Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Journal of Financial Stability 2 (2006) 28–54
Journal of FINANCIAL STABILITY
www.elsevier.com/locate/jfstabil

Assessing central bank credibility during the ERM crises: Comparing option and spot market-based forecasts☆

Markus Haas [a,*], Stefan Mittnik [b], Bruce Mizrach [c]

[a] *Institute of Statistics, University of Munich, Germany*
[b] *University of Munich, Ifo Institute for Economic Research, Center for Financial Studies, Frankfurt*
[c] *Department of Economics, Rutgers University, USA*

Received 4 January 2004; accepted 1 February 2005
Available online 24 February 2006

Abstract

Financial markets embed expectations of central bank policy into asset prices. This paper compares two approaches that extract a probability density of market beliefs. The first is a simulated moments estimator for option volatilities described in [Mizrach, B., 2002. When Did the Smart Money in Enron Lose Its' Smirk? Rutgers University Working Paper #2002-24]; the second is a new approach developed by [Haas, M., Mittnik, S., Paolella, M.S., 2004a. Mixed normal conditional heteroskedasticity, J. Financial Econ. 2, 211–250] for fat-tailed conditionally heteroskedastic time series. In an application to the 1992–1993 European Exchange Rate Mechanism crises, we find that both the options and the underlying exchange rates provide useful information for policy makers.

JEL classification: G12; G14; F31

Keywords: Options; Implied probability densities; GARCH; Fat-tails; Exchange Rate Mechanism

☆ Direct editorial correspondence to the third author at e-mail: mizrach@econ.rutgers.edu. Tel. +1 732 932 8261; fax: +1 732 932 7416, http://www.snde.rutgers.edu. This paper was prepared for the 2004 JFS-Carli Bank of Italy Conference on Derivatives and Financial Stability.

* Corresponding author.
E-mail addresses: haas@stat.uni-muenchen.de (M. Haas), finmetrics@stat.uni-muenchen.de (S. Mittnik).

doi:10.1016/j.jfs.2005.02.012

1. Introduction

A basic insight of financial economics is that asset prices should reflect views about the future. For this reason, many economists rely on market prices to make predictions. Even when these views are incorrect, policy makers may want to avoid changes that the market is not expecting.

In recent years, some novel techniques have been introduced to extract market expectations. This paper explores two of them: extracting implied probability densities from option prices and volatility modeling of the underlying. Both methods have the advantage of producing predictive densities rather than just point forecasts. These tools can, in principal, allow central bankers to examine the full range of risks facing their economies.

There are numerous approaches that generalize the Black–Scholes model. Merton (1976) and Bates (1991) allow sudden changes in the level of asset prices. Wiggins (1987), Hull and White (1987), Stein and Stein (1991) and Heston (1993) allow volatility to change over time. A related literature, with papers by Dumas et al. (1998) and Das and Sundaram (1999), has looked at deterministic variations in volatility with the level of the stock price or with time.

To extract market expectations of the exchange rate, we utilize a method first used in Mizrach (2002) that looks directly at the probability distribution. We parameterize the exchange rate process as a mixture of log normals, as in Ritchey (1990) and Melick and Thomas (1997), and fit the model to options prices. In an application to the Enron bankruptcy, Mizrach found that investors were far too optimistic about Enron until days before the stock's collapse.

Our second approach tries to extract information directly from the underlying currencies. We utilize a general mixture of two normal densities to extract information from the spot foreign exchange market. In this model, both the mixing weights as well as the parameters of the component densities, i.e., component means and variances, are time-varying and may depend on past exchange rates as well as further explanatory variables, such as interest rates. The dynamic mixture model we specify is a combination of the logistic autoregressive mixture with exogenous variables, or LMARX, model investigated in Wong and Li (2001) and the mixed normal GARCH process recently proposed by Haas et al. (2004a). The predictive densities generated from the resulting LMARX–GARCH model exhibit an enormous flexibility, and they may be multimodal, for example, in times where a realignment becomes more probable.

In this paper, we utilize the two approaches to explore market sentiment prior to the exchange rate crises of September 1992 and July–August 1993. In the first episode, the British Pound (BP) and Italian Lira withdrew from the Exchange Rate Mechanism (ERM) of the European Monetary System (EMS). The Pound had traded in a narrow range against the German Deutsche Mark (DM) for almost two years and the Lira for more than five. The crisis threw the entire plan for European economic and financial integration into turmoil. The French Franc (FF) remained in the mechanism, but speculative pressures against it remained strong. In the second crisis we examine, the Franc, in August 1993, had to abandon its very close link with the DM (the "Franc fort") and widen it's fluctuation band.

Campa and Chang (1996) have looked at ERM credibility using arbitrage bounds on option prices. They find that option prices reflected the declining credibility of the Lira and Pound in 1992 and the Franc in 1993. Malz (1996) finds an increasing risk of BP devaluation starting in late August 1992. Christoffersen and Mazzotta (2004) find useful predictive information in 10 European countries' over-the-counter currency options.

We first examine the options markets' implied probability of depreciation in the FF and BP prior to the ERM crises. The model estimates reveal that the market anticipated both events. The devaluation risk with the Franc rises significantly 11 days in advance of the crisis. With the Pound, the risk is subdued until only five days before it devalued on "Black Wednesday" September 16, 1992.

Vlaar and Palm (1993) were the first to use the normal mixture density to model EMS exchange rates against the DM, noting that, in contrast to freely floating currencies, these often show pronounced skewness, due to jumps which occur in case of realignments, but also, for example, as a result of expected policy changes or speculative attacks. Although Vlaar and Palm (1993) noted that making the jump probability a function of explanatory variables, such as inflation and interest rates, may be a promising task, they did not undertake such analysis.

Neely (1994) surveys research on forecasting realignments in the EMS and reports evidence for realignments to be predictable to some extent from information such as interest rates and the position of the exchange rate within the band. Building both on the results surveyed in Neely (1994) and the work of Vlaar and Palm (1993) and Palm and Vlaar (1997), the studies of Bekaert and Gray (1998), Neely (1999) and Klaster and Knot (2002) use more general dynamic mixture models of exchange rates in target zones. Thus, the model employed below has some similarities with those developed in these studies, as will be discussed below.

The dynamic mixture model provides, as in the options-based approach, estimates of the probability of a depreciation. For the FF, the model indicates a considerable increase of this probability one week in advance of the crisis, and a further increase immediately before the de facto devaluation of the FF, when the bands of the target zone were widened to $\pm 15\%$.

For the BP, we can, in contrast to the options-based approach, not develop a promising dynamic mixture model, because the BP joined the ERM only in October 1990 and withdrew in September 1992. During this period there were no realignments or large jumps within the band, so that the sample does not provide information that is necessary to fit a target zone mixture model. Consequently, the mixed normal GARCH model detects a rise in the devaluation probability only *after* the Pound was withdrawn from the ERM.

Both models provide a complete predictive density for the exchange rate, and the last part of the paper examines the fit of the entire density. We utilize the approach of Berkowitz (2001) to formally compare the model's density-forecasting performance. In the options market, the predictive density becomes indistinguishable from the post crisis density on July 21 for the FF, 11 days before the crisis. For the BP, there are some early warning signals in mid-August and the beginning of September. In the FF spot market, the predictive density is consistent with the post-crisis data from the outset. For the BP, the result is similar to the options. There are some brief early signals, but the densities statistically differ from the post-devalation period until September 10th.

The paper continues with some discussion of the ERM. Section 3 describes the theory of implied density extraction from options. It also proposes a mixture of log normals specification which nests the Black–Scholes model. We also develop a GARCH mixture model for the spot exchange rate. Section 4 contains some stylized features of the currency options, and some detailed issues in estimation for both models. From the two sets of parameter estimates, we compute implied devaluation probabilities. Section 5 compares the entire predictive density statistically. Section 6 concludes with directions for future research.

2. The ERM

The ERM began in 1979 with seven member countries.[1] The mechanism included a grid of fixed exchange rates with European Currency Unit (ECU) central parities and fluctuation

[1] Belgium, Denmark, France, Germany, Italy, Ireland, and The Netherlands.

bands. Prior to the crises, the FF had a target zone of ±2.25% and the BP ±6%. Maintaining the parities required policy coordination with the German Bundesbank, and when necessary, intervention.

By the Spring of 1992, the momenta towards a single European currency seemed irreversible. Spain had joined the ERM in June of 1989. Great Britain finally overcame its resistance in October 1990. Portugal joined in April 1992 bringing the total membership to 10. In addition, Finland and Sweden had been following indicative DM targets. All the major European currencies, save the Swiss Franc, were incorporated in a system of apparently stable exchange rate bands. Almost five years had passed without devaluations.[2] The financial sector seemed poised for monetary union, the next logical step in the blueprint of the Maastricht treaty signed on December 10, 1991.

A swift sequence of events left the idea of currency union almost irretrievably damaged. The Danes rejected the Maastricht treaty in June of 1992. The Finnish Markaa and the Swedish Krona faced devaluation pressures in August which the Bank of Finland and the Swedish Riksbank actively resisted. The Markaa was allowed to float on September 8, and it quickly devalued 15% against the DM. The Riksbank raised their marginal lending rate to 500% on September 16.

Then some of the core ERM currencies came under speculative attack. The Bank of England briefly raised their base lending rates, but the British chose to withdraw from the ERM on September 16 rather than expending additional reserves.[3] The Lira devalued by 7% on September 13 and withdrew from the mechanism on September 17.

A number of additional devaluations followed. The Krona was allowed to float on November 19. The Spanish Peseta (in September and November 1992), the Portuguese Escudo (in November 1992), and then the Irish Punt (in February 1993) subsequently adopted new parities. The ERM remained in turmoil into the summer. France faced continued pressure and went through a de facto devaluation when the ERM bands were widened to ±15% on August 2, 1993.

In retrospect, the origins of these crises were evident. The Finnish and Swedish economies were weakened by recession and the collapse of the Soviet Union. Britain had probably overvalued the Pound when it entered the ERM. The Lira had appreciated 30% in real terms against the DM since 1987. Germany had raised interest rates to fight off inflationary pressures from unification, weakening the entire European economy in the process.

The folklore of this period suggests that some market participants anticipated the crisis, and may even have precipitated it. The hedge fund trader George Soros is rumored to have made some US$ 1 billion speculating against the Pound and the Lira in 1992.

The question we ask here is how well diffused was this information. Did either the options market or spot market anticipate these events and can our models extract these expectations?

3. Models for currency options and the spot rate

3.1. Implied probability densities from options

The basic option pricing framework builds upon the Black–Scholes assumption that the underlying asset is log normally distributed. Let $f(S_T)$ denote the terminal risk neutral probability at time T, and let $F(S_T)$ denote the cumulative probability . A European call option at time t,

[2] There was a small devaluation of the Italian Lira when it moved to narrow bands in January 1990.

[3] The Bundesbank is reported to have spent DM92 billion defending the Pound and Lira during this crisis.

expiring at T, with strike price K, is priced

$$C(K,\tau) = \mathrm{e}^{-i_{\mathrm{d}}\tau} \int_K^{\infty} (S_T - K) f(S_T)\,\mathrm{d}S_T, \tag{1}$$

where $\tau = T - t$, and i_{d} and i_{f} are the annualized domestic and foreign risk-free interest rates. In the case where $f(\cdot)$ is the log–normal density and volatility σ is constant with respect to K, this yields the Black–Scholes formula,

$$\mathrm{BS}(S_t, K, \tau, i_{\mathrm{f}}, i_{\mathrm{d}}, \sigma) = S_t \mathrm{e}^{-i_{\mathrm{f}}\tau}\Phi(d_1) - K\mathrm{e}^{-i_{\mathrm{d}}\tau}\Phi(d_2), \tag{2}$$

$$d_1 = \frac{\ln(S_T/K) + (i_{\mathrm{d}} - i_{\mathrm{f}} + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \quad d_2 = d_1 - \sigma\sqrt{\tau},$$

where $\Phi(\cdot)$ denotes the cumulative standard normal distribution. In this benchmark case, implied volatility is a sufficient statistic for the entire implied probability density which is centered at the risk-free interest differential $i_{\mathrm{d}} - i_{\mathrm{f}}$.

