Dipartimento di Statistica
"Giuseppe Parenti"

Indirect estimation of α-stable stochastic volatility models

Marco J. Lombardi, Giorgio Calzolari

Università degli Studi di Firenze

Indirect estimation of α-stable stochastic volatility models

Marco J. Lombardi
European Central Bank

Giorgio Calzolari[*]
Università di Firenze

May 3, 2006

Abstract

The α-stable family of distributions constitutes a generalization of the Gaussian distribution, allowing for asymmetry and thicker tails. Its many useful properties, including a central limit theorem, are especially appreciated in the financial field. However, estimation difficulties have up to now hindered its diffusion among practitioners. In this paper we propose an indirect estimation approach to stochastic volatility models with α-stable innovations that exploits, as auxiliary model, a GARCH(1,1) with t-distributed innovations. We consider both cases of heavy-tailed noise in the returns or in the volatility. The approach is illustrated by means of a detailed simulation study and an application to currency crises.

1 Introduction

Heavy-tailedness of asset returns is one of the most prominent stylized facts in finance: studies questioning the Gaussian random-walk hypothesis and suggesting the use of α-stable distributions for the modelling of financial returns started appearing in the sixties, following the seminal works by Mandelbrot (1963) and Fama (1965). The features and the analytic properties of α-stable distributions are especially appreciated in the financial field: the fact that the family is closed under linear combination helps in portfolio analysis (Ortobelli Lozza, Huber & Schwartz 2002) and risk management (Khindanova, Rachev & Schwartz 2001); the possibility to accommodate for skewness and heavy tails allows to appropriately measure risk, avoiding to underestimate the probability of extreme losses; finally, the presence of a central limit theorem constitutes a theoretical basis which should lead to prefer the α-stable family over other heavy-tailed alternatives: since asset returns are commonly thought of as the result of the

aggregation of the asset allocation decisions of the market participants, the resulting distributions should arise, in the limit, from a central limit theorem.

However, practical application of models based on α-stable distributions has been hindered by estimation difficulties: the α-stable density function cannot be expressed in a closed form except for very few cases. This difficulty, coupled with the fact that moments of order greater than or equal to α do not exist whenever $\alpha \neq 2$, has made impossible the use of standard estimation methods such as maximum likelihood and method of moments. Researchers have thus proposed alternative estimation procedures, mainly based on quantiles (McCulloch 1986) or on the empirical characteristic function (Koutrouvelis 1980, Kogon & Williams 1998); those methods however can only estimate the parameters of the distributions, so that dealing with more complex models (both linear and nonlinear) based on α-stable disturbances would require a two-step estimation approach. In the recent years the availability of fast computing machines has made possible to employ computationally-intensive estimation; in particular, likelihood-based inference has been carried out by approximating the density with the FFT of the characteristic function (Mittnik, Doganoglu & Chenyao 1999) or with numerical quadrature (Nolan 1997); two indirect approaches have been proposed by Lombardi & Calzolari (2004) and Garcia, Renault & Veredas (2004). The Bayesian approach has also benefited from the introduction of modern computers: simulation-based MCMC methods have been proposed by Buckle (1995), Qiou & Ravishanker (1998), Lombardi (2006) and Casarin (2004*b*).

Several studies have highlighted that the heavy-tailedness of asset returns can be the consequence of conditional heteroscedasticity (Engle 1982). ARCH models have thus become very popular, given their ability to account for volatility clustering and, implicitly, heavy-tailedness at the same time. The introduction of this alternative way to deal with heavy-tailedness, coupled with the above-mentioned estimation difficulties, has somehow dampened the academic interest in α-stable distributions. A notable exception was an interesting analysis of the relation between GARCH models and α-stable distributions proposed by de Vries (1991) and Ghose & Kroner (1995). However, it must be remarked that, in practice, GARCH models are seldom able to accommodate for the excess of kurtosis: the standardized residuals are often found to be still leptokurtic. Thus, practitioners often use GARCH models with t-distributed innovations (Fiorentini, Sentana & Calzolari 2003), although it has to be remarked that GARCH models with α-stable innovations have been proposed by McCulloch (1985), Liu & Brorsen (1995) and Panorska, Mittnik, & Rachev (1995).

