH(1,1) scheme, the latter assuming that the logarithmic rates of change of an asset price (S) follow an autoregressive process of the first order, while their variance follows a GARCH(1,1) scheme. More complex schmes — better responding to real life — have also been suggested; in particular, if one supposes that the variance of the logarithmic rates of change of a price follows a Power ARCH scheme, i.e.

$$(7) \qquad \sigma_t^{\delta} = \omega + \alpha(|\varepsilon_{t-1}| - \gamma\varepsilon_{t-1})^{\delta} + \beta\sigma_{t-1}^{\delta},\ -1\leq\gamma\leq 1,\ (\omega,\alpha,\beta,\delta)\in R_+^4$$

with ε_t being an uncorrelated error term coming from (8), then Fornari and Mele (1997a) have shown that equation (7), coupled with

$$(8) \qquad r_t = \log(S_t/S_{t-1}) = \mu + \phi r_{t-1} + \varepsilon_t,$$

---

4 Note that also the Hull and White formula is based on a Taylor expansion around the average variance over the life to maturity of the option.

is the discrete time counterpart of the following system:

(9) $d\ln S_t = (\mu-\sigma_t^2/2)dt + \sigma_t dW_t^S$

(10) $d\sigma_t^\delta = (\omega-\varphi\sigma_t^\delta)dt + \zeta\sigma_t^\delta dW_t^\sigma$,

in which the concept of volatility is not restricted a priori, as in the traditional GARCH scheme, but can be estimated from the data; for example, when $\delta=1$ the relevant volatility concept is the standard deviation, while when $\delta = 2$ it is variance that matters, and so on. Within this framework one starts by estimating the parameters $(\mu, \phi, \omega, \alpha, \beta, \gamma, \delta)$ of the discrete time system (7) - (8), from which, in turn, the continuous time coefficients of (9)-(10) are recovered and employed in the simulation. The latter is based on a discretisation scheme of the Euler type for both models (5)-(6) and (9)-(10), which has the following form:

(11) $S_{t+1} = S_t \exp(r-\sigma_t^2/2+\sigma_t \hat{u}_{t+1})$

(12) $\sigma_{t+1}^2 = \sigma_t^2 \exp[-(\varphi+\zeta^2/2) + \zeta \hat{n}_{t+1}]$,

with $\hat{u}_t$ and $\hat{n}_t$ being drawn from a N(0,1) distribution and φ and ζ complex functions of $(\mu, \phi, \omega, \alpha, \beta, \gamma, \delta)$.

In this framework one can easily allow for the case of non-normality of the conditional distribution of ε_t by estimating a GARCH(1,1) model with non-normal density. This econometric specification becomes very general when one assumes that ε_t is conditionally described by the General Error distribution (g.e.d.), written as

(13) $\text{g.e.d.}_{(\upsilon)} = (\upsilon\exp(-0.5\Xi_\upsilon^{-\upsilon}|\xi|^\upsilon)/(2^{1+1/\upsilon}\Xi_\upsilon\Gamma(\upsilon^{-1}));\ \Xi_\upsilon^2 = (\Gamma(\upsilon^{-1}))/(2^{2/\upsilon}\Gamma(3/\upsilon));\ (\upsilon>0)$,

where $\Gamma(\cdot)$ is the Gamma distribution and υ its tail thickness parameter, which is able to provide a sufficiently high degree of flexibility in fitting asset price dynamics (Fornari and Mele, 1997a). The effect of the different distributional assumptions (i.e. normality or g.e.d.)

and the different volatility concepts (which depend on δ) will influence the simulated option prices by altering the values of the continuous time parameters (m and n hence φ and ζ).

2.3 *Derivative instruments as predictor of probability functions*

Among the parametric and semi-parametric models recently developed to extract information from option prices we chose to work with an extension of the approach of Söderlind and Svensson (1997); experiments have also been made with the implied volatility function of Dumas, Fleming and Whaley (1996). Such schemes relax the restrictive hypothesis of BS that volatility is fixed; in addition they enhance the econometric estimation of the PDFs, for they do not require an exact analytic expression for the pricing formula. To simplify the econometric specification, changes of the variability are specified in a non-parametric way and are attributed to the movements of the asset price or to some specific option features, such as their time to maturity or moneyness.

2.3.1 *The model of Söderlind and Svensson*

The model of Söderlind and Svensson (1997) assumes, as originally put forward in Merton (1976), the presence of different regimes characterised by high or low variability which give rise to jumps in the prices of financial assets. According to the authors, prices in each state would be generated by the combination of two lognormal bivariate distributions, which describe the evolution of the asset price itself (b) and of the interest rate (d), so that the price of a call option is evaluated as:

(14) $$C(X) = e^{-r\tau}\{\alpha_1[\exp(b^1+0.5\sigma_{bb}{}^1+\sigma_{bd})N[(b^1+\sigma_{bb}{}^1+\sigma_{bd}-\ln(X))/(\sigma_{bb}{}^1)^{0.5}]-$$
$$X\cdot N[(b_1+\sigma_{bd}-\ln(X))/(\sigma_{bb}{}^1)^{0.5}] + \alpha_2[\exp(b^2+0.5\sigma_{bb}{}^2+\sigma_{bd})$$
$$N[(b^2+\sigma_{bb}{}^2+\sigma_{bd}-\ln(X))/(\sigma_{bb}{}^1)^{0.5}] - X\cdot N[(b^2+\sigma_{bd}-\ln(X))/(\sigma_{bb}{}^2)^{0.5}],$$

where b^1 and b^2 are the means of the logarithm of the price of the asset in state 1 and 2, respectively; σ_{bb}^1 and σ_{bb}^2 are the variances of the mean of the logarithm of the price in state 1 and 2, respectively; σ_{bd}^1 is the covariance between the changes of the logarithm of the price and of the discount factor (d), equal in both states 1 and 2; X the exercise price of the option; α_1 and α_2 the relative weight of the two lognormal distributions, with $\alpha_1 + \alpha_2 = 1$.

It is easy to see that equation (14) is a generalisation of the BS formula. The parameters of interest $\phi = [b^1, b^2, \sigma_{bb}^1, \sigma_{bb}^2, \sigma_{bd}, \alpha_1, \alpha_2]$ are evaluated through nonlinear estimation techniques, minimising the squared deviations between the market prices of the calls and the theoretical prices generated by the model conditionally on the estimated parameters. In practice, the vector ϕ is found as:

(15) $$\mathrm{Min}_{\phi}[\Sigma_{i=1,m}(C_i - C(X_i,\tau_i,r_i))^2],$$

where $\{X_i\}_{i=1,..,m}$ represents the observed strike price range, $\{Ci\}_{i=1,...,m}$ the observed call prices, τ_i the life to maturity of the i-th option, r_i the interest rate for a maturity matching τ_i.

According to the original formulation of Söderlind and Svensson, the prices of the options are supposed to depend uniquely on their exercise prices, so that options with different maturity, or observed in different days, whose price is a function of different time horizons and values of the underlying asset, cannot be evaluated. We have therefore adapted equation (14) by plugging in these characteristics of the options, allowing the model to fit samples of option prices observed over time and across different maturities. We extend the model also by allowing for the existence of a term structure of volatility. This is accomplished by introducing a specification for the conditional variance developed in Jamshidian (1989), which provides more flexibility than the fixed parameters σ_{bb}^1 and σ_{bb}^2 of Söderlind and Svensson. Under these assumptions, the two variance functions are specified as:

(16) $$\sigma_{bb}^1 = \sigma_1[(1-\exp(-2a_1\tau))(1-\exp(-a_1\tau))^2]/(2a_1^3)$$

(17) $$\sigma_{bb}^2 = \sigma_2[(1-\exp(-2a_2\tau))(1-\exp(-a_2\tau))^2]/(2a_2^3).$$

Equations (16)-(17) are derived under the assumption of mean reversion for the price of the underlying asset.[5]

2.3.2 *Implied volatility functions*

In the model of Dumas, Fleming and Whaley (1996) particular emphasis is given

---

[5] Options written on 3-month eurodeposits differ from those written on bond prices, in that they regard yields and are therefore quoted in terms of price discount. For example, a futures on a 3-month rate of 7 percent quotes at 93, i.e. (100-7). Owing to this transformation a call option with strike price X is equivalent to a put option with strike price 100-X. In fact, the payoff of a put is

$$Max[X-B_{t,t-3},0] = Max[(1-B_{t,t+3})-(1-X),0],$$

where $B_{t,t+3}$ is the price of a 3-month bill; in the above relation, the right hand side is clearly the payoff of a call option written on the bond discount. The value of the call is given by

$$E(C_t) = \int_{100-X,\infty}[(1-B_{t,t+3})-(1-X)]f(1-B_{t,t+3})d(1-B_{t,t+3}),$$

where f is the risk-adjusted density function. If the bond discount is lognormal with mean μ and variance $\sigma^2\tau$, then the risk-adjusted distribution, f, is lognormal with mean μ* and variance $\sigma^2\tau$. Following Rubinstein (1976):

$$E(C_t) = e^{\mu^*+0.5\sigma^2\tau}N[(-\ln(1-X)+\mu^*+\sigma^2\tau)/(\sigma\tau^{0.5})]-(1-X)N[(-\ln(1-X)+\mu^*)/(\sigma\tau^{0.5})],$$

where the parameter μ* is determined by considering the value of a zero exercise price option. If X=1, i.e. 1-X = 0, the above relation becomes

$$E(C_t) = e^{\mu^*+0.5\sigma^2\tau}.$$

In this case the forward price of the option on the bond discount is the forward price of the discount itself. Hence

$$1 - F_{0,t}[B_{t,t+3}] = F_{0,t}[1-B_{t,t+3}] = e^{\mu^*+0.5\sigma^2\tau},$$

so that

$$\mu^* = \ln[1-F_{0,t}(B_{t,t+3})] - 0.5\sigma^2\tau.$$

Now one can substitute μ* to yield

$$C = B_{0,t}[(1-F_{0,t}(B_{t,t+3})]N[\ln(1-F_{0,t}(B_{t,t+3}))/(1-X) + 0.5\sigma^2\tau)/(\sigma\tau^{0.5})] - (1-X)N[\ln(1-F_{0,t}(B_{t,t+3}))/(1-X))+0.5\sigma^2\tau)/(\sigma\tau^{0.5})].$$

As in the preceding case, the specification of the variance σ^2 which we employ is that reported in equations (16)-(17); of course, this means we still retain the mixture of (log-)normal distribution assumption.

to simplifying the estimation of the *local volatility surface*, i.e. the volatility of an option with a given maturity and a given moneyness.