Mizrach (2002) surveys an extensive literature and finds that option prices in a variety of markets appear to be inconsistent with the Black–Scholes assumptions. In particular, volatility seems to vary across strike prices – often with a parabolic shape called the volatility "smile". The smile is often present on only one part of the distribution giving rise to a "smirk".

3.1.1. How volatility varies with the strike

Under basic no-arbitrage restrictions, we can consider more general densities than the log–normal for the underlying. Breeden and Litzenberger (1978) show that the first derivative is a function of the cumulative distribution,

$$\left.\frac{\partial C}{\partial K}\right|_{K=S_T} = -\mathrm{e}^{-i_{\mathrm{d}}\tau}(1 - F(S_T)). \tag{3}$$

The second derivative then extracts the density,

$$\left.\frac{\partial^2 C}{\partial K^2}\right|_{K=S_T} = \mathrm{e}^{-i_{\mathrm{d}}\tau} f(S_T). \tag{4}$$

The principal problem in estimating f is that we do not observe a continuous function of option prices and strikes. Early attempts in the literature, like Shimko (1993), simply interpolated between option prices. Bliss and Panigirtzoglou (2002) find that implied volatility functions fit well when the strikes are dense, but as Mizrach (2002) notes, this often leads to arbitrage violations in the tails. Later attempts turned to either specifying a density family for f or a more general stochastic process for the spot price. Dupire (1994) shows that both approaches are equivalent; for driftless diffusions, there is a unique stochastic process corresponding to a given implied probability density. This paper follows Ritchey (1990) and Melick and Thomas (1997) by specifying f as a mixture of log–normal distributions. The advantage of this specification is that the option prices are just probability weighted averages of the Black–Scholes prices for each mixture component.

3.1.2. Mixture-of-log–normals specification

We assume that the stock price process is a draw from a mixture of three (non-standard) normal distributions, $\Phi(\mu_j, \sigma_j)$, $j = 1, 2, 3$, with $\mu_3 \geq \mu_2 \geq \mu_1$. Three additional parameters λ_1, λ_2 and λ_3 define the probabilities of drawing from each normal. To nest the Black–Scholes, we restrict the mean to equal the interest differential, $\mu_2 = i_\mathrm{d} - i_\mathrm{f}$. Risk neutral pricing then implies restrictions on either the other means or the probabilities. We chose to let μ_1, λ_1 and λ_3 vary, which implies

$$\mu_3 = \frac{\mu_1\lambda_1}{\lambda_{3,}}, \tag{5}$$

and

$$\lambda_2 = 1 - \lambda_1 - \lambda_3. \tag{6}$$

For estimation purposes, this leaves six free parameters $\theta = (\theta_1, \theta_2, \ldots, \theta_6)$. We take exponentials of all the parameters because they are constrained to be positive. The left-hand mixture is given by

$$\Phi(\mu_1, \sigma_1) = \Phi(i_\mathrm{d} - i_\mathrm{f} - \mathrm{e}^{\theta_1}, 100 \times \mathrm{e}^{\theta_2}). \tag{7}$$

The only free parameter of the middle normal density is the standard deviation,

$$\Phi(\mu_2, \sigma_2) = \Phi(i_\mathrm{d} - i_\mathrm{f}, 100 \times \mathrm{e}^{\theta_3}). \tag{8}$$

We use the logistic function for the probabilities to bound them on [0, 1],

$$\lambda_1 = \frac{\mathrm{e}^{\theta_4}}{1 + \mathrm{e}^{\theta_4}}, \tag{9}$$

$$\lambda_3 = \frac{\mathrm{e}^{\theta_5}}{1 + \mathrm{e}^{\theta_5}}. \tag{10}$$

The probability specification implies the following mean restrictions on the third normal,

$$\Phi(\mu_3, \sigma_3) = \Phi\left((i_\mathrm{d} - i_\mathrm{f} + \mathrm{e}^{\theta_1}) \times \frac{\mathrm{e}^{\theta_4}/(1 + \mathrm{e}^{\theta_4})}{\mathrm{e}^{\theta_5}/(1 + \mathrm{e}^{\theta_5})}, 100 \times \mathrm{e}^{\theta_6}\right). \tag{11}$$

Mizrach (2002) shows that this data generating mechanism can match a wide range of shapes for the volatility smile.

3.2. GARCH mixture model for the spot exchange rate

The mixed normal GARCH process is the building block of our models for the spot rate.[4] It was recently proposed by Haas et al. (2004a) and generalizes the classic normal GARCH model of Bollerslev (1986) to the mixture setting. The percentage change of the log-exchange rate, $r_t = 100 \times \log(S_t/S_{t-1})$, where S_t is the exchange rate at time t, is said to follow a k-component mixed normal (MN) GARCH(p, q) process if the conditional distribution of r_t is a k-component MN, that is,

$$r_t|\Psi_{t-1} \approx \mathrm{MN}(\lambda_{1,t}, \ldots, \lambda_{k,t}, \mu_{1,t}, \ldots, \mu_{k,t}, \sigma^2_{1,t}, \ldots, \sigma^2_{k,t}), \tag{12}$$

where Ψ_t is the information at time t, and the mixing weights sastisfy $\lambda_j \in (0, 1)$, $j = 1, \ldots, k$, and $\sum_j \lambda_j = 1$. The $k \times 1$ vector of component variances, denoted by $\sigma_t^{(2)} = [\sigma^2_{1,t}, \ldots, \sigma^2_{k,t}]'$,

[4] For an application of a related model class, the Markov-switching GARCH model, to predicting exchange rate densities, see Haas et al. (2004b).

evolves according to

$$\sigma_t^{(2)} = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^{(2)}, \tag{13}$$

where α_0 is a positive $k \times 1$ vector; $\alpha_i, i = 1, \ldots, q$, are nonnegative $k \times 1$ vectors; and β_i, $i = 1, \ldots, p$, are nonnegative $k \times k$ matrices, and

$$\epsilon_t = r_t - E(r_t|\Psi_{t-1}) = r_t - \sum_{j=1}^{k} \lambda_{j,t}\mu_{j,t}. \tag{14}$$

Haas et al. (2004a) considered the case where the mixing weights, $\lambda_{j,t}$, and the component means, $\mu_{j,t}$, $j = 1, \ldots, k$, are constant over time, but the generalization considered in (12)–(14), with these quantities being time-varying, is straightforward conceptually. The mixing weights and the component means may depend both on lagged values of r_t and on further explanatory variables, as in the LMARX model of Wong and Li (2001). Thus, the dynamic mixture model employed in the present paper is a combination of the MN–GARCH and the LMARX models, which will be termed LMARX–GARCH.

As with the classic GARCH model, the MN–GARCH(1,1) specification will usually be sufficient, and in most applications it will be reasonable to impose certain restrictions on the α_i's and β_i's in (13). However, the general formulation will be useful in discussing different versions of the MN–GARCH process corresponding to different restrictions imposed on the parameters.

The conditional moments of the LMARX–GARCH model depend nonlinearly on the mixing weights and the parameters of the component densities. Their dynamics will thus be quite complicated. For example, the conditional mean is immediately seen to be the weighted average of the component means,

$$\bar{\mu}_t := E(r_t|\Psi_{t-1}) = \sum_{j=1}^{k} \lambda_{j,t}\mu_{j,t}, \tag{15}$$

while the conditional variance is

$$\bar{\sigma}_t^2 := \mathrm{Var}(r_t|\Psi_{t-1}) = \sum_{j=1}^{k} \lambda_{j,t}(\sigma_{j,t}^2 + \mu_{jt}^2) - \left(\sum_{j=1}^{k} \lambda_{j,t}\mu_{j,t}\right)^2$$

$$= \sum_{j=1}^{k} \lambda_{j,t}\sigma_{j,t}^2 + \sum_{j=1}^{k} \lambda_{j,t}(\mu_{j,t} - \bar{\mu}_t)^2 \tag{16}$$

$$= \sum_{j=1}^{k} \lambda_{j,t}\sigma_{j,t}^2 + \frac{1}{2}\sum_{i=1}^{k}\sum_{j=1}^{k} \lambda_{i,t}\lambda_{j,t}(\mu_{i,t} - \mu_{j,t})^2. \tag{17}$$

Thus, the conditional variance is the weighted average of the component variances plus a term that measures the distance between the means of the mixture components. Note that the second term in (16) can be interpreted as the variance of the conditional mean. In the two-component model considered below, the term involving the means in (17) becomes $\lambda_{1,t}(1 - \lambda_{1,t})(\mu_{1,t} - \mu_{2,t})^2$. The variance increases, for example, if the expected devaluation in case of a realignment is large. The coefficient of the squared distance between the means equals $\lambda_{1,t}(1-\lambda_{1,t})$, which is the variance of the conditional Bernoulli distribution over the mixture components.

Due to the different histories of the currencies within the EMS, the conditional densities differ for the Franc and the Pound. We discuss the model for the Franc first and subsequently outline the modifications that are necessary for the Pound.

3.2.1. Conditional density for the Franc

We assume that the conditional density of the exchange rate return process, r_t, is a two-component normal mixture density, that is,

$$f(r_t|\Psi_{t-1}) = \frac{\lambda_t}{\sigma_{1,t}\sqrt{2\pi}} e^{\{-(r_t-\mu_{1,t})^2/2\sigma_{1,t}^2\}} + \frac{1-\lambda_t}{\sigma_{2,t}\sqrt{2\pi}} e^{\{-(r_t-\mu_{2,t})^2/2\sigma_{2,t}^2\}}, \tag{18}$$

where information set Ψ_{t-1} consists of the exchange rates as well as further explanatory variables, such as interest rates.

With probability λ_t, there is a jump in the exchange rate, due to a realignment or a relatively large movement within the target zone. As in Bekaert and Gray (1998) and Neely (1999), the mixing weight, or probability of a jump, λ_t, depends on the slope of the yield curve, $yc_t = i_t^3 - i_t^1$, where i_t^3 and i_t^1 denote the three- and one-month interest rates, respectively. The functional relationship is specified in a logistic fashion. More specifically, we assume that

$$\lambda_t = \frac{1}{1+e^{\gamma_0+\gamma_1 yc_{t-1}^\star}}, \tag{19}$$

where $yc_t^\star = \operatorname{sign}(yc_t)\log(1+|yc_t|)$. We have also considered a probit specification in (19), where $\lambda_t = \Phi(\gamma_0+\gamma_1 yc_{t-1}^\star)$, and $\Phi(z) = (2\pi)^{-1/2}\int_{-\infty}^{z} e^{-\xi^2/2}\,d\xi$, which is used in Mizrach (1995), Bekaert and Gray (1998) and Neely (1999). Here, for the data at hand, it leads to virtually the same relation between λ_t and yc_{t-1}.[5] Beine and Laurent (2003) and Beine et al. (2003) use the logistic specification in modeling returns of the US$ against other major currencies, where the mixing weight depends on central bank interventions. In addition to using the probit specification, Bekaert and Gray (1998) and Neely (1999) work in terms of the untransformed variable yc_t, that is, they set $\lambda_t = \Phi(\gamma_0+\gamma_1 yc_{t-1})$.[6] The motivation for our use of the contracting transformation $yc_t^\star$ is illustrated in the left panel of Fig. 1, which plots r_t against the once-lagged slope measures yc_{t-1} and $yc_{t-1}^\star$, respectively, for the 172 monthly observations that form our estimation period. Obviously, using yc_{t-1} directly, estimated relationships between yc_{t-1} and the next period's density of r_t will suffer from the single large "outlier" $\min\{yc_t\} = -40$.