A widely employed alternative to ARCH-type models is represented by stochastic volatility models (Taylor 1986): their close relationship with continuous-time diffusions makes them particularly appreciated in what they can bridge the most recent results of the theoretical finance literature. However, even in the simplest Gaussian SV case, the estimation is complicated by the latent structure of the model; indirect estimation approach were proposed by Gallant, Hsieh & Tauchen (1997), Monfardini (1998) and Calzolari, Fiorentini & Sentana (2004)

In this paper we show how stochastic volatility models with α-stable innovations can be estimated using an indirect estimation approach. Traditionally, heavy tails in the setting of SV models have been accounted for using t-distributed innovations (Chib, Nardari & Shephard 2002); a recent exception is Casarin (2004*a*), where a simulation-

based Bayesian approach to a stochastic volatility model with symmetric α-stable noise is proposed. Our opinion is that the use of α-stable distributions should be preferred: first because of the presence of the generalized central limit theorem, second because of the availability of formal option pricing schemes. For example, Hurst, Platen & Rachev (1999) consider log-symmetric prices for the assets, and in McCulloch (2003) and Carr & Wu (2003) the symmetry assumption is relaxed; Cartea & Howison (2006) also consider the more appealing case of time-varying volatility with α-stable shocks.

As auxiliary model, we will employ a GARCH model with skew-t-distributed innovations; this in a sense mimics the approach followed by Calzolari et al. (2004) for the indirect estimation of stochastic volatility models. We will first examine the compliance of the auxiliary model with the conditions required to ensure consistency; then, a detailed simulation study aimed at assessing the properties of the estimators will be conducted. An application to exchange rate crises will conclude the paper.

2 α-Stable distributions

The α-stable family of distributions is identified by means of the characteristic function

$$
\phi_1(t) = \begin{cases} \exp\left\{ \mathrm{i}\delta t - \gamma^\alpha |t|^\alpha \left[1 - \mathrm{i}\tilde{\beta}\mathrm{sgn}(t)\tan\frac{\pi\alpha}{2} \right] \right\} & \text{if } \alpha \neq 1 \\ \exp\left\{ \mathrm{i}\delta t - \gamma|t| \left[1 + \mathrm{i}\tilde{\beta}\frac{2}{\pi}\mathrm{sgn}(t)\ln|t| \right] \right\} & \text{if } \alpha = 1 \end{cases}
\tag{1}
$$

which depends on four parameters: $\alpha \in (0, 2]$, measuring the tail thickness (thicker tails for smaller values of the parameter), $\tilde{\beta} \in [-1, 1]$ determining the degree and sign of asymmetry, $\gamma > 0$ (scale) and $\delta_1 \in \mathbb{R}$ (location). The distribution will be denoted as $\mathcal{S}_1(\alpha, \tilde{\beta}, \gamma, \delta_0)$.

While the characteristic function (1) has a quite manageable expression and can straightforwardly produce several interesting analytic results, it unfortunately has a major drawback for what concerns estimation and inferential purposes: it is not continuous with respect to the parameters, having a pole at $\alpha = 1$.

An alternative way to write the characteristic function that overcomes this problem, due to Zolotarev (1986), is the following:

$$
\phi_0(t) = \begin{cases} \exp\left\{ \mathrm{i}\delta_0 t - \gamma^\alpha |t|^\alpha \left[1 + \mathrm{i}\tilde{\beta}\tan\frac{\pi\alpha}{2}\mathrm{sgn}(t)\left(|\gamma t|^{1-\alpha} - 1\right) \right] \right\} & \text{if } \alpha \neq 1 \\ \exp\left\{ \mathrm{i}\delta_0 t - \gamma|t| \left[1 + \mathrm{i}\tilde{\beta}\frac{2}{\pi}\mathrm{sgn}(t)\ln(\gamma|t|) \right] \right\} & \text{if } \alpha = 1 \end{cases}
\tag{2}
$$