The aim is to find a specification for the volatility function, $\sigma(\cdot,\cdot)$, that is not based on any general equilibrium consideration or derived from a specific law of motion of the volatility, such that the prices of the options predicted with this estimated volatility match the observed prices. The term *deterministic* is employed to denote that only observable variables, perfectly known at time t, enter the estimates of the variance; the authors propose the following four specifications:

(18) $\sigma = a_0$

(19) $\sigma(X) = \sigma_0 + \sigma_1 X + \sigma_2 X^2$

(20) $\sigma(X,\tau) = \sigma_0 + \sigma_1 X + \sigma_2 X^2 + \sigma_3 \tau + \sigma_4 X\tau$

(21) $\sigma(X,\tau) = \sigma_0 + \sigma_1 X + \sigma_2 X^2 + \sigma_3 \tau + \sigma_4 \tau^2 + \sigma_5 X\tau,$

where X denotes the moneyness of the options and τ their time to expiration. Of course, model (18) corresponds to the standard BS analysis.

As in the case analysed in the previous section, we suppose that the overall distribution of long term interest rates is made up as a linear combination of two normal distributions. This implies that when volatility is reasonably approximated by equation (19) we will have two separate volatility functions, each referring to a specific state of the economy:

(19.1) $\sigma_1(X) = \sigma^1_0 + \sigma^1_1 X + \sigma^1_2 X^2$

(19.2) $\sigma_2(X) = \sigma^2_0 + \sigma^2_1 X + \sigma^2_2 X^2.$

The quadratic expressions (19.1) and (19.2) have provided the best performance in the original application of Dumas, Fleming and Whaley (1996), so they will be employed as natural candidates in our analysis.

3. Distributions based on the simulation of stochastic volatility schemes

As was recalled in the preceding section, a feature of Garch schemes that is relevant for financial applications is their ability to approximate models developed in continuous time as systems of stochastic differential equations. In this respect, recall that the evaluation of the price of stock options according to the Hull and White model can be based on the estimation of a GARCH(1,1) scheme, whose coefficients are employed to recover the parameters of the continuous time model; at this point, one simulates the latter to get the final value of the stock at the expiration of the option (S_T); at this moment the call is worth Max[S_T-X;0], with X being the exercise price. By so doing one also derives a number of final (i.e. at time T) prices of the underlying asset along with their variance, which can be used to approximate the probability distributions $f(S_T|S_t)$ and $g(\sigma^2_T|\sigma^2_t)$.[6]

Further, the class of GARCH models is rich enough to allow one to simulate the price and the variance of an asset according to a wide range of distributions and using different volatility concepts. For example, the Power Arch scheme of Ding, Granger and Engle (1993) easily allows for non-normality, by directly assuming that the forecast error of equation (8) follows a General Error distribution, which has the normal, the Laplace and the uniform as particular cases, and also for different volatility concepts, by modelling σ_t^δ rather than σ_t^2 (eq. 7).

We will apply this technique to Italian BTP futures, for which the simulation of

[6] The reader must be aware that all the methods developed and employed hereafter are suited for European options only. Unlike these, the options traded at LIFFE, which we use in the remainder of the paper, are American options, exercisable at any instant before maturity. It can be shown that early exercise of American call options is never optimal if the underlying asset does not pay dividends; even under these circumstances, early exercise of put options can become profitable. Futures, the underlying asset used in this paper, can be considered as assets paying a continuous dividend, equal to $F(1+e^{r\tau})$, with F being the price of the futures, r the risk-free interest rate, τ the appropriate time span. It can be shown (Melick and Thomas, 1997) that the price of an American call option falls between an upper and a lower bound and is higher than the price of an equivalent European option. The maximum width of the upper-lower bound interval is $e^{r\tau}$. American option pricing based on simulation schemes is not easily practicable, since one should take into account (at any step of the simulation) the likelihood of early exercise. On this aspect Hull (1988, ch. 9) illustrates how the simulation scheme could be set up. This further step falls outside the aim of the paper so that, in the remainder of this paragraph, the reader should consider the simulation scheme as providing the price of a European call option on interest rate futures.

the probability distribution will proceed as follows. System (22) below is made up of two equations, one describing the conditional mean of the price changes and one their conditional volatility, which is raised to the power δ, the latter estimated from the data. Further, there is a distributional assumption for the conditional errors, ε_t, which are treated as coming from a General Error function (g.e.d.(υ), with υ the tail-thickness parameter; see equation (13)). The econometric estimation of this system, known as Power ARCH, has provided the following results, where all the parameters are statistically significant, with the exception of the mean of the *variance* equation):

$$
\begin{aligned}
& r_t = \log(Btp_t/Btp_{t-1}) = 5.6710^{-5} + \varepsilon_t \\
(22) \qquad & \varepsilon_t \mid I_{t-1} \sim \text{g.e.d.}(1.263) \\
& \sigma_t^{1.033} = 5.9610^{-5} + 0.116964(\left|\varepsilon_t\right| - 0.402\varepsilon_{t-1})^{1.033} + 0.9025\sigma_{t-1}^{1.033}.
\end{aligned}
$$

The conditional distribution of the forecast errors departs from normality, as highlighted by υ, the tail-thickness parameter, which is significantly different from two, being nearly a Laplace. The volatility concept that shows the highest autocorrelation — hence is most useful in providing information about the riskiness of the market — is the standard deviation, δ being very close to unity. Under such conditions the Hull and White hypotheses are clearly rejected. However, the latter model, whose discrete time version is a GARCH(1,1) with conditionally normal errors, is estimated and employed as a benchmark. The values of the parameters that maximise the likelihood function in this case are:

$$
\begin{aligned}
& r_t = \log(Btp_t/Btp_{t-1}) = 5.6710^{-5} + \varepsilon_t \\
(23) \qquad & \varepsilon_t \mid I_{t-1} \sim N(0, \sigma_t^2) \\
& \sigma_t^2 = 3.136 \cdot 10^{-9} + 0.065536\varepsilon_{t-1}^2 + 0.931225\sigma_{t-1}^2.
\end{aligned}
$$

The continuous time parameters employed in the simulations of the option prices based on the two models, i.e. (m, n, φ and ξ) are derived as combinations of the discrete time coefficients (γ, υ, δ, α, β); they are reported in Table 1.

Table 1

PARAMETERS OF MODEL (11)-(12) IMPLIED BY THE CONTINUOUS-TIME PARAMETERS OF A GARCH(1,1) AND A POWER ARCH(1,1) *

	GARCH(1,1) - Hull and White option pricing model	Power ARCH(1,1) - option pricing under non-standard assumptions
γ (dt)	Restricted to 0	-0.402
υ (dt)	Restricted to 2	1.263
δ (dt)	Restricted to 2	1.033
m (ct)	**1.5**	**0.51668**
n (ct)	**0.5**	**0.3732**
α (dt)	0.065536	0.11696
β (dt)	0.93123	0.9025
φ **(ct)**	**0.003234**	**0.009951**
ξ **(ct)**	**0.065536**	**0.067028**

* γ captures the asymmetric reaction of volatility to shocks of opposite sign; υ is the tail-thickness parameter of the g.e.d. distribution: when it equals 2 the g.e.d. is a normal; δ captures different volatility concepts; α and β are the parameters of the discrete time equation for the conditional variance. The estimation is performed on BTP futures prices, observed daily from October 15, 1991 to February 20, 1996, a sample which includes 1100 observations; dt and ct denote discrete time and continuous time, respectively. The expression of the parameters m, n, φ and ξ are reported in Fornari and Mele (1997a).

The distributions of the BTP futures prices and their standard deviation have been evaluated for the final day of the sample, February 20, 1996; that day, the closing price of the futures was 107.48, while the conditional variance of the BTP futures price changes, evaluated under the two models, i.e. Power ARCH and GARCH, was $6.868 \cdot 10^{-5}$ and $1.187 \cdot 10^{-4}$, respectively. For each of the two cases we performed 2000 replications of the experiment. The distributions of S_T and σ_T^2 are observed at T= 5, 10 and 20 days after February 20 and are reported in Figures 1-4, along with sensitivity analyses performed with concern for the initial level of the conditional variance (the level at which the simulation is started off); the parameters (δ, υ, γ); the persistence of the conditional heteroskedastic model (α+β). The first column in Figure 1 shows the PDFs of the BTP futures prices based on an initial level for the variance which is four times higher than the maximum value recorded in the sample; top to bottom, the distributions are evaluated at T = t+20, t+10 and t+5. Given the high persistence of the conditional variance equations (the sum of the relevant parameters was close to unity), the terminal distributions tend to have increasing variance; for example after 5 days the futures price is bounded between 80 and 160 whereas after 20 days it ranges between 50 and 200, while. Of course, this experiment is heavily dependent on an extremely high value for the initial volatility. In the second column, where the volatility is started off at the maximum value recorded in the sample, the BTP futures prices range between 84 and 132 after 20 days and between 95 and 120 after 5 days. When the variance is started off at the minimum value of the sample (third column) the price range is very narrow: from 106.2 to 109 after 5 days and from 104 to 112 after 20 days. The fourth column has to be compared to the second, being based on the same initial level of the conditional variance; it differs from the second column in that the parameters (γ,δ,υ) have been fixed at (-1,0.5,0.5) which is an extremely rare occurrence in real life, denoting maximum leverage effect (the correlation between price and volatility), as well as heavy peaks and heavy tails for the distribution of S_T. Even under such extreme hypotheses, the effects on the tails and on the asymmetry are evident, yet not so strong. It is important here to recall that the distributions of S_T are analytically unknown under the hypotheses of system (22).