The mean of the jump-component, μ_{1t}, is also assumed to depend on yc_{t-1}, namely

$$\mu_{1,t} = \phi_0 + \phi_1 yc_{t-1}^\star. \tag{20}$$

The second mixture component in (18) represents the density of the exchange rate when the target zone is credible, so that, as in Neely (1999), it is plausible to let $\mu_{2,t}$ depend on the position of

---

5 A generalization of the probit approach to more than two mixture components is considered in Lanne and Saikkonen (2003).

6 Actually, Neely (1999) uses short-term interest rate differentials as a second explanatory variable. The latter and the slope of the yield curve are highly correlated, however, with a correlation coefficient of −0.8216 in our training sample. Engel and Hakkio (1996) let the transition probabilities in a Markov-switching model depend on the position of the exchange rate within its EMS band, but this did not lead to any improvement with our data.

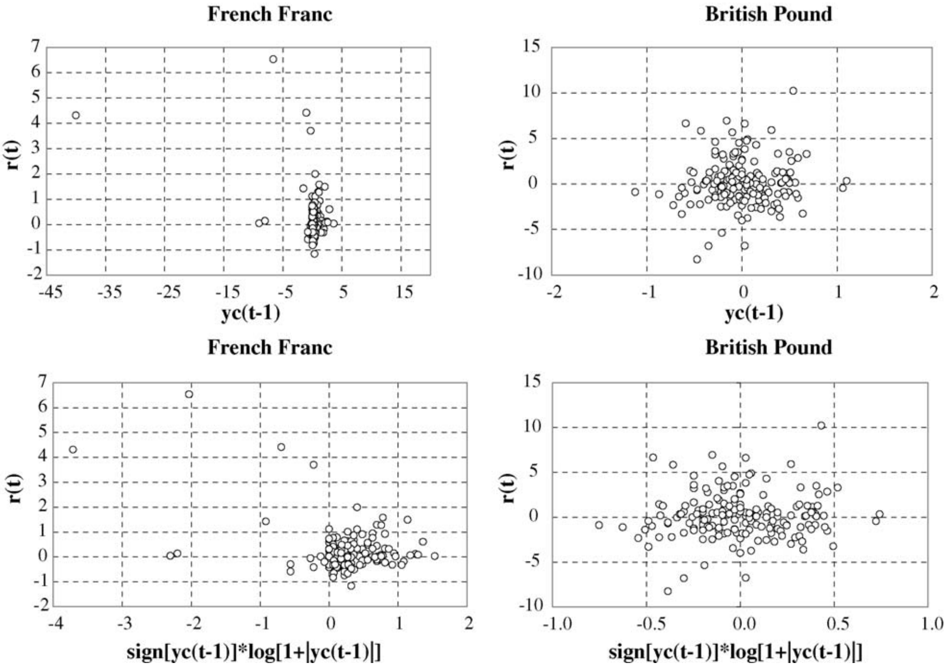

Fig. 1. Scatter plot of returns against slope of yield curve.

the exchange rate within the target zone. More specifically,

$$\mu_{2,t} = \psi_0 + \psi_1(S_{t-1} - P_{t-1}), \tag{21}$$

where P_t is the central parity at date t.

Finally, we discuss the conditional heteroskedasticity in the component variances $\sigma^2_{1,t}$ and $\sigma^2_{2,t}$. To do so, we reproduce the defining equation of the MN–GARCH process specified by Haas et al. (2004a) for the two-component GARCH(1,1) case, where (13) is of the form

$$\begin{bmatrix} \sigma^2_{1,t} \\ \sigma^2_{2,t} \end{bmatrix} = \begin{bmatrix} \alpha_{01} \\ \alpha_{02} \end{bmatrix} + \begin{bmatrix} \alpha_{11} \\ \alpha_{12} \end{bmatrix} \epsilon^2_{t-1} + \begin{bmatrix} \beta_{11} & \beta_{12} \\ \beta_{21} & \beta_{22} \end{bmatrix} \begin{bmatrix} \sigma^2_{1,t-1} \\ \sigma^2_{2,t-1} \end{bmatrix} \tag{22}$$

with $\epsilon_t = r_t - E(r_t|\Psi_{t-1}) = r_t - \lambda_t\mu_{1,t} - (1-\lambda_t)\mu_{2,t}$. Vlaar and Palm (1993) assume that, for all t, the difference between $\sigma^2_{1,t}$ and $\sigma^2_{2,t}$ is equal to a constant jump size, δ^2; that is, they restrict, in (22), $\alpha_{01} = \alpha_{02} + \delta^2$, $\alpha_{11} = \alpha_{12}$, $\beta_{12} = \beta_{22}$, and $\beta_{21} = \beta_{11} = 0$, so that $\sigma^2_{1,t} = \sigma^2_{2,t} + \delta^2$ for all t. Vlaar and Palm (1993) argue that "this procedure is preferred to that of independent variances, since it seems reasonable to assume that the same GARCH effect is present in all variances". This specification is also adopted in Neely (1999) and Beine and Laurent (2003). We will, however, not use this for the Franc, but rather employ the restricted version of (22), termed "partial MN–GARCH" in Haas et al. (2004a), which sets $\alpha_{11} = 0$, and $\beta_{11} = \beta_{12} = \beta_{21} = 0$, so that $\sigma^2_{1,t} = \sigma^2_1 = \alpha_{01}$ for all t. That is, only the variance in the "credibility regime" is driven by a GARCH process, while the variance in the jump component is constant. This specification seems more reasonable, given that, in a system of target zones, jumps are not expected to come clustered, so that a dynamic behavior of the jump component's variance would be difficult to interpret.

3.2.2. Modifications for the Pound

For the Pound, we use the model of the previous section with two modifications, which are enforced by the short duration of the Pound's membership in the EMS.

The first modification concerns the conditional mean in the second mixture component, given by (21). As we use monthly data from January 1978 to December 1991 to fit the model, there is no central parity for most of the data points. Thus, we replace (21) with a simple first-order autoregressive specification, i.e., for the BP (21) becomes

$$\mu_{2,t} = \psi_0 + \psi_1 r_{t-1}.$$

Secondly, we use a different specification for the conditional heteroskedasticity. In the previous section, we argued for the partial MN–GARCH structure because the first mixture component was suited to capture large jumps in the exchange rate, particularly due to realignments, which do not come clustered. The pound, however, did not join the EMS before October 1990, and so, this argument is not valid for this currency. Instead, we treat the components symmetrically and assume a GARCH(1,1) process in both components, where, for parsimony, we adopt the restricted specification of Vlaar and Palm (1993), where both components are driven by the same GARCH process. That is, the first component's variance is given by $\sigma_{2,t}^2 + \delta^2$, and[7]

$$\sigma_{2,t}^2 = \alpha_{02} + \alpha_{12}\epsilon_{t-1}^2 + \beta_{22}\sigma_{2,t-1}^2$$

describes the evolution of the variance in the second component.

4. Data and estimation results

4.1. Options market

4.1.1. Data

The majority of the intra-ERM derivatives trading is in the over-the-counter markets, and the data is not generally available to non-traders. The best publicly available data are for US dollar (US$) exchange rates which are traded in Philadelphia. We focus on the US Dollar/British Pound (US$/BP) and Dollar/French Franc (US$/FF) contracts. We have data for the years 1992 and 1993, which encompass both major ERM realignments.

The US$ appears to be an adequate proxy for the DM. During September 1992, the DM depreciated by −1.47% against the US$, while the BP depreciated −11.51%. From July 1 to August 5, 1993, the DM was similarly stable, depreciating −0.83%, while the Franc devalued by −3.59% against the US$.

Both American[8] and European options are traded. The BP options are for 31,250 Pounds, and the FF options are for 250,000 Francs. We use daily closing option prices that are quoted in cents. Spot exchange rates are expressed as US$ per unit of foreign currency and are recorded contemporaneously with the closing trade. Foreign currency appreciation (depreciation) will increase the

[7] It is, of course, not necessarily the case that the first mixture component has the higher variance, as implied by this specification. This is not just a labelling problem and may be a serious restriction in general, because the component means are modeled differently. However, it is not restrictive for the present data, as we confirmed by switching the roles of $\sigma_{1,t}^2$ and $\sigma_{2,t}^2$.

[8] Currency options may be thought of as options on a dividend paying stock where the dividend is equal to the foreign risk free rate. Early exercise is relevant for call options where the foreign risk free rate is high because this indicates that the currency is likely to devalue. The risk of devaluation will then be priced into American options of all maturities.

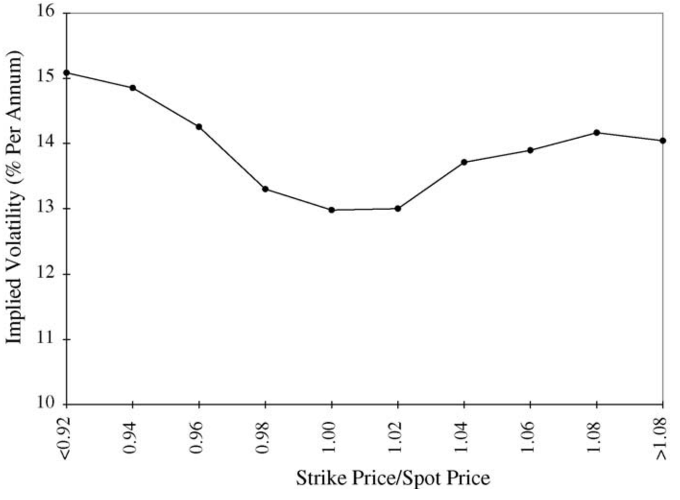

Fig. 2. Averages of implied volatility US$/FF options 1992 and 1993.

moneyness of a call (put) option. Interest rates are the Eurodeposit rates closest in maturity to the term of the option.

To obtain a rough idea about the implied volatility pattern in the currency options, we look at sample averages. We sort the data into bins based on the strike/spot ratio, S/K, and compute implied volatilities using the Black–Scholes formula. In Figs. 2 and 3, we plot the data for all of 1992 and 1993, for the FF and BP, respectively. Both appear to display the characteristic pattern, with the minima of the implied volatility at the money, and with higher implied volatilities in the two tails.

For estimation purposes, we excluded options that were more than 10% in or out of the money and with volumes less than 5 contracts. This seemed to eliminate most data points with unreasonably high implied volatilities. For the Pound, we looked at options from 5 to 75 days to maturity. Because the data were thinner with the Franc, we utilized all maturities greater than five days.

We will now try to infer whether changes in the smile signalled an impending crisis in the ERM.

4.1.2. Implied density estimation

There are two key issues in fitting the model. The first is to extend the analysis to American options which can be exercised before expiration. The second is choosing the loss function for estimation.

We approximate American puts and calls using the Bjerksund and Stensland (1993) approach. Hoffman (2000) shows that the Bjerksund–Stensland algorithm compares favorably in accuracy and computational efficiency to the Barone-Adesi and Whaley (1987) quadratic approximation. Our estimates were also quite similar using implied binomial trees.