Figure 1

PDF OF THE BTP FUTURES: PARCH MODEL (*)

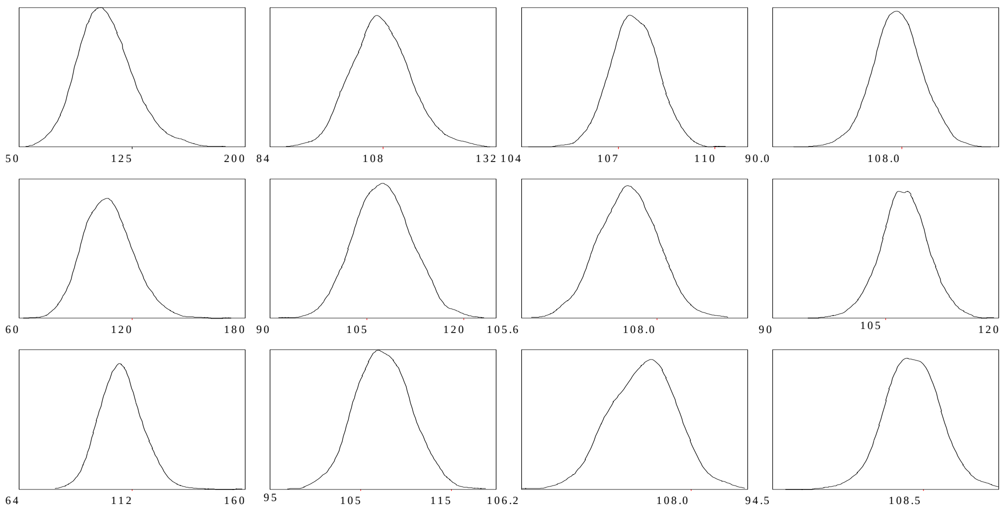

(*) The above simulation is carried out starting the BTP futures price at 107.48 the value of February 20, 1996. In the first column the conditional variance is started off at four times its maximum value in the sample; in the second column the conditional variance is started off at the maximum value recorded in the sample; in the third the conditional variance is instead started off at the minimum value; in the fourth the parameters γ, δ, υ are fixed at -1, 0.5, 0.5, respectively. Bottom to top the distributions are evaluated 20, 15 and 5 days after time t, respectively

Figure 2 shows the conditional distributions of the standard deviation of the BTP futures price changes evaluated with the preceding simulation based on the Power ARCH scheme. Unlike the distribution of S_T, the PDF of σ_T^2 is analytically known and converges to a lognormal, as evidenced in (12). In the first column, where the variance is started off at a very high level, the distribution of the conditional standard deviation is extremely skewed to the right, reaching values up to 140 percent per year 20 days after the beginning of the simulation; in the other cases the distribution of the standard deviation is rather stable through time, ranging from 1 to 40 percent in all of the remaining 9 cases.

Figure 2

PDF OF THE BTP FUTURES VARIABILITY: PARCH MODEL (*)

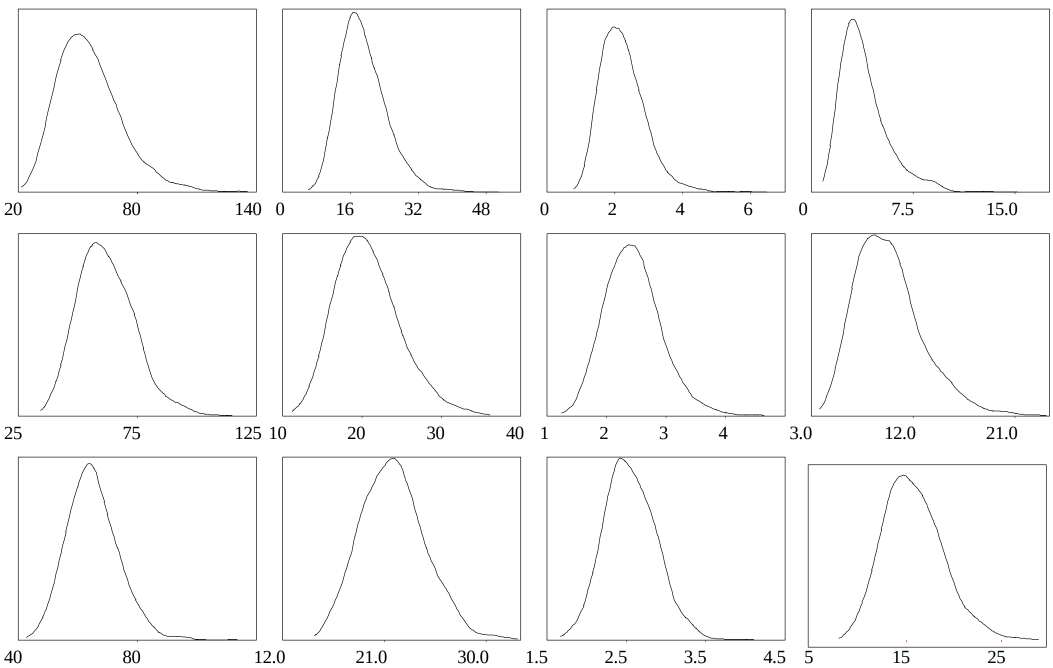

(*) The above simulation is carried out starting the BTP futures price at 107.48, the value for February 20, 1996. In the first column the conditional variance is started off at four times its maximum value in the sample; in the second column the conditional variance is started off at the maximum value recorded in the sample; in the third the conditional variance is started off at the minimum value; in the fourth the parameters γ, δ, υ are fixed at -1, 0.5, 0.5, respectively. Bottom to top, the distributions are evaluated 20, 15 and 5 days after time t, respectively.

Figures 3 and 4 report analogous simulations based on the coefficients of the GARCH(1,1) model, which corresponds to simulating the Hull and White scheme. Here, also to highlight other aspects of the technique, the distributions have been graphed in a different way. Figure 3 shows four terminal distributions of the underlying price, corresponding to 20 and 5 days after time t and having four times the maximum and the maximum volatility as starting points, respectively.

The effects of the initial level of the variance are clearly detectable by comparing curves a) and c), which differ only in the starting conditions. Figure 4 also unambiguously evidences this effect: the means of curves a) and c) are significantly different and the standard deviation of the latter is much lower.

Figure 3

PDFS OF THE BTP FUTURES: GARCH MODEL (*)

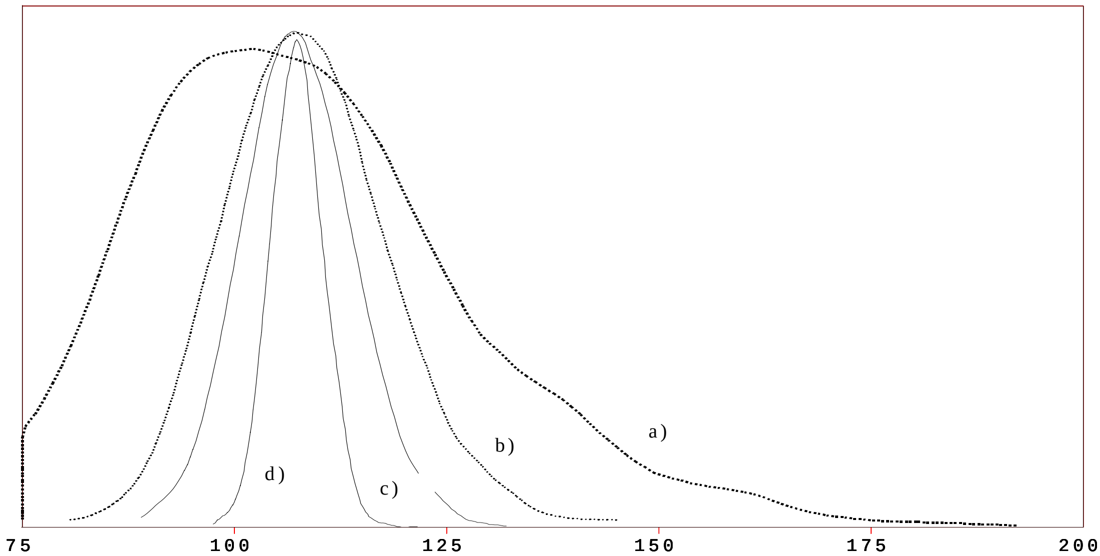

(*) The above simulation is carried out starting the BTP futures price at 107.48 the value for February 20, 1996. Curve a) corresponds to 20 days ahead, starting off the variance at four times its maximum level in the sample; curve b) corresponds to 5 days ahead, starting off the conditional variance at four times its maximum value in the sample; curve c) corresponds to 20 days ahead, starting off the conditional variance at the maximum value recorded in the sample; curve d) corresponds to 5 days ahead, starting off the variance at its minimum value.

Figure 4

PDFS OF THE BTP FUTURES VARIABILITY: GARCH MODEL (*)

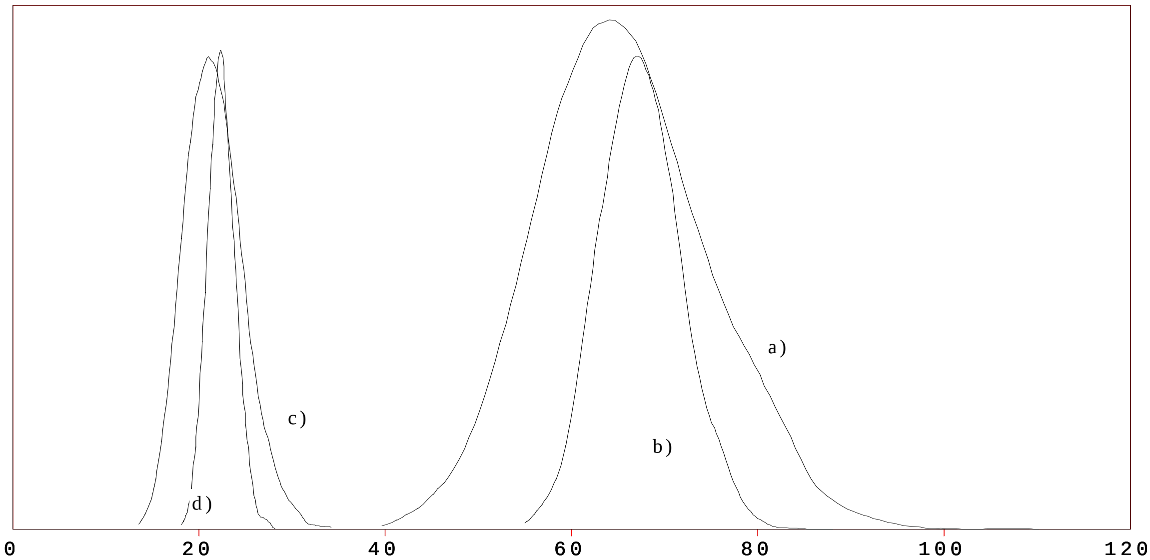

(*) a) is the 20-day-ahead distribution obtained starting off the variance at four times the maximum; b) is the 5-day-ahead distribution obtained starting off the variance at four times the maximum; c) is the 20-day-ahead distribution obtained starting off the variance at the maximum; d) is the 5-day-ahead distribution obtained starting off the variance at the maximum.

Figure 5

PDFS OF THE BTP FUTURES: GARCH MODEL (*)

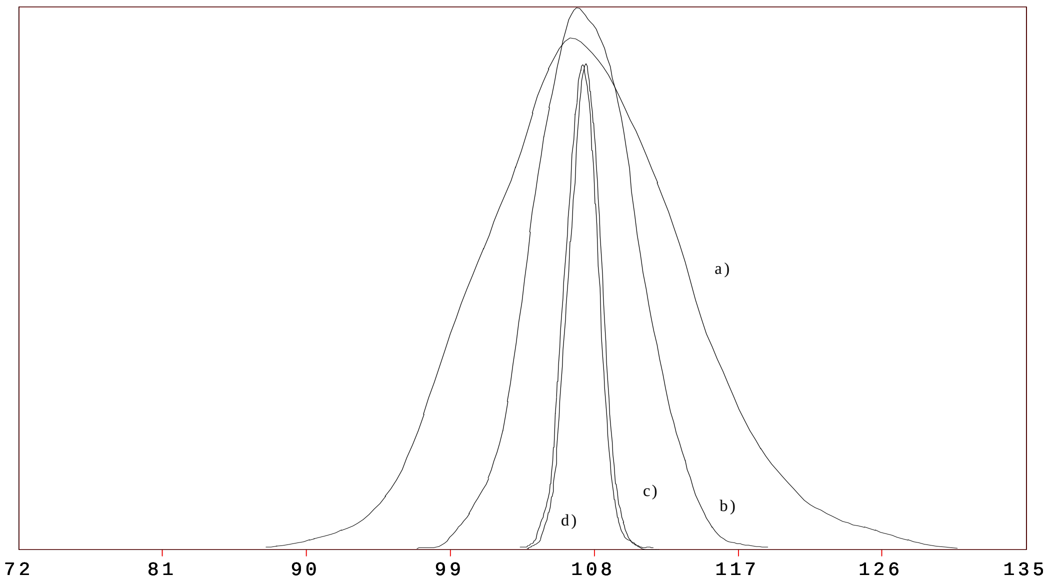

(*) a) is the 20-day-ahead distribution obtained starting off the variance at four times the maximum with the estimated persistence; b) is the 5-day-ahead distribution obtained starting off the variance at four times the maximum; c) is the 20-day-ahead distribution obtained starting off the variance at the maximum with the persistence set equal to 0.5; d) is the 5-day-ahead distribution obtained starting off the variance at the maximum with the persistence set equal to 0.5.

Figure 6

PDFS OF THE BTP FUTURES VARIABILITY: GARCH MODEL (*)

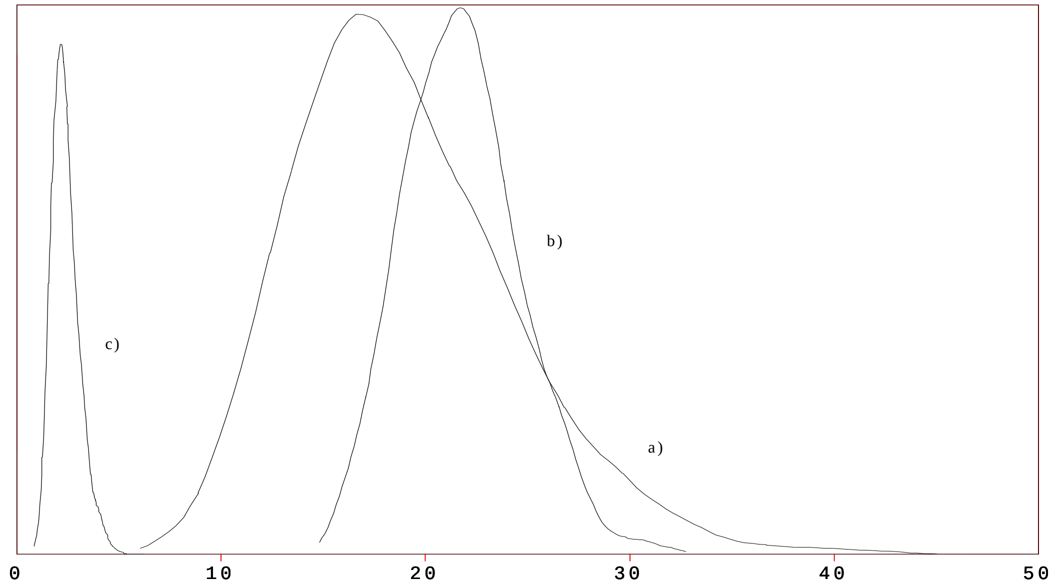

(*) a) is the 20-day-ahead distribution obtained starting off the variance at four times the maximum with the estimated persistence; b) is the 5-day-ahead distribution obtained starting off the variance at four times the maximum; c) is the 5-day-ahead distribution obtained starting off the variance at the maximum with the persistence set equal to 0.5; the 20-day-ahead distribution collapses to a constant close to nil.

The last set of simulations is reported in Figures 5 and 6. Here the conditional variance is always started off at the maximum, but in the cases labeled c) and d) the persistence of the stochastic process that generates the variance has been reduced to 0.5, while the true value of the parameter was close to unity. The effect of the persistence is clearly evidenced in Figure 5: in 20 days the price of the BTP futures is expected to range between 90 and 135 under normal circumstances, while the range narrows to just (104, 112) when the persistence is reduced. Why this happens is clear when one observes Figure 6: the lower persistence forces the standard deviation to collapse to a constant, small, value after 20 days; yet after 5 days it can range between 0 and 6 only, dramatically less than the true interval of 15-33.

4. Implied PDFs

Section 3 has shown how the PDF of financial prices can be estimated within a rigorous theoretical setting. In this case, however, one has to resort to Monte Carlo schemes, since closed form solutions are not available. This problem can also be overcome by considering derivative instruments, which are able to back some of the unknown parameters as evaluated by economic agents at the time of setting the price. In this section we use options written on lira-denominated long and short term interest rate futures traded on LIFFE, the main European exchange for this kind of contract, between October 1991 and January 1997.[7] No attempt is made to compare the results of this section to those of the previous one; there, a theoretical scheme for pricing options under general hypotheses is presented and, being analytically intractable, solved by simulation; here, empirical procedures are developed to estimate the PDFs from market data on interest rate options.

4.1 *Black and Scholes-based PDFs*

Now we briefly show the potential information that can be derived through the BS model (see Section 2.1) at selected dates, characterised by potentially destabilising events; we have examined the changes in the shape of the PDFs and their ability to anticipate future movements of interest rates. We leave the examination of other dates, characterized by official interest rate movements, to the more reliable model of Söderlind and Svensson, in terms of ability to reproduce the features of the data.

It is important to recall here that the BS model is heavily dependent on the assumption of constant volatility, which is strongly rejected in empirical investigations,;

[7] The valuation schemes presented hereafter are suitable for European options only, though LIFFE trades American-style interest rate options. Melick and Thomas (1997), with concern for crude oil options, derive the width of the interval which embraces the price of the option. The figure that they derive is negligible, reaching a maximum of 0.8 percent of the lower limit of the option price. Analytically, the width of the interval, i.e. the ratio between upper and lower bound of the price of an American call option is at most $e^{r\tau}$. If τ equals one day (i.e. 1/250 = 0.004) and the risk-free rate of interest is 7 per cent (i.e. 0.07) then the upper bound exceeds the lower bound by a factor of 1.00028. Thus if the lower bound for the price of the option is 4, the upper bound is 4.00112.

actually volatility changes over time. In the theoretical models developed in Sections 2.3.1 and 2.3.2 the problem was overcome by assuming that the variance is a function of observable characteristics of the options, but this strategy failed to unveil the assumptions underlying the pricing model employed by economic agents. When one deals with the inversion of the BS formula, however, the variance is implicitly taken as stochastic, yet without making any assumption about the properties and the characteristics of the latter; as the data reported hereafter evidence, despite the ex-ante assumption of normality for the log-changes of the asset price, the distributions are not normal, and their shape is analytically unknown.

The distributions evaluated in the period that led Italy and the UK to withdraw from participation in the ERM are reported in figure 7a. The PDF observed on August 6, 1992 was close to normality, though sharply asymmetric towards values lower than the median, equal to approximately 96; this feature implied that, at the beginning of August, the probability of the 10-year BTP futures price going below 89 was 3 percent, while the symmetric event under normality (that it would exceed 103) had no probability. On August 21 the shape of the distribution had changed significantly, with a reduction of the probability of observing values close to the current level, a shape which remained unchanged until September 14; the 24th of the month, after the lira's exit from the ERM, the PDF gradually moved back towards the original configuration. This change is more clearly evidenced by the data reported in Table 2, on the probability of the BTP price exceeding a given threshold by December 1992.

Despite similar values of the median, ranging from 96 to 97 at the four dates, the probability that the BTP price would go below 90 increased substantially from 5 percent on August 6 to 21 percent on September 11; after the lira's exit from the ERM the probability declined, reaching 14 percent on September 21, the same level as in early September.

Analogous shifts in the PDFs can be observed in the wake of the interest rate rise of August 12, 1994 (Figure 7b and Table 3). The distribution observed on August 11, compared with that of August 5, appears to be more asymmetric and with more probabilities in the left tail; a shift to the left hand side is partly attributable to the fall of BTP prices in those days, as confirmed by comparing the medians of the two distributions, 103.5 and 100.5, respectively.

Figure 7

PDF OF THE BTP FUTURES

(Black and Scholes method)

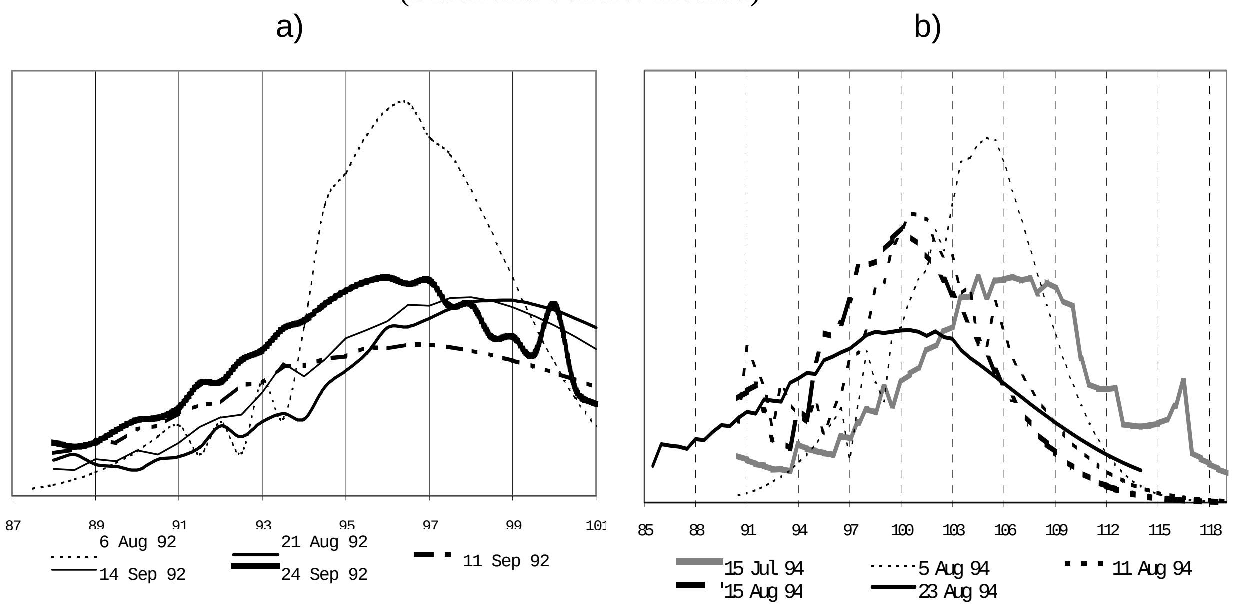

Table 2

PROBABILITY OF THE BTP PRICE EXCEEDING GIVEN THRESHOLDS

	Aug. 6, 1992	Sept. 1, 1992	Sept. 11, 1992	Sept. 17, 1992
median	96	96.5	96	97
P(BTP≥90)	95	88	79	86
P(BTP<90)	5	12	21	14
P(BTP≥92,5)	90	78	71	79
P(BTP<92,5)	10	22	29	21
P(BTP≥95)	68	62	58	65
P(BTP<95)	32	38	42	35
P(BTP≥97,5)	29	41	42	46
P(BTP<97,5)	71	59	58	54
P(BTP≥100)	6.3	22	27	27
P(BTP<100)	93.7	78	73	73