Because $f(S_t)$ is the risk neutral density and is not directly observable, we must find a way to treat the options prices as sample “moments”. Let

$$\{d_{j,t}\}_{j=1}^{n} = [c(\tau_1, K_1), \dots, c(\tau_m, K_m), p(\tau_{m+1}, K_{m+1}), \dots, p(\tau_n, K_n)]$$

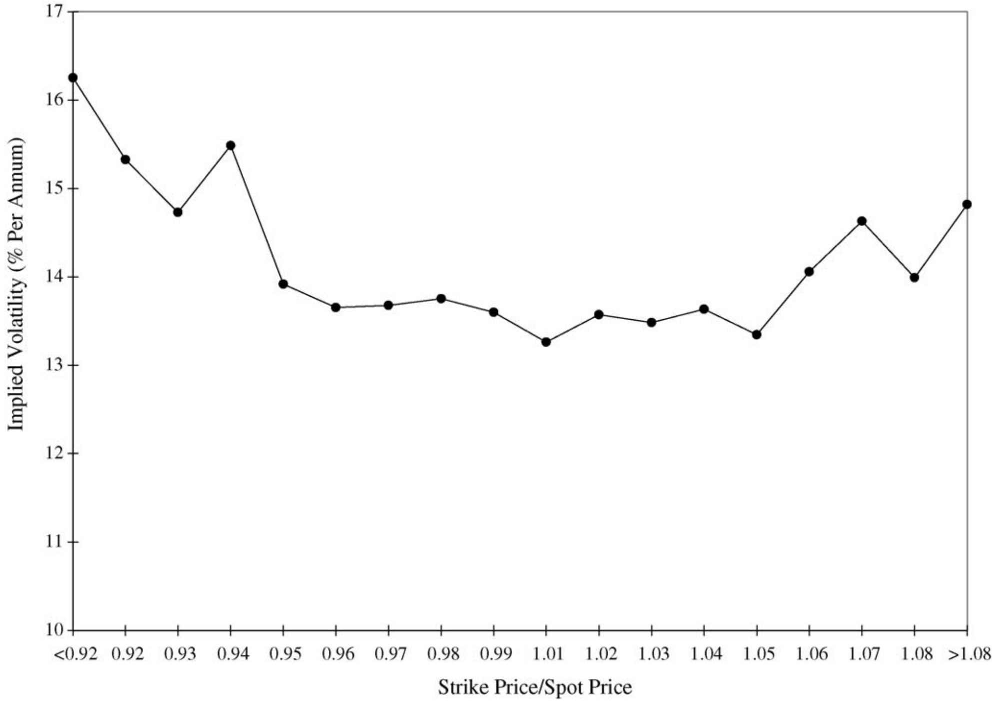

Fig. 3. Averages of implied volatility US$/BP options 1992 and 1993.

denote a sample of size n of the calls c and puts p traded at time t, with strike price K_j and expiring in τ_j years, and denote the pricing estimates from the model by $\{d_{j,t}(\theta)\}$.

In matching model to data, Christoffersen and Jacobs (2001) emphasize that the choice of loss function is important. Bakshi et al. (1997), for example, match the model to data using option prices. This can lead to substantial errors among the low priced options though. Since these options are associated with tail probability events, this is not the best metric for our exercise. We obtained the best fit overall using the implied Bjerksund–Stensland implied volatility,

$$\sigma_{j,t} = \mathrm{BJST}^{-1}(d_{j,t}, S_t, i_t). \tag{23}$$

Let the estimated volatility be denoted by

$$\sigma_{j,t}(\theta) = \mathrm{BJST}^{-1}(d_{j,t}(\theta), S_t, i_t). \tag{24}$$

We then minimize the sum of squared deviations from the implied volatility in the data,

$$\min_{\theta} \sum_{j=1}^{n} (\sigma_{j,t}(\theta) - \sigma_{j,t})^2. \tag{25}$$

As Christoffersen and Jacobs note, this is just a weighted least squares problem that, with the monotonicity of the option price in θ , satisfies the usual regularity conditions.

We next fit (25) to daily option prices for the FF and BP in intervals around the two crises. We first estimate the probability of a depreciation of at least 3% in a four-week horizon. We chose the jump size to be large enough for the BP to escape from the midpoint of the upper half of the band.

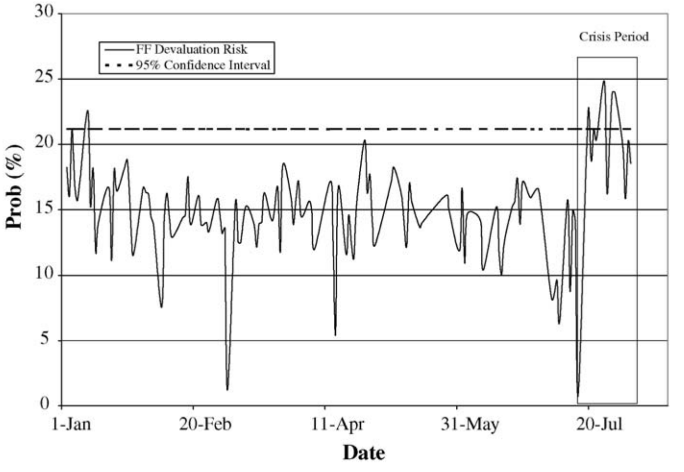

Fig. 4. Options market risk of 3% devaluation in the FF: January–August 1993.

We defer discussion of the entire predictive density until Section 5 after we develop forecasts using both options and spot market models.

4.1.3. French Franc options estimates

We estimate the six parameter model day-by-day from January 1 to August 5, 1993 for the FF. We report coefficient estimates, *t*-ratios, and R^2 in Table 1 for the crisis period, July 16 to August 5, 1993. The model describes the option prices well with an average goodness of fit of 97%.

From the fitted model, we back out an implied distribution for the spot exchange rate returns over a four week interval. We plot in Fig. 4 the 3% devaluation risk for January–August 1993.

We also compute empirical 95% confidence intervals based on the sampling distribution of the devaluation risk. A risk above 21.15% is in the upper 5% tail. All of the highest risk occur in the period immediately before and after the crisis. The one exception is the 22.57% spike on January 11, 1993 that quickly diminished.

In the period leading up to the crisis, the devaluation risk, depicted in Fig. 5, starts at less than 1% on July 18, quickly rises to nearly 23% on July 20 and peaks at nearly 25% on July 26. The risk stays above 20% for 6 of the 7 days prior to the FF's de facto devaluation.

This exercise, we feel, is largely successful. The model fits the data well and provides a sharp increase in devaluation risk 11 days before the FF bands widen. In principal, this could provide sufficient time for the central bank to react to market expectations.

4.1.4. British Pound options estimates

We next estimate the model for January 1 to September 29, 1992 for the BP. We report coefficient estimates, *t*-ratios, and R^2 in Table 2 for the crisis period August 19 to September 29, 1992. The model again captures the data well with an average R^2 of 96%.

The option implied devaluation risk is consistently under 20% and below the upper 5% risk level of 20.97% for all but 3 days prior to the crisis. On January 16, 17 and 24, 1992, the devaluation risk in Fig. 6 rises above 21%.

At the beginning of the crisis period displayed in Fig. 7, the devaluation risk on August 19, 1992 is below the sample average of 16.20%. It rises steadily into the crisis, except for two steep

Table 1
French Franc options mode

Date	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	R^2	LR
16-July-1993	−8.000(0.00)	−5.346(0.04)	3.766(0.00)	−4.908(12.53)	3.232(6.50)	−7.740(0.00)	0.910	
19-July-1993	−0.174(0.00)	1.928(0.80)	−5.876(1.82)	5.848(0.00)	−2.269(1.27)	−1.945(0.00)	0.990	2.3752(0.12)
20-July-1993	−1.019(0.69)	−1.720(108.26)	−0.191(1.79)	−3.441(0.00)	−4.006(0.00)	−4.233(0.00)	0.990	2.7213(0.10)
21-July-1993	−1.125(0.00)	−1.042(0.01)	−3.579(0.01)	−0.474(0.00)	−2.567(0.13)	−2.056(0.05)	0.999	1.1477(0.28)
22-July-1993	−1.505(0.02)	−3.058(0.29)	−4.188(0.00)	−6.801(28.36)	9.809(0.00)	−2.147(0.00)	0.997	1.7291(0.19)
23-July-1993	−2.504(0.03)	−2.460(0.22)	−0.325(0.00)	−2.405(0.01)	−1.034(0.01)	−2.040(0.07)	0.993	0.2145(0.64)
26-July-1993	−0.393(0.01)	−1.835(0.09)	−1.047(0.01)	−0.014(0.01)	−2.218(0.16)	−2.064(0.05)	0.990	1.1618(0.28)
27-July-1993	−3.945(0.06)	−2.111(36.96)	17.480(0.00)	−4.701(0.03)	−6.418(0.00)	−4.795(0.00)	0.994	2.3117(0.14)
28-July-1993	0.394(0.00)	−2.183(0.09)	−1.007(0.00)	−1.524(0.01)	−2.798(0.06)	−1.955(0.09)	0.999	2.1560(0.14)
29-July-1993	−1.231(0.00)	−2.162(0.02)	−1.040(0.00)	−0.328(0.00)	−2.110(0.03)	−1.826(0.01)	0.945	0.6379(0.42)
30-July-1993	−1.355(0.01)	−1.637(0.05)	−1.490(0.01)	−0.464(0.00)	−2.386(0.14)	−1.873(0.05)	0.977	1.1151(0.29)
2-August-1993	−1.029(0.01)	−0.035(0.01)	−4.371(0.73)	−0.324(0.00)	−2.139(0.73)	−2.222(1.20)	0.991	3.2378(0.11)
3-August-1993	3.192(0.23)	−0.562(0.06)	−3.473(0.21)	−0.332(0.07)	−3.207(1.52)	−1.930(1.14)	0.986	2.2335(0.14)
4-August-1993	−1.007(0.01)	−1.687(0.02)	−1.547(0.00)	−0.455(0.00)	−2.472(0.07)	−2.060(0.04)	0.912	1.6081(0.20)
5-August-1993	0.083(0.01)	−3.169(2.10)	−0.776(0.20)	−1.132(0.01)	−1.880(0.06)	−1.938(0.12)	0.974	1.3703(0.24)

The θ's are estimates of the model (25), t-ratios are in parentheses. The LR statistic, with p-values in parentheses, is given by (26) and is distributed $\chi^2(1)$.

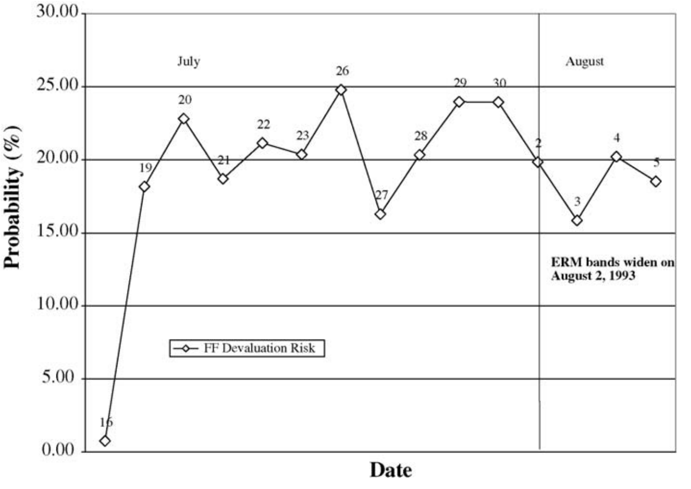

Fig. 5. Options market risk of 3% devaluation in the FF during ERM crisis: July–August 1993.

declines on September 4 and 11, 1992. The risk exceeds 20% for 17 out of 18 trading days prior to the BP devaluation on September 17, 1992.

The options again provide a potential early warning signal to policy makers. The devaluation risk exceeds the 5% limit on August 20, 1992, 25 days before the British Pound leaves the ERM.

We now turn to the spot market volatility to search for possible signals of the crises.