Table 3

PROBABILITY OF THE BTP PRICE EXCEEDING GIVEN THRESHOLDS

	Aug. 5, 1994	Aug. 11, 1994	Aug. 23, 1994
median	103.5	100.5	99
P(BTP≥90)	98.7	91	81
P(BTP<90)	1.3	9	19
P(BTP≥92,5)	97.8	83	75
P(BTP<92,5)	2.2	17	25
P(BTP≥95)	94.9	76	66
P(BTP<95)	5.1	24	34
P(BTP≥100)	79.6	53	44
P(BTP<100)	20.4	47	56
P(BTP≥105)	39.5	20	22
P(BTP<105)	60.5	80	78
P(BTP≥110)	6	4	9
P(BTP<110)	94	96	91

4.2 *The extended version of the Söderlind and Svensson model*

As reported in Section 2.3.1, the Söderlind and Svensson (1997) model has been extended to be used with *panels* of options prices, by inserting in equation (14) the current value of the underlying *futures* price and the time to maturity of the option.[8] In this framework, the parameters b^1 and b^2 are no longer the means of the two normal distributions, whose weighted average yields the overall probability distribution of the underlying asset; rather, they become a measure of the average deviation between the theoretical *futures* price and its current value.

This model has been estimated on option price samples that include 5 trading days before and after a selected event, avoiding in any case overlapping of the data. For the most recent reductions of the discount rate, when the market had reached satisfactory levels of liquidity, the estimation has been carried out with data spanning a single trading day. In the first case one has to assume that parameters are constant within the sample, though they are

[8] To provide an example, the term $\exp(b^1+0.5\sigma_{bb}{}^1+\sigma_{bd})$ in the original formulation has been modified as $\exp(\log(S)+ b^1+0.5\sigma_{bb}{}^1+\sigma_{bd})\tau$, where S is the current level of the futures and τ the time to maturity of the option.

allowed to change across samples. Estimates of the relevant coefficients are obtained by maximum likelihood, which under normality of pricing errors reduces to minimising the deviation between theoretical and observed option prices.

Table 4 shows the estimated coefficients of the option pricing model for the BTP futures at various dates. In all cases the null hypothesis that the true probability distribution is a mixture of two normals cannot be rejected at any level of significance. The two combining distributions often have zero and positive mean, respectively; this implies that one of them is centred on the current value of the futures (i.e. its mean equals the futures price), the other on higher values, so that the resulting mixture is positively skewed. The two components contribute to the overall distribution to different extent, with the weight of the less important ranging between 15 and 45 percent. However, as it also has higher standard deviation, the latter controls the size of the area in the upper tail of the overall distribution. This effect is especially evident in Figures 9a-b, which show the distributions estimated on October 23 and 27, 1992, and expected over a 1-month time horizon.

The extended Söderlind and Svensson model provides a much more regular shape for the PDFs, compared to that obtained via BS; the shapes are also more stable over time, partly also because the coefficients are estimated on a larger sample (10 trading days against a single trading day).

Figure 8

PROBABILITY DISTRIBUTION OF BTP PRICES

(forecast horizon: 1 month; extended Söderlind and Svensson method)

11 September 1992 a) 18 September 1992 b)

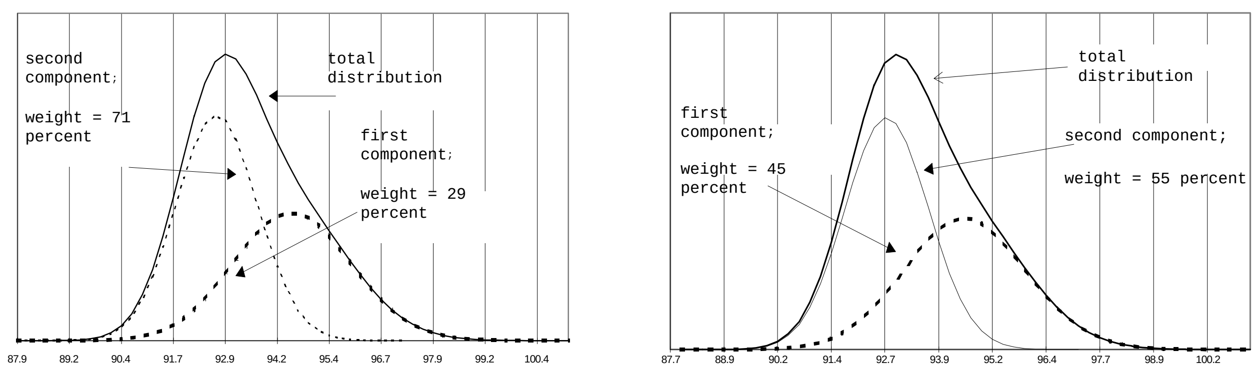

Figure 9

PROBABILITY DISTRIBUTION OF BTP PRICES

(forecast horizon: 1 month; extended Söderlind and Svensson method)

23 October 1992 a)

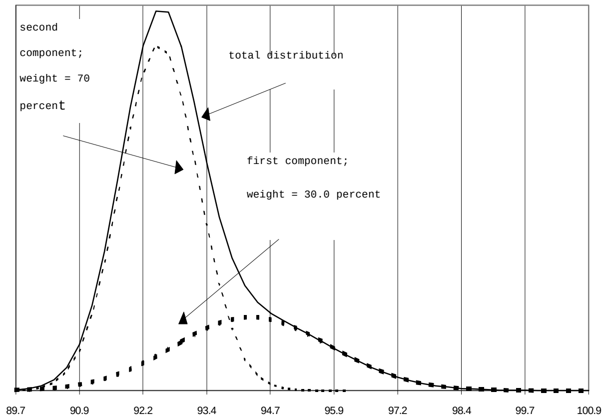

27 October 1992 b)

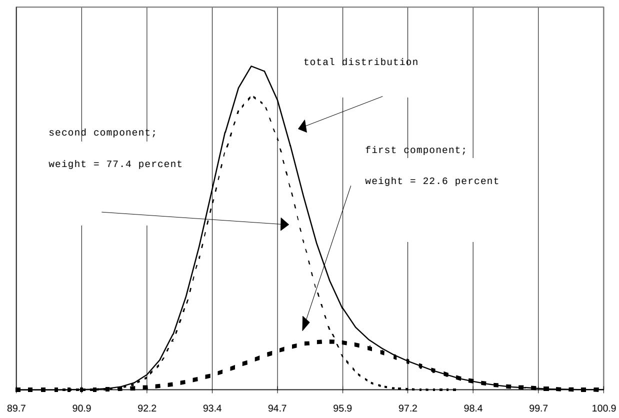

During the ERM crisis in 1992 (Figures 8a-b), i.e. in the week between September 11 and September 18, kurtosis increased by nearly 10 percent (from 4.4 to 4.7). The decline of the probability of observing reductions of the BTP price in the period following the withdrawal from the ERM can be detected by noticing the greater weight of the first component in the overall distribution on September 18 (45 percent compared to 29 percent of September 11) and the corresponding higher values expected for the futures compared to its current level. As a result, the (positive) asymmetry of the overall distribution increased slightly, from 1.31 to 1.38.

4.2.1 *The official rate reductions in 1996-97: evidence from long and short term rate options.*

4.2.1.1 *Options on long-term rates*

This section evidences the impact of recent reductions of the official discount rate (July 23 and October 23, 1996; January 21, 1997) on long term interest rates. In all cases, option prices predicted via the extended model of Söderlind and Svensson fit observed prices fairly well, as is evidenced by the coefficient of determination of the nonlinear regressions according to which the parameters of the pricing functions are identified, which approach unity, and by the standard error of the model, which is reasonably small (between 5 and 15 basis points). The estimated parameters and their significance are reported in Table 5; the estimated probability distributions are drawn in Figures 10a-i.

Table 4

OPTIONS ON BTP FUTURES:
EXTENDED SÖDERLIND AND SVENSSON MODEL (*)
(in brackets Student's-t)

sample → parameters ↓	from 20.8.1992 to 11.9.1992	18.9.1992	from 12.10.1992 to 23.10.1992	from 26.10.1992 to 9.11.1992	from 19.7.1993 to 30.7.1993
σ_1	0.1056 (8.15)	0.2150 (56.23)	0.2021 (158.52)	0.1933 (148.72)	0.0711 (15.26)
σ_2	0.0371 (29.61)	0.0836 (19.09)	0.0439 (281.81)	0.04698 (232.33)	0.0210 (20.59)
a_1	1.5978 (134.72)	1.628 (226.19)	1.2880 (206.55)	1.4157 (237.42)	0.9665 (46.09)
a_2	1.1081 (91.92)	1.2789 (95.48)	0.9917 (481.36)	0.9223 (268.17)	1.1078 (104.08)
α_1	0.2869 (5.50)	0.4512 (94.56)	0.299 (73.99)	0.2266 (57.38)	0.2110 (82.03)
b_1	-0.2838 (-6.34)	0.0194 (47.41)	0.01949 (164.50)	0.01529 (112.90)	0.01141 (86.12)
Characteristics of the mixture of distributions					
σ	2.31	2.27	2.14	2.06	1.47
Skewness	1.31	1.38	1.60	1.45	1.84
Kurtosis	4.37	4.69	3.55	3.09	0.78

sample → parameters ↓	from 4.8.1993 to 16.8.1993	from 4.2.1994 to 17.2.1994	from 19.2.1994 to 4.3.1994
σ_1	0.1763 (50.42)	0.1955 (104.72)	0.1794 (194.82)
σ_2	0.0282 (56.63)	0.02375 (58.30)	0.0386 (512.11)
a_1	1.1840 (40.97)	1.0569 (85.31)	1.4864 (347.11)
a_2	1.0651 (170.11)	1.0523 (128.39)	1.0971 (1213.6)
α_1	0.1912 (47.49)	0.1519 (35.90)	0.2798 (70.69)
b_1	0.0114 (86.12)	0.0085 (40.03)	0.0210 (329.34)
Characteristics of the mixture of distributions			
σ	1.62	1.88	1.99
Skewness	1.93	3.46	4.21
Kurtosis	2.52	1.59	1.97

(*) See equations (19), (21) and (22) for the meaning of the parameters.

The interest rate reductions of October 1996 and especially that of January 1997 appear to have been anticipated by the market, although not fully discounted. Probability distributions before the reduction were significantly skewed towards price increases, consistently with expectations of lower interest rates (Figure 10b-c). In both cases, after the official rate reduction asymmetry and variability declined.

Table 5

PARAMETERS OF THE MODIFIED MODEL OF SVENSSON (*)

(in brackets, Student's-t)

	BTP futures July 1996		BTP futures October 1996		BTP futures January 1997	
	July 22, 1996	July 23, 1996	October 22, 1996	October 23, 1996	January 20, 1997	January 23, 1997
σ_1	0.042	0.028	0.012	0.054	0.073	0.064
	(28.0)	(5.54)	(5.53)	(77.31)	(29.37)	(70.60)
σ_2	0.016	0.017	0.021	0.023	0.018	0.023
	(29.1)	(20.55)	(81.32)	(79.88)	(77.68)	(47.82)
a_1	1.22	1.08	1.71	1.15	1.41	1.13
	(46.3)	(14.28)	(21.05)	(285.6)	(78.16)	(299.66)
a_2	1.00	1.00	0.83	0.97	1.09	1.01
	(72.7)	(75.98)	(91.25)	(162.0)	(108.6)	(132.68)
α_1	0.713	0.962	0.127	0.217	0.276	0.203
	(7.53)	(9.35)	(10.48)	(108.1)	(40.09)	(80.49)
b_1	0.005	0.003	0.029	0.015	0.003	0.013
	(6.19)	(9.12)	(15.68)	(82.91)	(26.35)	(46.60)

(*) See equations (19), (21) and (22) for the meaning of the parameters.

In the case of the July 1996 reduction it is not easy to interpret market expectations. Long rate pdf_s estimated before and after the rate cut were symmetric and did not change significantly (Figure 10a), a circumstance which may be related to the anomalous increase, after the discount rate cut, in short term market rates, in contrast with the direction of the policy move. This behaviour of short term rates may have made it more difficult for economic agents to anticipate the future path of interest rates, which influenced their expectations.

Figure 10

REDUCTIONS OF THE OFFICIAL DISCOUNT RATE
PROBABILITY DISTRIBUTION OF THE FUTURES ON THE 10-YEAR BTP
(time horizon: 3 months; all the reductions by 75 basis points) (*)

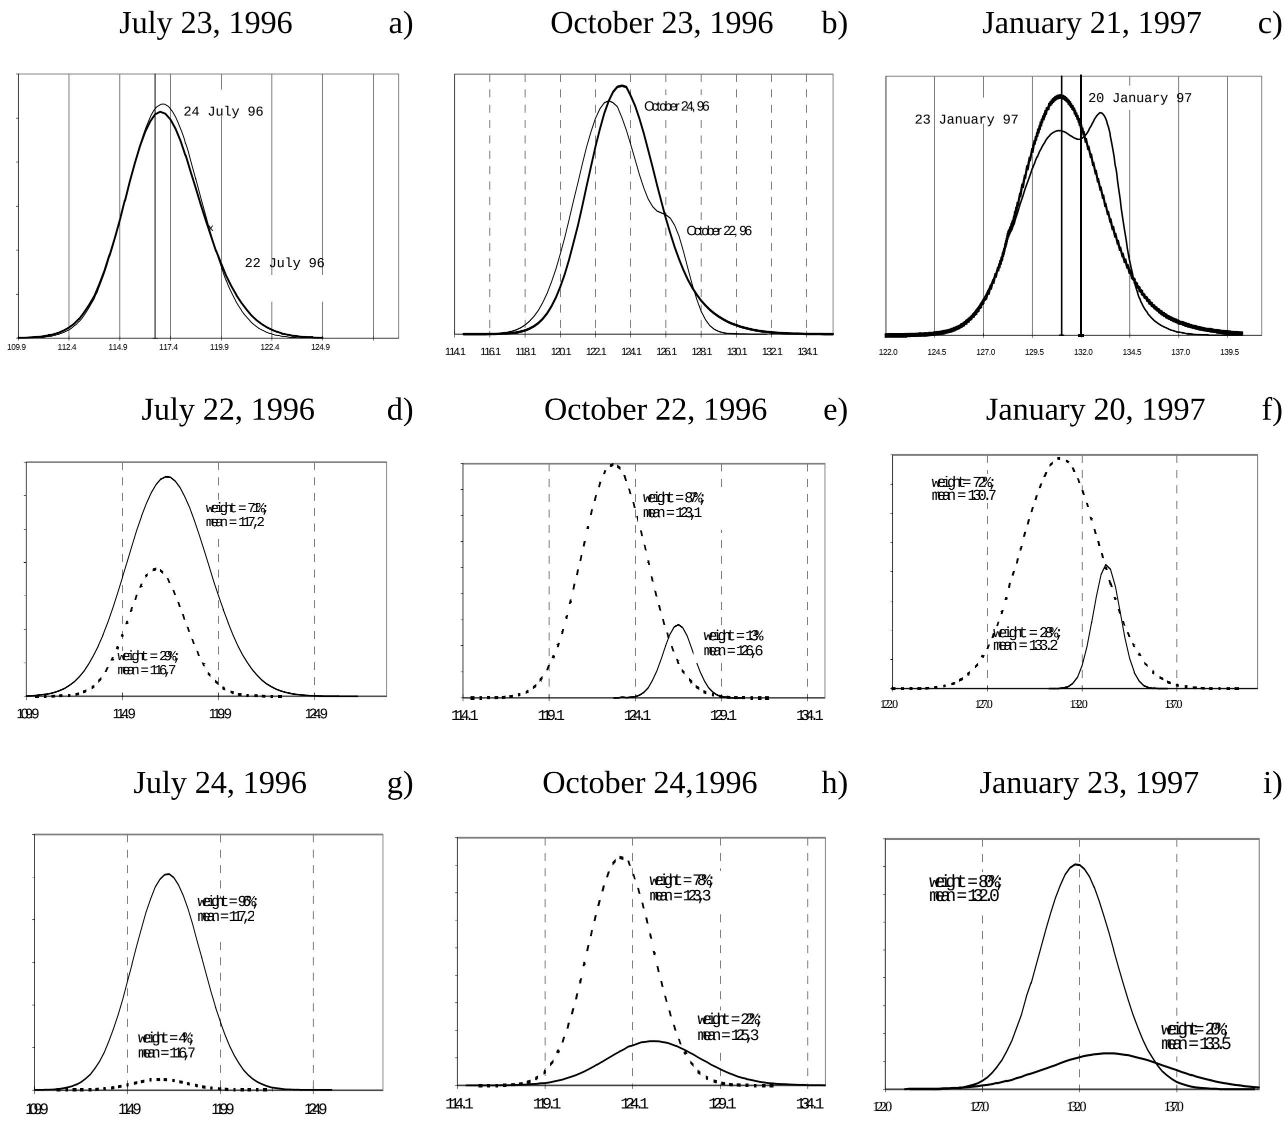

(*) The vertical bar in panels a-c denotes the value of the futures in the specified day.

4.2.1.2 *Options on short-term rates*

The information derived from options written on 3-month eurolira rates appears to be fairly consistent with that obtained from long-term interest rate options. Probability distributions are estimated through the Söderlind and Svensson model adapted for yield option pricing (see footnote 5), for the same official rate reductions as for long term rates; results are reported in Table 6. The official rate cut of July 1996 seems to have been fully anticipated by the market; this evidence appears more clearly from short rate options than from long rate options. The probability distribution obtained as a mixture of the two individual components, evaluated for July 22 (the day before the rate reduction), remains almost unchanged after the policy move; the same happens for the two components (Figure 11a,d,g). A similar pattern emerges for the October 23 rate cut (Figure 11b,e,h). In this case, unlike July, there is a wide difference in the shape of the two components, one of them being almost uniformly distributed in the 2-11 range. This shape reveals persisting uncertainty about the future evolution of short rates, quite independent of the occurrence of the official rate cut. A similar shape is also found in the wake of the January 1997 rate cut, just the day before the easing of monetary policy (Figure 11c,f,i). Contrary to the October rate cut, however, this almost uniformly distributed component of the overall distribution changes after the policy move, in such a way as to reduce the uncertainty surrounding the expected level of the future short rate (in the direction of favouring the likelihood of future reductions).

Table 6

PARAMETERS OF THE MODIFIED MODEL OF SVENSSON (*)

(in brackets, Student's-t)

	Eurolira futures: July 1996		Eurolira futures: October 1996		Eurolira futures: January 1997	
	July 22, 1996	July 23, 1996	October 22, 1996	October 23, 1996	January 20, 1997	January 23, 1997
σ_1	0.296	0.300	100.60	106.26	64.26	0.064
	(36.9)	(10.97)	(5.38)	(4.91)	(3.76)	(70.60)
σ_2	$= \sigma_1$	$= \sigma_1$	$= \sigma_1$	$= \sigma_1$	$= \sigma_1$	$= \sigma_1$
a_1	0.900	0.964	3.57	3.616	3.47	2.99
	(4.06)	(6.45)	(34.61)	(29.32)	(27.10)	(115.26)
a_2	1.61	1.61	2.58	2.632	2.36	2.27
	(18.4)	(58.99)	(25.18)	(22.59)	(9.48)	(54.56)
α_1	0.538	0.526	0.629	0.618	0.75	0.533
	(4.95)	(12.74)	(6.79)	(6.06)	(6.36)	(5.99)
b_1	0.00329	0.0026	-0.0335	-0.03348	-0.0174	-0.301
	(1.96)	(1.97)	(-6.56)	(-5.82)	(-3.26)	(-5.07)

(*) See equations (19), (21) and (22) for the meaning of the parameters.