4.2. The spot market

4.2.1. French Franc

As we do not model the dynamics of the interest rates, and are interested in one-month-ahead density forecasts, we estimate the LMARX–GARCH model with monthly data. For the FF, we

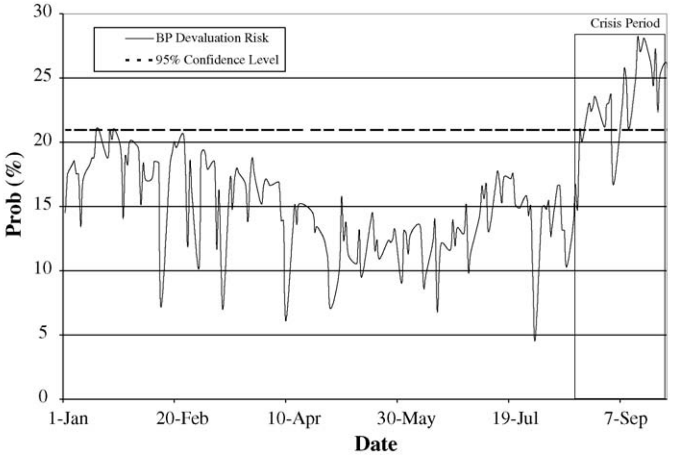

Fig. 6. Options market risk of 3% devaluation in the BP: January–September 1992.

Table 2
British Pound options model

Date	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	R^2	LR
19-August-1992	−30.857(0.00)	−2.468(5.96)	6.351(0.02)	−5.617(5.51)	1.879(1.84)	−1.982(0.00)	0.995	
20-August-1992	−2.179(0.01)	−1.908(0.15)	−0.816(0.01)	−1.855(0.01)	−1.908(0.07)	−2.440(0.63)	0.935	1.5008 (0.22)
21-August-1992	−1.620(0.01)	−1.623(0.07)	−2.679(0.02)	−1.099(0.01)	−1.859(0.14)	−2.435(0.90)	0.973	2.4618 (0.12)
24-August-1992	−2.374(0.26)	−2.609(4.98)	−0.452(0.27)	−1.805(0.08)	−1.409(0.37)	−1.897(1.01)	0.967	3.5960 (0.06)
25-August-1992	−17.296(0.00)	−2.183(15.64)	4.381(0.01)	−4.843(2.67)	0.987(0.63)	−2.163(0.00)	0.974	7.8066 (0.01)
26-August-1992	−2.306(0.12)	−2.614(1.10)	−0.467(0.05)	−2.018(0.05)	−1.595(0.30)	−1.865(0.60)	0.988	6.5960 (0.01)
27-August-1992	−2.601(0.05)	−2.465(0.57)	−0.479(0.02)	−1.722(0.02)	−1.706(0.26)	−1.934(0.29)	0.990	6.8458 (0.04)
31-August-1992	−2.314(0.00)	14.497(0.00)	−7.251(11.92)	−4.856(0.27)	−3.710(0.04)	−2.204(0.00)	0.908	4.6372 (0.03)
1-September-1992	−2.165(0.05)	−2.541(0.45)	−0.595(0.01)	−1.656(0.01)	−1.567(0.06)	−1.992(0.09)	0.948	2.6884 (0.10)
2-September-1992	−2.541(0.13)	−2.194(0.95)	0.930(0.01)	−5.933(2.28)	1.135(0.64)	−2.000(0.52)	0.974	2.6996 (0.10)
3-September-1992	−4.724(0.02)	−2.178(0.59)	−0.205(0.16)	−3.812(0.18)	−0.685(0.07)	−2.077(0.55)	0.945	4.1894 (0.04)
4-September-1992	2.824(0.66)	−0.324(0.27)	−3.620(1.15)	0.390(0.09)	−2.823(1.88)	−1.757(1.28)	0.990	3.3074 (0.07)
8-September-1992	−3.384(0.13)	−2.216(1.40)	(0.01)0.734	−4.205(1.85)	−0.028(0.02)	−2.040(0.51)	0.992	5.4434 (0.02)
9-September-1992	−0.651(0.01)	43.422(0.00)	−8.564(16.27)	−2.607(6.56)	−10.131(0.00)	−1.936(0.00)	0.960	4.9878 (0.03)
10-September-1992	−2.968(0.02)	−0.836(0.13)	−2.859(0.17)	0.017(0.11)	−2.226(0.94)	−2.019(0.50)	0.990	3.4620 (0.06)
11-September-1992	−2.259(0.06)	−3.154(1.22)	−0.597(0.52)	−1.889(0.79)	−0.879(1.43)	−1.891(3.37)	0.965	3.9484 (0.05)
14-September-1992	−1.592(0.06)	−2.197(0.84)	−0.517(0.01)	−3.083(0.16)	−0.737(0.09)	−2.011(0.55)	0.978	2.9796 (0.08)
15-September-1992	−3.353(0.01)	−1.991(0.17)	−1.054(0.01)	−1.719(0.01)	−1.801(0.09)	−1.898(0.21)	0.980	1.7626 (0.18)
16-September-1992	−128.074(0.00)	−2.008(40.88)	16.878(0.00)	−5.017(53.64)	5.456(0.00)	−3.079(0.00)	0.980	0.8910 (0.35)
17-September-1992	−0.882(0.07)	−0.943(0.48)	−3.634(0.48)	−3.108(0.02)	−2.051(0.09)	−1.933(2.12)	0.985	0.7756 (0.38)
18-September-1992	−5.051(0.05)	−2.165(0.49)	−0.453(0.01)	−3.130(0.17)	−0.889(0.19)	−1.735(0.49)	0.928	0.8970 (0.34)
21-September-1992	−4.066(0.01)	−2.235(0.43)	−0.436(0.03)	−1.379(0.08)	−2.102(0.13)	−1.601(1.48)	0.908	0.8004 (0.37)
22-September-1992	−6.819(0.00)	−2.284(1.50)	−0.241(0.07)	−2.746(0.15)	−0.900(0.16)	−1.886(0.68)	0.975	0.4570 (0.50)
23-September-1992	−1.142(0.00)	−1.678(0.05)	−4.145(0.01)	−6.490(4.69)	3.387(0.01)	−1.918(2.35)	0.987	0.7244 (0.39)
24-September-1992	−2.741(0.05)	−2.980(1.07)	−0.844(0.18)	−2.448(0.10)	−0.977(0.14)	−1.787(0.75)	0.955	0.4260 (0.51)
25-September-1992	−3.089(0.03)	−2.012(0.38)	0.772(0.00)	−4.469(0.24)	0.121(0.01)	−1.956(0.17)	0.984	0.7130 (0.40)
28-September-1992	−2.325(0.04)	−2.460(0.53)	−0.781(0.04)	−1.612(0.12)	−1.536(0.30)	−1.721(0.38)	0.970	1.4356 (0.23)
29-September-1992	−1.886(0.09)	−0.884(0.39)	−2.202(0.29)	−0.640(0.15)	−2.703(2.00)	−1.846(1.16)	0.975	2.3652 (0.12)

The θ's are estimates of the model (25), t-ratios are in parentheses. The LR statistic, with p-values in parentheses, is given by (26) and is distributed $\chi^2(1)$.

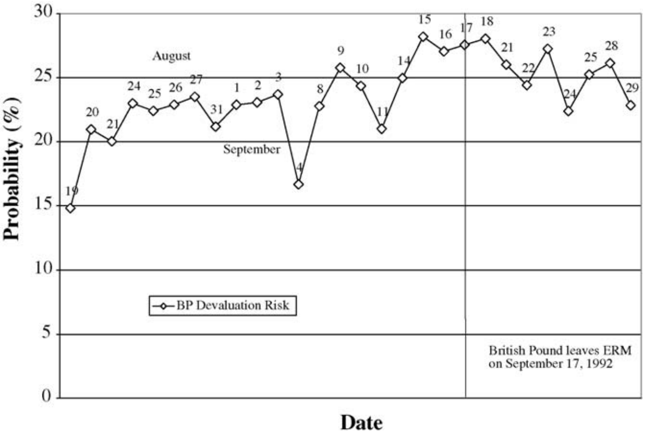

Fig. 7. Options market risk of 3% devaluation in BP during ERM crisis August–September 1992.

Table 3
French Franc Spot Exchange Rate Model

σ_1^2	α_{02}	α_{12}	β_{22}	γ_0
2.945 (1.705)	0.021 (0.011)	0.019 (0.015)	0.836 (0.071)	2.988 (0.807)
γ_1	ϕ_0	ϕ_1	ψ_0	ψ_1
1.711 (0.621)	2.048 (1.297)	−1.049 (0.574)	0.093 (0.036)	−2.820 (1.066)

Shown are the parameter estimates for the LMARX–GARCH for the French Franc (with approximate standard errors in parentheses), which is given by the following equations for the conditional density $f(r_t|\Psi_{t-1})$:

$$\begin{aligned}
f(r_t|\Psi_{t-1}) &= \lambda_t\phi(r_t;\mu_{1,t},\sigma_{1,t}^2) + (1-\lambda_t)\phi(r_t;\mu_{2,t},\sigma_{2,t}^2),\\
r_t &= 100 \times \log(S_t/S_{t-1}),\\
\phi(y;\mu,\sigma^2) &= \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{(y-\mu)^2}{2\sigma^2}\right\},\\
\lambda_t &= \left(1+\exp\{\gamma_0+\gamma_1 yc_{t-1}^\star\}\right)^{-1},\\
\mu_{1,t} &= \phi_0 + \phi_1 yc_{t-1}^\star,\\
\mu_{2,t} &= \psi_0 + \psi_1(S_{t-1}-P_{t-1}),\\
\sigma_{2,t}^2 &= \alpha_{02} + \alpha_{12}\epsilon_{t-1}^2 + \beta_{22}\sigma_{2,t-1}^2,\\
\epsilon_t &= r_t - \lambda_t\mu_{1,t} - (1-\lambda_t)\mu_{2,t}.
\end{aligned}$$

S_t is the exchange rate at time t, $yc_t^\star = sign(yc_t)\log(1+|yc_t|)$, where yc_t is the slope of the French yield curve, i.e., the difference between the three- and one-month interest rates, and P_t is the central parity at time t.

use monthly percentage returns, $r_t = 100 \times \log(S_t/S_{t-1})$, from May 1979 to December 31, 1992, a total of 172 monthly observations. Maximum likelihood estimates[9] of the model described in Section 3.2.1 are reported in Table 3.

[9] See Haas et al. (2004a) for a discussion of maximum likelihood estimation. See also Alexander and Lazar (2004) for the special GARCH(1,1)-mixture case.

Table 4
French Franc spot exchange rate densities

Date	λ_t	$\mu_{1,t}$	$\mu_{2,t}$	$\sigma^2_{1,t}$	$\sigma^2_{2,t}$	$\bar{\mu}_t$	$\bar{\sigma}^2_t$	LR
16-July-1993	0.073	2.323	−0.084	2.945	0.206	0.092	0.799	
19-July-1993	0.074	2.331	−0.068	2.945	0.186	0.109	0.785	2.7435 (0.10)
20-July-1993	0.068	2.274	−0.075	2.945	0.188	0.085	0.724	1.7474 (0.19)
21-July-1993	0.048	2.048	−0.085	2.945	0.190	0.017	0.530	0.8250 (0.36)
22-July-1993	0.068	2.274	−0.087	2.945	0.182	0.073	0.722	0.3478 (0.56)
23-July-1993	0.232	3.147	−0.087	2.945	0.178	0.664	2.686	0.0003 (0.99)
26-July-1993	0.248	3.201	−0.079	2.945	0.170	0.735	2.867	0.0365 (0.85)
27-July-1993	0.248	3.201	−0.078	2.945	0.170	0.736	2.865	0.0185 (0.89)
28-July-1993	0.182	2.958	−0.064	2.945	0.164	0.485	2.028	0.1819 (0.67)
29-July-1993	0.168	2.899	−0.082	2.945	0.162	0.419	1.871	0.1413 (0.71)
30-July-1993	0.375	3.566	−0.121	2.945	0.175	1.261	4.399	0.0032 (0.96)
02-August-1993	0.168	2.899	−0.299	2.945	0.303	0.238	2.177	0.0044 (0.95)
03-August-1993	0.209	3.062	−0.318	2.945	0.334	0.387	2.767	0.1947 (0.66)
04-August-1993	0.262	3.245	−0.184	2.945	0.217	0.715	3.207	0.0040 (0.95)
05-August-1993	0.242	3.179	−0.222	2.945	0.226	0.601	3.006	0.0175 (0.89)

Column 1 shows the day when the four-week-ahead forecast density is computed. Columns 2–6 report the parameters of the predictive four-week-ahead normal mixture density for the respective trading days. Columns 7 and 8 report the overall mean and variance, $\bar{\mu}_t := E(r_t|\Psi_{t-1}) = \lambda_t\mu_{1,t} + (1-\lambda_t)\mu_{2,t}$ and $\bar{\sigma}^2_t := \mathrm{Var}(r_t|\Psi_{t-1}) = \lambda_t\sigma^2_{1,t} + (1-\lambda_t)\sigma^2_{2,t} + \lambda_t(1-\lambda_t)(\mu_{1,t}-\mu_{2,t})^2$. The last column shows the LR statistic (26), with p-values in parentheses, which is distributed $\chi^2(1)$.

As expected, $\gamma_1 > 0$ and $\phi_1 < 0$, so that both the probability of a jump, λ_t, as well as the expected jump size, $\mu_{1,t}$, increase when the yield curve inverts. Also, $\psi_1 < 0$, that is, there is mean reversion when the target zone is credible.

From the fitted model, we compute the four-week-ahead densities for the period from July 16 to August 5, 1993. The implied densities of the percentage log-change of the FF against the DM four weeks from the trading date are summarized in Table 4.

To illustrate the flexibility of the density forecasts resulting from the LMARX–GARCH model, the top panel of Fig. 8 shows the predictive densities calculated for July 21 and 30, respectively. While the density forecast of July 21 is somewhat skewed to the right, the predictive density for July 30, shortly before the de facto devaluation of the Franc, exhibits a pronounced bimodality. Fig. 8 also shows the probability weighted mixture components, i.e., $\lambda_t\phi(r_t;\mu_{1,t},\sigma^2_{1,t})$ and $(1-\lambda_t)\phi(r_t;\mu_{2,t},\sigma^2_{2,t})$, in the middle panel, as well as the raw densities, $\phi(r_t;\mu_{1,t},\sigma^2_{1,t})$ and $\phi(r_t;\mu_{2,t},\sigma^2_{2,t})$, in the bottom panel. The weighted densities document the contribution of each component to the overall mixture density. Hence, the middle panel illustrates the increasing importance of the first component a few days before the crisis. From the bottom panel, we note that the probability mass of the second (credibility) density is essentially concentrated between −2 and 2, as implied by the exchange rate mechanism.

The normal mixture densities extracted from the time series of currency prices demonstrate a considerable increase in downside risk at least a week before the de facto devaluation of the Franc, with a further sharp increase immediately before the widening of the target zone, that is, on July 30. The evolution of the probabilities is shown in Figs. 9 and 10.

4.2.2. British Pound

Given the short period of time of the BP belonging to the EMS, we do not necessarily expect to fit a meaningful model as we did for the FF. This is already evident from the right panel of

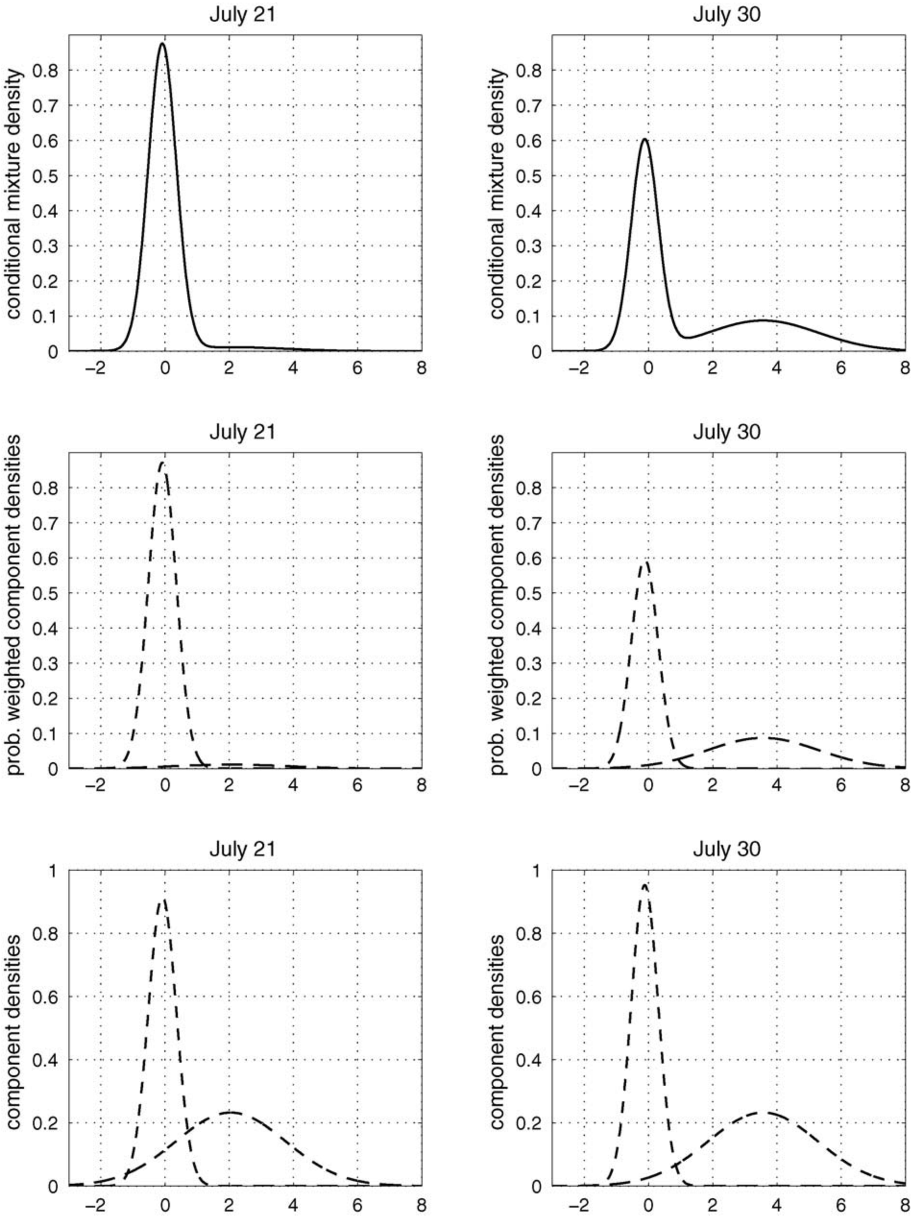

Fig. 8. Spot market density predictions for the FF. The Figure depicts the predicted densities for the FF from July 21 and 30, 1993 (top panel), probability weighted components (middle) and raw components (bottom).

Fig. 1, where the relation between the yield curve and the next period's return is shown for the BP. Obviously, there is much less information in the British yield curve than is in the French. This, of course, was expected due to the well-known differences between exchange rates in the EMS and floating systems (Neely, 1994).

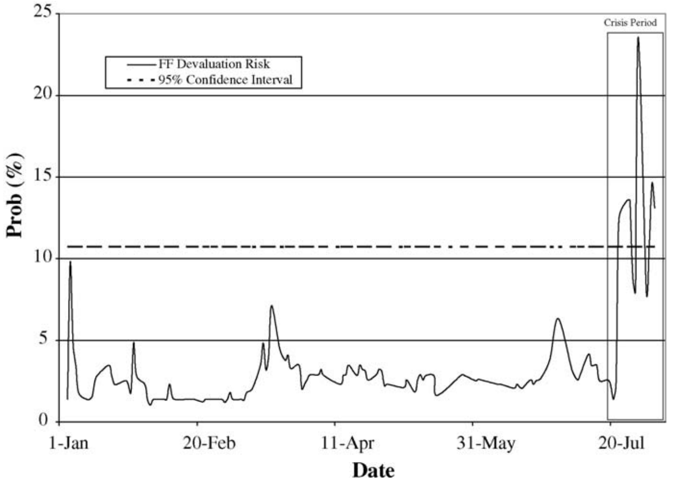

Fig. 9. Spot market risk of 3% devaluation in the FF: January–August 1993.

We make use of pre-ERM data, that is, we use monthly returns from January 1978 to December 31, 1991 (176 observations), to fit the MN–GARCH model discussed in Section 3.2.2. The parameter estimates for this model are reported in Table 5. The implied densities of the percentage log-change of the BP against the DM four weeks from the trading date, for the period August 19 to September 29, 1992 are summarized in Table 6.

The conditional four-week-ahead densities for a pre- and a post-crisis day are shown in Fig. 11. The mixture components are both centered around zero, only their variances differ, so there is a considerable overlap. This is in contrast to the results for the Franc, where the components are very well separated, because their means are far enough apart, relative to the variances. The mixture for the Pound, thus, mainly captures the kurtosis in the data, but incorporates no information about regime-specific conditional means.

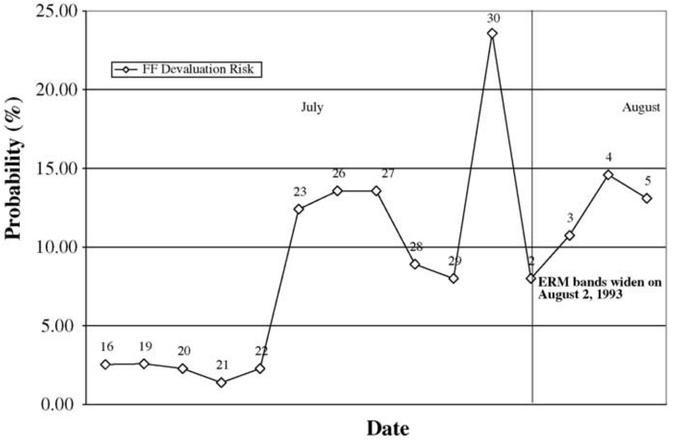

Fig. 10. Spot market risk of 3% devaluation in the FF during ERM crisis: July–August 1993.

Table 5
British Pound spot exchange rate model

δ^2	α_{02}	α_{12}	β_{22}	γ_0
9.644 (2.612)	0.141 (0.150)	0.099 (0.054)	0.676 (0.087)	0.539 (0.433)

γ_1	ϕ_0	ϕ_1	ψ_0	ψ_1
0.896 (1.132)	0.960 (0.626)	2.776 (2.018)	−0.153 (0.190)	−0.048 (0.096)

Shown are the parameter estimates for the LMARX–GARCH model for the British Pound (with approximate standard errors in parentheses), which is given by the following equations for the conditional density $f(r_t|\Psi_{t-1})$:

$$f(r_t|\Psi_{t-1}) = \lambda_t\phi(r_t; \mu_{1,t}, \sigma^2_{1,t}) + (1-\lambda_t)\phi(r_t; \mu_{2,t}, \sigma^2_{2,t}),$$
$$r_t = 100 \times \log(S_t/S_{t-1}),$$
$$\phi(y; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{(y-\mu)^2}{2\sigma^2}\right\},$$
$$\lambda_t = \left(1 + \exp\{\gamma_0 + \gamma_1 yc^\star_{t-1}\}\right)^{-1},$$
$$\mu_{1,t} = \phi_0 + \phi_1 yc^\star_{t-1},$$
$$\mu_{2,t} = \psi_0 + \psi_1 r_{t-1},$$
$$\sigma^2_{2,t} = \alpha_{02} + \alpha_{12}\epsilon^2_{t-1} + \beta_{22}\sigma^2_{2,t-1},$$
$$\sigma^2_{1,t} = \sigma^2_{2,t} + \delta^2,$$
$$\epsilon_t = r_t - \lambda_t\mu_{1,t} - (1-\lambda_t)\mu_{2,t}.$$

S_t is the exchange rate at time t, and $yc^\star_t = sign(yc_t)\log(1 + |yc_t|)$, where yc_t is the slope of the British yield curve, i.e., the difference between the three- and one-month interest rates.