Figure 11

REDUCTIONS OF THE OFFICIAL DISCOUNT RATE (75 BASIS POINTS)
PDF OF THE FUTURES ON THE 3-MONTH EUROLIRA
(time horizon: 3 months)

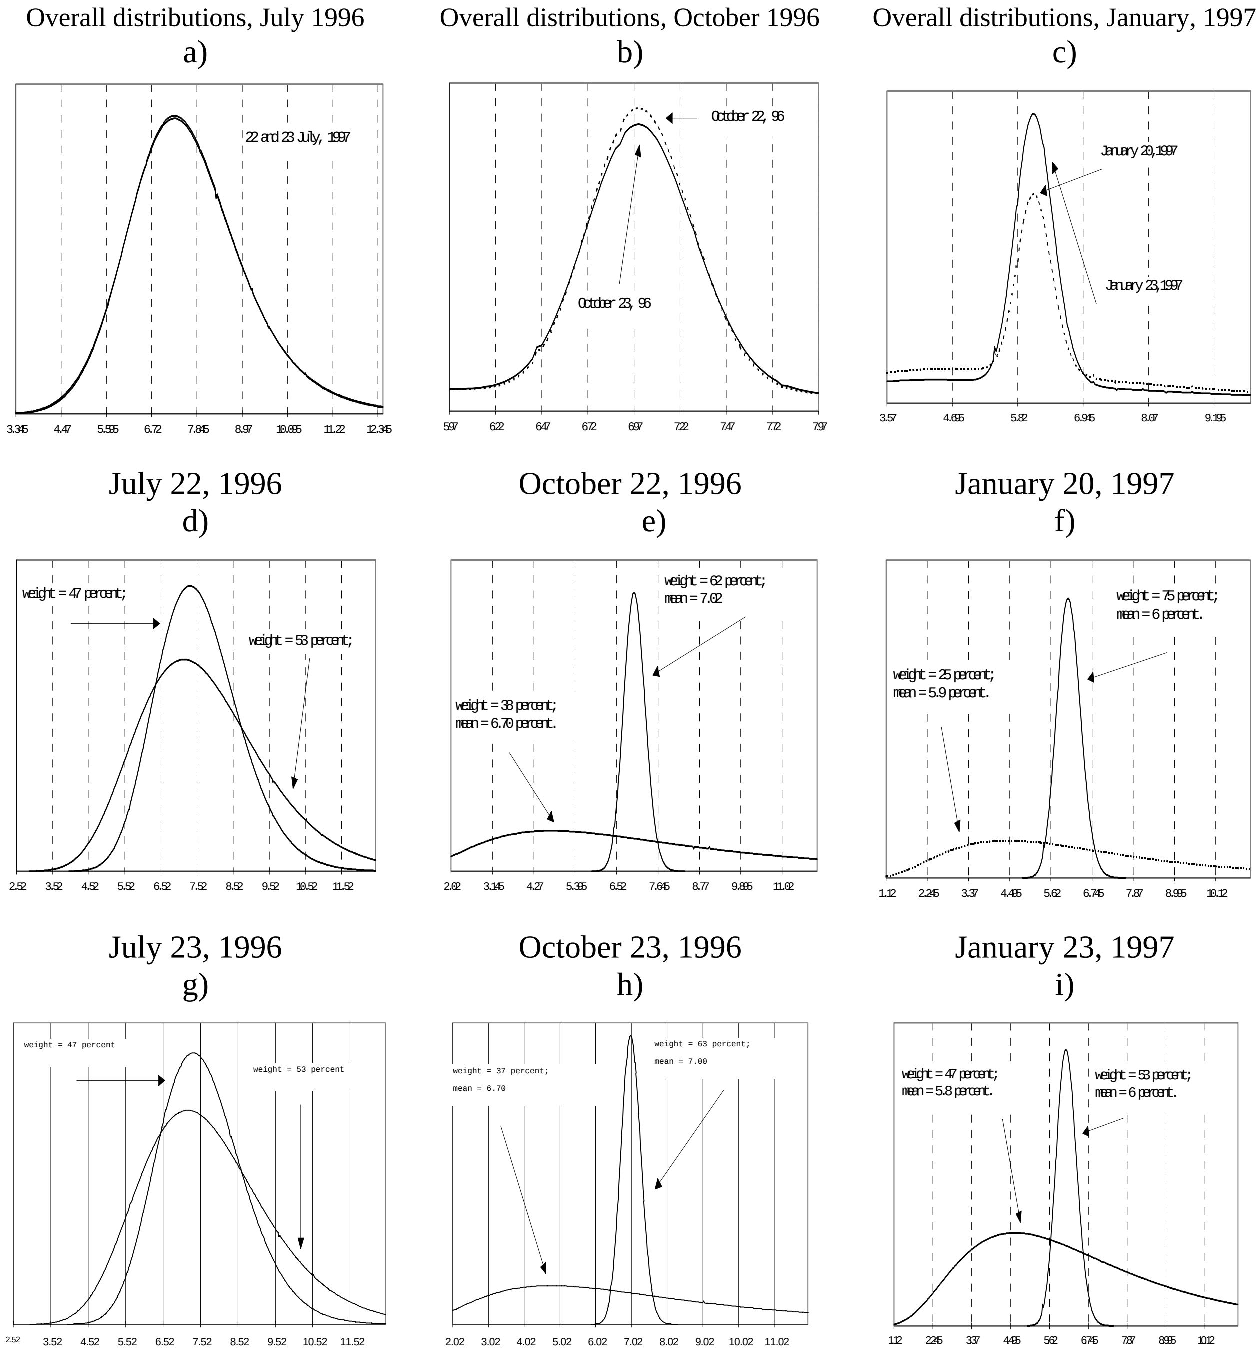

4.3 *Implied volatility functions sensitivity analysis*

As is noted in the Introduction, the extraction of interest rate PDFs from option prices requires the specification of an option pricing formula that can provide, at best, only an approximation of the true price if, as normally happens, the traditional hypotheses are rejected. Here we outline a strategy to assess the validity of the results based on the extended Söderlind and Svensson model. Since volatility has been shown to be the parameter of the option pricing formula to which most of the mispricing can be attributed, we replace Jamshidian's conditional volatility specification with the volatility function specification suggested by Dumas, Fleming and Whaley (1996), reported in (19.1)-(19.2). This exercise is intended to provide a tool to test the sensitivity of the estimated PDFs to the assumption made about the term structure of volatility. As for the Söderlind and Svensson scheme, we obtain the estimates of the coefficients in the model by Dumas, Fleming and Whaley model by maximum likelihood, minimising the squared difference between observed and theoretical option prices. Here, the price of an option, i.e. the one-to-one mapping with its volatility, depends only upon its moneyness. For brevity's sake, comparison of results from the two models has been limited to four dates only; all in September and October 1992, a period of market turbulence when changes in market volatility could have played a prominent role. The results, reported in Table 7, show that the parameters of the volatility function may react significantly as a result of official rate changes; the most sizeable of these variations affecting σ^1_2 and σ^2_2, i.e. the parameters that regulate the quadratic term of equations (19.1) - (19.2). Significant changes also concern b_1, the parameter that determines the dispersion of the mean of the second component of the overall distribution; the relative weight of the two PDFs, which ranges from 8.4 to 39.7 percent in the first two subperiods and from 18.3 to 26.9 in the second, also displays large variability (Figure 12a-d).

The implied $PDFs_s$ based on the implied volatility functions have shapes similar to those derived from the extended Söderlind and Svensson model. Like the latter, they tend to be skewed towards prices higher than the current value of the futures, albeit to a lesser extent. In particular, the value of b_1 implies that the estimated mean of the compounded distribution is closer to the current price of the futures. The most similar shapes for the overall distributions are obtained for the estimates carried out for October 23, 1992, except for the relative weights of the two components (70 and 30 percent in the first case, 82 and 18 in the second). Further analysis may be required by examining more complex functional forms for the implied volatility functions, and by extending this technique to longer time periods.

Table 7

PARAMETERS OF THE IMPLIED VOLATILITY FUNCTIONS (*)

(t-ratios in parentheses)

	from 20.8.1992 to 11.9.1992	18.9.1992	23.10.1992	from 26.10.1992 to 27.10.1992
σ^1_0	0.0073	0.0081	-0.0036	0.0081
	(114.51)	(62.45)	(-2.60)	(43.22)
σ^2_0	0.0078	0.0084	0.00267	0.0144
	(186.56)	(65.91)	(11.74)	(12.93)
σ^1_1	-0.3618	-0.3003	-0.1604	-0.1909
	(-69.70)	(-18.10)	(-1.72)	(-7.87)
σ^2_1	-0.2728	-0.2667	-0.0765	-0.6614
	(-146.10)	(-22.73)	(-3.48)	(-6.98)
σ^1_2	-35.27	-30.75	73.221	-42.062
	(-90.00)	(-14.92)	(8.02)	(-11.10)
σ^2_2	-22.46	-27.67	0.0925	-77.178
	(-59.76)	(-26.12)	(.02)	(-10.37)
α_1	0.0839	0.3968	0.1825	0.2691
	(5.67)	(18.98)	(2.51)	(1.97)
b_1	0.00504	0.01547	0.0268	0.00299
	(9.00)	(16.92)	(3.18)	(2.48)

(*) See equations (19), (29.1) and (29.2) for the meaning of the parameters.