The model's probabilities of a devaluation of at least 3% are shown in Figs. 12 and 13, with Fig. 12 displaying empirical 95% confidence limits.

As was expected in view of the right panel of Fig. 1, the parameter estimates in Table 5 do not represent meaningful economic relationships. The positive γ_1 implies that the weight of the first component increases when the yield curve inverts, but its mean, $\mu_{1,t}$, decreases. This means that the probability of a devaluation decreases when the yield curve inverts, and gives rise to the very strange result that the probability of a large devaluation initially decreases on September 17, 1992 when the Pound left the EMS and the yield curve is negative with $yc_t = i^3_t - i^1_t = -2.5$. Subsequently, the probability increases only as a result of the GARCH effects in the component variances. Clearly the numbers describing the conditional mean dynamics cannot be given any interpretation, due to their very large (approximate) standard errors. However, the MN–GARCH structure is reasonable for the data, given that the jump size δ^2 is rather large, implying a considerable difference between the component variances.

Summarizing the results for the Pound, we conclude that lacking characteristic information in the sample used for estimation, the mixed normal GARCH model does not anticipate the withdrawal of the pound from the EMS, and so, the probability of a large devaluation rises only ex-post, due to the GARCH effects in the mixture model.

5. Comparison of predictive densities

Next, we evaluate the forecast densities produced by our two models. The approach we take is the one originally proposed by Berkowitz (2001). Let $f(s_t)$ be the probability density of the spot

Table 6
British Pound spot exchange rate densities

Date	λ_t	$\mu_{1,t}$	$\mu_{2,t}$	$\sigma^2_{1,t}$	$\sigma^2_{2,t}$	$\bar{\mu}_t$	$\bar{\sigma}^2_t$	LR
19-August-1992	0.328	1.512	−0.191	10.779	1.135	0.368	4.938	
20-August-1992	0.323	1.580	−0.208	10.785	1.141	0.370	4.957	2.0796 (0.15)
21-August-1992	0.323	1.580	−0.226	10.751	1.107	0.357	4.938	3.1264 (0.08)
24-August-1992	0.323	1.580	−0.231	10.691	1.046	0.355	4.881	2.8652 (0.09)
25-August-1992	0.316	1.689	−0.243	10.802	1.158	0.367	5.008	1.7725 (0.18)
26-August-1992	0.328	1.512	−0.234	10.674	1.030	0.339	4.866	2.5806 (0.11)
27-August-1992	0.310	1.773	−0.243	10.815	1.171	0.381	5.027	2.6934 (0.10)
28-August-1992	0.292	2.030	−0.245	10.909	1.265	0.420	5.154	2.9932 (0.08)
1-September-1992	0.289	2.086	−0.257	10.925	1.280	0.419	5.191	2.2503 (0.13)
2-September-1992	0.270	2.367	−0.238	10.820	1.176	0.467	5.121	2.7256 (0.10)
3-September-1992	0.310	1.773	−0.208	10.622	0.978	0.405	4.805	4.2503 (0.04)
4-September-1992	0.340	1.348	−0.198	10.462	0.818	0.328	4.631	6.0795 (0.01)
7-September-1992	0.325	1.558	−0.200	10.426	0.782	0.371	4.592	7.3170 (0.01)
8-September-1992	0.328	1.512	−0.221	10.494	0.850	0.348	4.676	3.7090 (0.05)
9-September-1992	0.306	1.834	−0.225	10.541	0.897	0.404	4.744	3.0282 (0.08)
10-September-1992	0.319	1.646	−0.215	10.591	0.947	0.378	4.771	2.3305 (0.13)
11-September-1992	0.314	1.710	−0.208	10.479	0.835	0.394	4.657	1.5798 (0.21)
14-September-1992	0.362	1.043	−0.161	10.487	0.843	0.275	4.672	2.4595 (0.12)
15-September-1992	0.338	1.372	−0.216	10.582	0.938	0.321	4.763	1.4172 (0.23)
16-September-1992	0.405	0.478	−0.214	10.612	0.967	0.066	4.992	1.9766 (0.16)
17-September-1992	0.642	−2.517	−0.476	14.474	4.830	−1.786	11.978	2.6203 (0.11)
18-September-1992	0.465	−0.274	−0.507	15.322	5.678	−0.399	10.175	1.7021 (0.19)
21-September-1992	0.475	−0.396	−0.625	19.275	9.631	−0.516	14.222	1.3931 (0.24)
22-September-1992	0.456	−0.165	−0.615	18.857	9.212	−0.410	13.662	0.6636 (0.42)
23-September-1992	0.407	0.454	−0.565	17.138	7.494	−0.150	11.672	1.2587 (0.26)
24-September-1992	0.411	0.409	−0.618	18.959	9.315	−0.196	13.531	0.4803 (0.49)
25-September-1992	0.394	0.621	−0.626	19.293	9.649	−0.134	13.823	0.2549 (0.61)
28-September-1992	0.413	0.386	−0.653	20.359	10.715	−0.225	14.955	0.3857 (0.53)
29-September-1992	0.396	0.597	−0.625	19.244	9.600	−0.141	13.777	0.7949 (0.37)
30-September-1992	0.416	0.341	−0.659	20.430	10.786	−0.243	15.041	0.3935 (0.53)

Column 1 shows the day when the four-week-ahead forecast density is computed. Columns 2–6 report the parameters of the predictive four-week-ahead normal mixture density. Columns 7 and 8 report the overall mean and variance, $\bar{\mu}_t := E(r_t|\Psi_{t-1}) = \lambda_t\mu_{1,t} + (1-\lambda_t)\mu_{2,t}$ and $\bar{\sigma}^2_t := \mathrm{Var}(r_t|\Psi_{t-1}) = \lambda_t\sigma^2_{1,t} + (1-\lambda_t)\sigma^2_{2,t} + \lambda_t(1-\lambda_t)(\mu_{1,t}-\mu_{2,t})^2$. The last column shows the LR statistic (26), with p-values in parentheses, which is distributed $\chi^2(1)$.

exchange rate, and let $F(s_t)$ be the cumulative distribution

$$F(s_t) = \int_{-\infty}^{s_t} f(u)\,\mathrm{d}u.$$

Berkowitz notes that estimates $\hat{F}(s_t)$ are uniform, independent and identically distributed under fairly weak assumptions.

Testing for an independent uniform density in small samples can be problematic, so Berkowitz suggests transforming the data into normal random variates,

$$z_t = \Phi^{-1}(\hat{F}(s_t)),$$

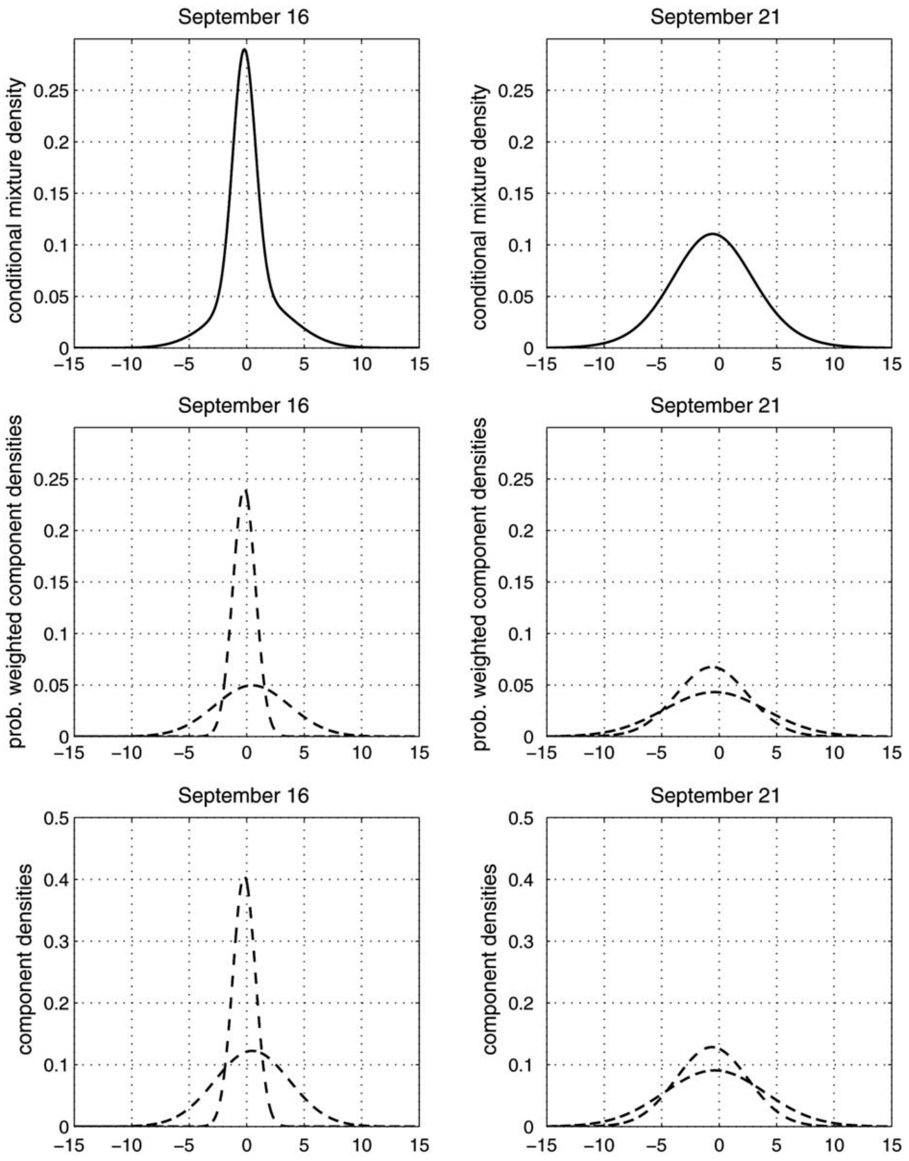

Fig. 11. Spot market density predictions for the BP. The Figure depicts the predicted densities for the BP from September 16 and 21, 1992 (top panel), probability weighted components (middle) and raw components (bottom).

where $\Phi(\cdot)$ denotes again the standard normal distribution.[10] The likelihood ratio,

$$\mathrm{LR} = \sum_{t=1}^{20} \left(\frac{z_t^2}{\hat{\sigma}^2} - 1 \right), \tag{26}$$

[10] We use the numerical transformation for the inverse normal proposed by Wichura (1988).