Figure 12

PROBABILITY DISTRIBUTION OF BTP PRICES
(forecast horizon: 1 month; Dumas, Fleming and Whaley's volatility)

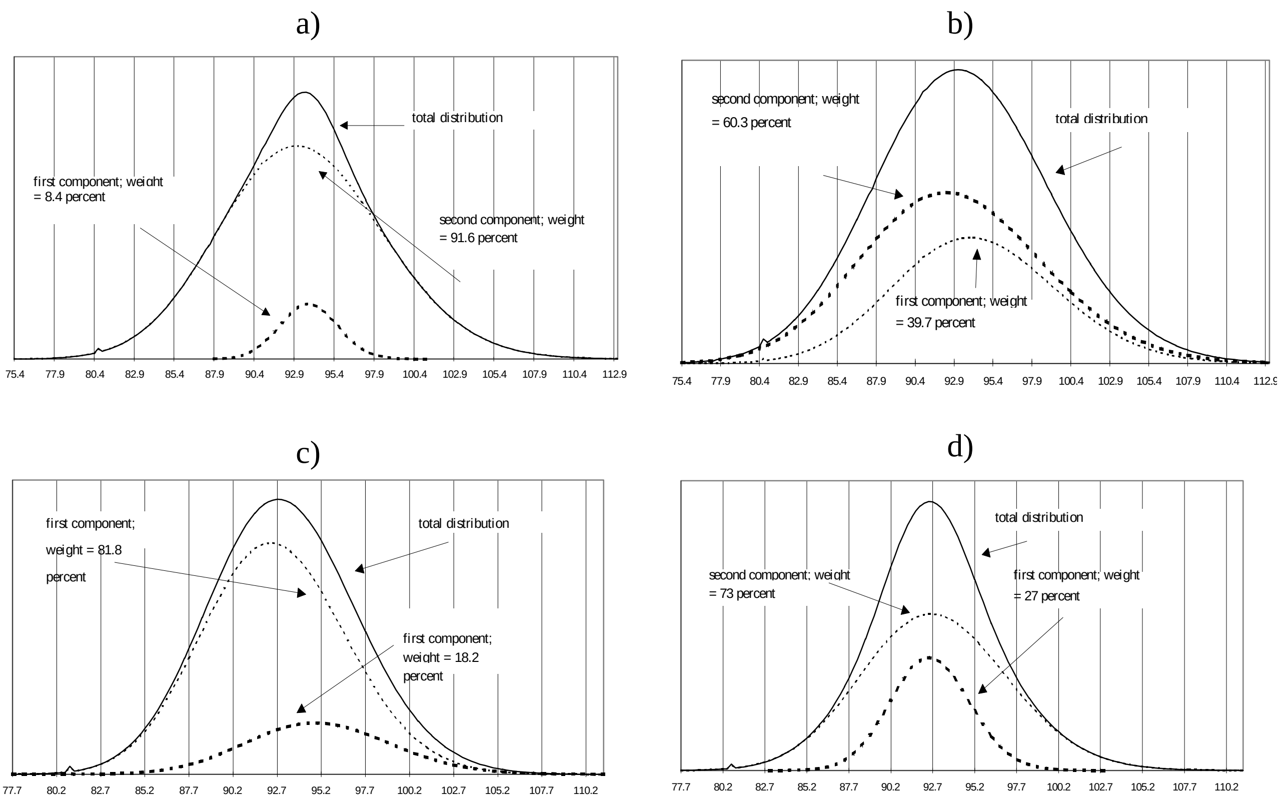

5. Conclusions

This paper presents and applies several methods for extracting implied probability distributions of asset prices from the prices of traded options. The empirical shortcomings of BS framework calls for substantial modifications of its underlying assumptions, above all the constant variance hypothesis. We show how it is possible to derive the implied PDF of an asset by simulation schemes, when the traditional hypotheses are rejected; the unavailability of closed-form solutions in this case requires the adoption of Monte Carlo simulations. From a more practical viewpoint we have obtained the PDFs of long and short term asset prices from two versions of an option pricing schemes, which enrich the Söderlind and Svensson model, extending it to more general specifications for the volatility of asset prices. In one case the asset price dynamics are assumed to conform to the hypothesis of Jamshidian (1989), where the volatility of the asset underlying the option contract is driven by a mean-reverting stochastic process. This volatility specification is then compared to a more general model, suggested by Dumas, Fleming and Whaley (1996) and based on a broader parametric specification of the implied volatility function, where the underlying asset variability evolves according to a nonlinear function of observable features of the options.

Empirical results provide evidence that the PDFs implicit in the prices of options tend to anticipate the future trend of the underlying asset price; hence they can provide valuable information to policy makers about asset price reactions to changes in monetary stance, as perceived by economic agents. Our analysis, though carried out for a limited number of periods and events, has shown that notable changes may occur in the shape of probability distributions for long and short rates, shedding light on the shifts that occur in the expectations of market participants in reaction to moves in the monetary policy stance.

The period that led to the withdrawal of the Italian lira from the ERM is shown to be a clear case study for analysing rapid shifts in expectations: it was indeed characterised by a sharp reduction of the probability of asset price stability and an increase of the likelihood of sharp falls. The distribution estimated for August 17, 1992, displaying a gaussian shape, had completely changed by mid-September, when the Lira left the ERM, driving the probability of observing bond prices above the level of 93 to nil.

The sequence of three recent official rate cuts in Italy (July 23, 1996; October 23, 1996; January 21, 1997) was a useful benchmark to test the ability of our methodology of extraction the implied PDFs from option prices. Evidence based on the estimated distributions for short and long rates suggest that the first cut in the discount rate (23 July) was partially anticipated by the market. In the two later episodes, however, short and long rates point to slightly different conclusions. The distributions observed on these dates for the long rate were skewed towards an increase of the *futures* price before the official rate cut (e.g. consistent with an expectation of lower rates); the mean of the distribution was also higher than the value of the futures price then prevailing, whereas the probability distribution of short rates displayed a lesser degree of asymmetry. After the monetary easing, little reaction took place in the short rate distribution, as if the expectation of a rate change had been fulfilled, whereas for long rates, probably as a result of reduced asymmetry as well as variability, a significant adjustment in the shape of the distribution was observed.

References

Aït-Sahalia, Y., Y. Wang and F. Yared (1998), "Do Option Markets Correctly Price the Probabilities of Movements of the Underlying Asset?", University of Chicago, mimeo.

Backus, D. and S. Zin (1994), "Reverse Engineering the Yield Curve", NBER Working Paper, No. 4676.

Bahra, B. (1996), "Probability Distributions of Future Asset Prices Implied by Option Prices", *Bank of England Quarterly Bullettin*, August, pp. 299-311.

Black, F. (1976), "Studies of Stock Price Volatility Changes", Proceedings of the American Statistical Association, Business and Economic Statistics Section, pp. 177-81.

Black, F. and M. Scholes (1972), "The Valuation of Option Contracts and a Test of Market Efficiency", *Journal of Finance*, Vol. 27, pp. 399-417.

Black, F. and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, Vol. 81, pp. 637-59.

Bossaert, P. and P. Hillion (1997), "Local Parametric Analysis of Hedging in Discrete Time", *Journal of Econometrics*, Vol. 81, pp. 243-72.

Breeden, D. and R. Litzenberger (1978), "Price of State Contingent Claims Implicit in Option Prices", *Journal of Business*, Vol. 51, pp. 621-51.

Butler, C. and H. Davis (1998), "Assessing Market Views on Monetary Policy: The Use of Implied Risk Neutral Probability Distributions", paper presented at the BIS-CEPR Conference on Asset Prices and Monetary Policy, 26-27 January 1998.

Campa, J. M. and K. Chang (1995), "Testing the Expectations Hypothesis on the Term Structure of Volatilities in Foreign Exchange Options", *Journal of Finance*, Vol. 50, pp. 529-47.

Derman, E., I. Kani and J. Z. Zhou (1996), "The Local Volatility Surface: Unlocking the Information in Index Option Prices", *Financial Analyst Journal*, July-August, pp. 25-36.

Ding, Z., C. Granger and R. Engle (1993), "A Long Memory Property of Stock Market Returns and a New Model", *Journal of Empirical Finance*, Vol. 1, pp. 83-106.

Duan, J. C. (1995), "The GARCH Option Pricing Model", *Mathematical Finance*, Vol. 5, pp. 13-32.

Dumas, B., J. Fleming and R. Whaley (1996), "Implied Volatility Functions: Empirical Tests", CEPR Discussion Paper, No. 1369.

Engle, R. and C. Mustafa (1992), “Implied ARCH Models from Option Prices”, *Journal of Econometrics*, Vol. 52, pp. 289-311.

Fornari, F. and A. Mele (1997a), “Weak Convergence and Distributional Assumptions for a General Class of Nonlinear Arch Models”, *Econometric Reviews*, Vol. 16, pp. 205-29.

Fornari, F. and A. Mele (1997b), “Pricing Options under Non Standard Assumptions”, paper presented at the ESEM98 meeting in Berlin (available on request from the first named author).

Fornari, F. and C. Monticelli (1998), “The Information Content of Interest Rate Options: Evidence from LIFFE”, *Applied Financial Economics*, forthcoming.

Heston, S. (1993), “A Closed Form Solution for Options with Stochastic Volatility with Application to Bond and Currency Options”, *Review of Financial Studies*, Vol. 6, pp. 327-44.

Hull, J. (1988), *Options, Futures and Other Derivative Securities*, Englewood Cliffs, Prentice-Hall.

Hull, J. and A. White (1987), “The Pricing of Options on Assets with Stochastic Volatility”, *Journal of Finance*, Vol. 42, pp. 281-300.

Jamshidian, F. (1989), “An Exact Bond Option Formula”, *Journal of Finance*, Vol. 44, pp. 205-9.

Lamoureux, C. and W. Lastrapes (1993), “Forecasting Stock-Return Variance: Towards an Understanding of Stochastic Implied Volatilities”, *Review of Financial Studies*, Vol. 6, pp. 293-326.

Longstaff, F. (1990), “The Valuation of Options on Yields”, *Journal of Financial Economics*, Vol. 26, pp. 97-121.

Longstaff, F. and E. Schwartz (1992), “Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model”, *Journal of Finance*, Vol. 47, pp. 1259-82.

Malz, A. (1996), “Using Option Prices to Estimate Realignment Probabilities in the European Monetary System: The case of Sterling-Mark”, *Journal of International Money and Finance*, Vol. 15, pp. 717-48.

Melick, W. R. and C. P. Thomas (1997), “Recovering an Asset’s Implied PDF from Option Prices: An Application to Crude Oil During the Gulf Crisis”, *Journal of Financial and Quantitative Analysis*, Vol. 32, pp. 91-115.

Merton, R. C. (1973), “The Theory of Rational Option Pricing”, *Bell Journal of Economics and Management Science*, Vol. 4, pp. 141-83.

Merton, R. C. (1976), “Option Pricing When Underlying Stock Returns Are Discontinuous”, *Journal of Financial Economics*, Vol. 3, pp. 125-44.

Nelson, D. (1990), “ARCH Models as Diffusion Approximation”, *Journal of Econometrics*, Vol. 45, pp. 7-38.

Neuhaus, H. (1995), “The Information Content of Derivatives for Monetary Policy”, Deutsche Bundesbank, Discussion Paper, No. 3.

Rubinstein, M. (1976), “The Valuation of Uncertain Income Streams and the Pricing of Options”, *Bell Journal of Economics and Management Science*, Vol. 7, pp. 407-25.

Söderlind, P. and L. Svensson (1997), “New Techniques to Extract Market Expectations from Financial Instruments”, *Journal of Monetary Economics*, Vol. 40, pp. 383-429.

Stapleton, R. C. and M. G. Subrahmanyam (1993), “The Analysis and Valuation of Interest Rate Options”, *Journal of Banking and Finance*, Vol. 17, pp. 1079-95.

RECENTLY PUBLISHED “TEMI” (*)

(*) Requests for copies should be sent to:
Banca d’Italia – Servizio Studi – Divisione Biblioteca e pubblicazioni – Via Nazionale, 91 – 00184 Rome (fax 39 6 47922059)