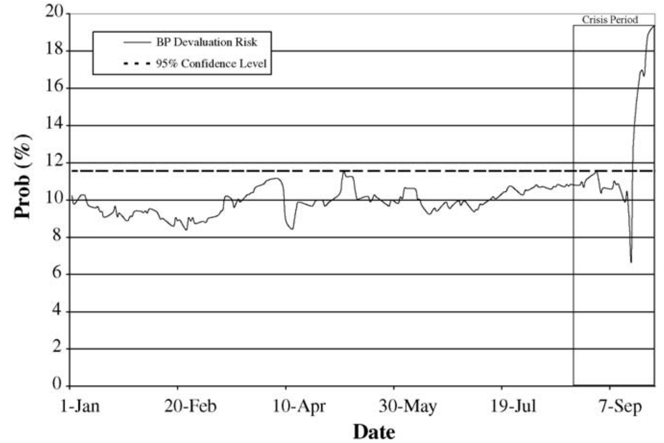

Fig. 12. Spot market risk of 3% devaluation in the BP: January–September 1992.

where $\hat{\sigma}$ is the forecast standard deviation, is then approximately distributed $\chi^2(1)$ for the null hypothesis that the transformed forecast statistics, z_t, have mean zero.

5.1. Option forecasts

We test the forecast densities for the FF from July 20 to August 29, 1993. Likelihood-ratio statistics are in the last column of Table 1. At the 10% level, we can accept the null that our forecast could have generated the subsequent four weeks of trading data from July 21 through the rest of the crisis. After that point, our model is statistically indistinguishable from the post-crisis density except for two days in August.

We do the same exercise for the BP for the period August 20 to September 29, 1992. There are stronger rejections prior to this crisis. Nonetheless, on August 20, 21 and September 1 and 2, we have a forecast consistent with the four-week returns data at the 10% level.

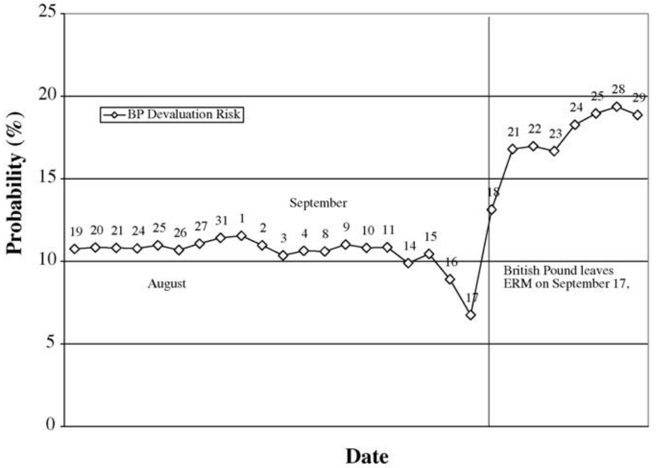

Fig. 13. Spot market risk of 3% devaluation in the BP during EMS crisis: August–September 1992.

5.2. Spot market forecasts

Ignoring non-trading days, as we do in model specification and estimation, the 20-trading day-ahead forecast density is given by a mixture of two normals, namely,

$$f(r_{t+20}|\Psi_t) = \sum_{j=1}^{2} \lambda_{j,t+20} \frac{1}{\sqrt{2\pi}\sigma_{j,t+20}} \mathrm{e}^{\{-(r_{t+20}-\mu_{j,t+20})^2/2\sigma_{j,t+20}^2\}}, \tag{27}$$

where $r_{t+20} = 100 \times (\log S_{t+20} - \log S_t)$. Under constancy assumptions, we can scale the 20-day ahead densities to obtain daily log-changes $r_{t+\tau}^{\mathrm{d}} := 100 \times (\log S_{t+\tau} - \log S_{t+\tau-1})$, $\tau = 1, \ldots, 20$, implying a two-component normal mixture distribution, given by

$$f(r_{t+\tau}^{\mathrm{d}}|\Psi_t) = \sum_{j=1}^{2} \lambda_{j,t+20} \frac{1}{\sqrt{2\pi}\sigma_{j,t+20}/\sqrt{20}} \mathrm{e}^{\{-(r_{t+\tau}^{\mathrm{d}}-\mu_{j,t+20}/20)^2/2\sigma_{j,t+20}^2/20\}}. \tag{28}$$

Expression (28) can be used to compute the cumulative distribution function $F(r_{t+\tau}^{\mathrm{d}}|\Psi_t)$ and transformation $z_t = \Phi^{-1}(F(r_{t+\tau}^{\mathrm{d}}|\Psi_t))$, $\tau = 1, \ldots, 20$. Then, provided our density predictions are correctly specified, the likelihood ratio (26) again has an approximate $\chi^2(1)$ distribution. Using (26), we test for a correct specification of the mean of our forecast distribution. In principle, we could test for additional properties of the forecast density, such as skewness – reflecting in our mixture models some sense of the realignment risk – or kurtosis. However, with only 20 data points at hand, any test involving higher-order forecast moments is rather questionable.

The test results are reported in Tables 4 and 6 for the FF and BP, respectively. In terms of the LR test (26), the dynamic mixture model performs well for the FF, but, as expected, exhibits a poor performance in predicting the crisis of the BP.

6. Conclusion

Analyzed with some recently developed modeling techniques, asset prices can provide insights into the entire probability distribution of future events. This paper has utilized the mixture of log normals in two separate contexts: with options and with the underlying currencies.

The crisis of the European Exchange Rate Mechanism case was certainly an epochal event for the markets, where central bankers became aware – perhaps for the first time – that the markets might be an irresistible force.

Policy makers may find these tools and inference worthwhile in a variety of contexts. Their subjective weights between types I and II errors should not only be tested ex-post but incorporated directly in the estimation. Both Skouras (2001) and Christoffersen and Jacobs (2001) have made progress along these lines. Loss aversion on the part of investors and traders may give them similar preferences.

Whether or not the accuracy of density predictions can be improved by combining options and spot-market information is the subject of future research. One possible strategy in this direction, employed in Claessen and Mittnik (2002), is to use implied volatility as an explanatory variable in the GARCH equation. Alternatively, the predicted density could be formed by a mixture of options- and the spot market-based density predictions.

Acknowledgements

The research of Haas and Mittnik was supported by a grant from the Deutsche Forschungsgemeinschaft (SFB 386). Mittnik conducted part of the research while visiting the Department of Economics at Washington University in St. Louis with a grant from the Fulbright Commission. We thank our discussant Paul Wachtel for his valuable comments.

References

Alexander, C., Lazar, E, 2004. Normal Mixture GARCH(1,1), Applications to Exchange Rate Modelling, ISMA Centre Discussion Papers in Finance 2004–2006, The Business School for Financial Markets at the University of Reading.

Bakshi, C., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. J. Finance 52, 2003–2049.

Barone-Adesi, G., Whaley, R.E., 1987. Efficient analytic approximation of American option values. J. Finance 42, 301–320.

Bates, D., 1991. The crash of '87: was it expected? The evidence from options markets. J. Finance 46, 1009–1044.

Beine, M., Laurent, S., 2003. Central bank interventions and jumps in double long memory models of daily exchange rates. J. Emp. Finance 10, 641–660.

Beine, M., Laurent, S., Lecourt, C., 2003. Official central bank interventions and exchange rate volatility: evidence from a regime-switching analysis. Eur. Econ. Rev. 47, 891–911.

Bekaert, G., Gray, S.F., 1998. Target zones and exchange rates: an empirical investigation. J. Int. Econ. 45, 1–35.

Berkowitz, J., 2001. Testing density forecasts with applications to risk management. J. Business Econ. Stat. 19, 466–474.

Bjerksund, P., Stensland, G., 1993. Closed-form approximation of American options. Scand. J. Manage. 9, 87–99.

Bliss, R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. J. Bank. Finance 26, 381–422.

Breeden, D., Litzenberger, R., 1978. State contingent prices implicit in option prices. J. Business 51, 621–651.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. J. Econ. 31, 307–327.

Campa, J., Chang, P.H.K., 1996. Arbitrage based tests of target zone credibility: evidence from ERM cross rate options. Am. Econ. Rev. 86, 726–740.

Christoffersen, P., Jacobs, K., 2001. The Importance of the Loss Function in Option Pricing, CIRANO Working Paper 2001-45.

Christoffersen, P., Mazzotta, S., 2004. The Informational Content of Over the Counter Currency Options, CIRANO Working Paper 2004-16.

Claessen, H., Mittnik, S., 2002. Forecasting stock market volatility and the informational efficiency of the DAX-index options market. Eur. J. Finance 8, 302–321.

Das, S.R., Sundaram, R.K., 1999. Of smiles and smirks: a term-structure perspective. J. Financial Quantit. Anal. 34, 211–239.

Dumas, B., Fleming, J., Whaley, R., 1998. Implied volatility functions: empirical tests. J. Finance 53, 2059–2106.

Dupire, B., 1994. Pricing with a smile. Risk 7, 18–20.

Engel, C., Hakkio, C.S., 1996. The distribution of exchange rates in the EMS. Int. J. Finance Econ. 1, 55–67.

Haas, M., Mittnik, S., Paolella, M.S., 2004a. Mixed normal conditional heteroskedasticity. J. Financial Econometrics 2, 211–250.

Haas, M., Mittnik, S., Paolella, M.S., 2004b. A new approach to Markov-switching GARCH models. J. Financial Econometrics 2, 493–530.

Heston, S., 1993. A closed form solution for options with stochastic volatility with applications to bond and currency options. Rev. Financial Stud. 6, 327–343.

Hoffman, C., 2000. Valuation of American Options, Oxford University thesis.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatility. J. Finance 42, 281–300.

Klaster, M.A., Knot, K.H.W., 2002. Toward an econometric target zone model with endogenous devaluation risk for a small open economy. Econ. Modell. 19, 509–529.

Lanne, M., Saikkonen, P., 2003. Modeling the U.S. short-term interest rate by mixture autoregressive processes. J. Financial Econometrics 1, 96–125.

Malz, A., 1996. Using option prices to estimate realignment probabilities in the European monetary system. J. Int. Money Finance 15, 717–748.

Melick, W., Thomas, C., 1997. Recovering an asset's implied PDF from option prices: an application to crude oil during the gulf crisis. J. Financial Quantit. Anal. 32, 91–115.
Merton, R., 1976. Option pricing when underlying stock returns are discontinuous. J. Financial Econ. 3, 124–144.
Mizrach, B., 1995. Target zone models with stochastic realignments: an econometric evaluation. J. Int. Money Finance 14, 641–657.
Mizrach, B., 2002. When Did the Smart Money in Enron Lose Its' Smirk? Rutgers University Working Paper #2002-24.
Neely, C.J., 1994. Realignments of target zone exchange rate systems: what do we know?. Federal Reserve Bank of St. Louis Rev. 76, 23–34.
Neely, C.J., 1999. Target zones and conditional volatility: the role of realignments. J. Emp. Finance 6, 177–192.
Palm, F.C., Vlaar, P.J.G., 1997. Simple diagnostic procedures for modeling financial time series. Allgemeines Stat. Arch. 81, 85–101.
Ritchey, R., 1990. Call option valuation for discrete normal mixtures. J. Financial Res. 13, 285–295.
Shimko, D., 1993. Bounds of probability. Risk 6, 33–37.
Skouras, S., 2001. Decisionmetrics: A Decision-Based Approach To Econometric Modelling, Santa Fe Institute Working Paper No. 01-10-64.
Stein, E.M., Stein, J.C., 1991. Stock price distributions with stochastic volatility: an analytic approach. Rev. Financial Stud. 4, 727–752.
Vlaar, P.J.G., Palm, F.C., 1993. The message in weekly exchange rates in the European monetary system: mean reversion, conditional heteroskedasticity, and jumps. J. Business Econ. Stat. 11, 351–360.
Wichura, M.J., 1988. Algorithm AS 241: the percentage points of the normal distribution. Appl. Stat. 37, 477–484.
Wiggins, J.B., 1987. Option values under stochastic volatility: theory and empirical estimates. J. Financial Econ. 19, 351–372.
Wong, C.S., Li, W.K., 2001. On a logistic mixture autoregressive model. Biometrika 88, 833–846.