In this case, the distribution will be denoted as $\mathcal{S}_0(\alpha, \tilde{\beta}, \gamma, \delta_0)$. The formulation of the characteristic function is, in this case, more cumbersome, and the analytic properties have less intuitive meaning; but it is much more useful for statistical purposes and, unless otherwise stated, we will refer to it in the following. The only parameter that needs to be "translated" according to the following relationship is δ:

$$
\delta_0 = \begin{cases} \delta_1 + \tilde{\beta}\gamma\tan\frac{\pi\alpha}{2} & \text{if } \alpha \neq 1 \\ \delta_1 + \tilde{\beta}\frac{2}{\pi}\gamma\ln\gamma & \text{if } \alpha = 1 \end{cases}
\tag{3}
$$

3

On the basis of the above equations, a $\mathcal{S}_1(\alpha, \tilde{\beta}, 1, 0)$ distribution corresponds to a $\mathcal{S}_0(\alpha, \tilde{\beta}, 1, -\tilde{\beta}\gamma \tan \frac{\pi\alpha}{2})$, provided that $\alpha \neq 1$.

Unfortunately, (1) and (2) cannot be analytically inverted to yield a closed-form density function except for very few cases: $\alpha = 2$, corresponding to the normal distribution, $\alpha = 1$ and $\tilde{\beta} = 0$, yielding the Cauchy distribution, and $\alpha = \frac{1}{2}, \tilde{\beta} = \pm 1$ for the Lévy distribution. We remark that, in the case of the normal distribution, $\tilde{\beta}$ becomes unidentified.

Despite the computational burden associated with the evaluation of the probability density function, stably distributed pseudo-random numbers can be straightforwardly simulated using the algorithm proposed in Chambers, Mallows & Stuck (1976) and Chambers, Mallows & Stuck (1987). Let W be a random variable with exponential distribution of mean 1 and let U be an uniformly distributed random variable on $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$. Furthermore, let $\zeta = \arctan\left(\tilde{\beta} \tan \frac{\pi\alpha}{2}/\alpha\right)$. Then

$$
Z = \begin{cases}
\dfrac{\sin \alpha(\zeta + U)}{\sqrt[\alpha]{\cos \alpha\zeta \cos U}} \left[\dfrac{\cos\left(\alpha\zeta + \alpha U - U\right)}{W}\right]^{\frac{1-\alpha}{\alpha}} & \text{if } \alpha \neq 1 \\[3mm]
\dfrac{2}{\pi}\left[\left(\dfrac{\pi}{2} + \tilde{\beta}U\right)\tan U - \tilde{\beta}\ln\dfrac{\frac{\pi}{2}W\cos U}{\frac{\pi}{2} + \beta U}\right] & \text{if } \alpha = 1.
\end{cases}
\tag{4}
$$

has $\mathcal{S}_0(\alpha, \tilde{\beta}, 1, 0)$ distribution. Random numbers for the general case containing also the location and scale parameters δ and γ may be straightforwardly obtained exploiting the fact that, if $X \sim \mathcal{S}_1(\alpha, \tilde{\beta}, \gamma, \delta)$, then $Z = \frac{X-\delta}{\gamma} \sim \mathcal{S}_1(\alpha, \tilde{\beta}, 1, 0)$. Similarly, random numbers from an α-stable distribution expressed in parametrization (1) can be readily obtained using (3). In what follows, we will often omit the subscript and the parameters γ and δ; we will use the shorthand notation

$$
\mathcal{S}(\alpha, \tilde{\beta}) = \mathcal{S}_1(\alpha, \tilde{\beta}, 1, 0).
$$

2.1 α-Stable stochastic volatility models

Stochastic volatility models have been studied extensively by Taylor (1986) as an alternative to ARCH models. Their main advantage is that they can be regarded as the discrete time analog of the continuous time stochastic processes for instantaneous log volatility frequently used in the theoretical finance literature. A standard stochastic volatility model is composed of a latent volatility equation and of an observed return equation:

$$
\begin{aligned}
\ln h_t &= \delta + \varphi \ln h_{t-1} + \sigma_h v_t, \\
r_t &= w_t \sqrt{h_t};
\end{aligned}
\tag{5}
$$

in the most simple case, the noise terms v_t and w_t are assumed to be Gaussian and uncorrelated. The latent structure of the model makes inference troublesome, as the likelihood cannot be expressed in closed form. The estimation is therefore carried out using QML and the Kalman filter (Harvey, Ruiz & Shephard 1994), Bayesian MCMC

techniques (Jacquier, Polson & Rossi 1994, Kim, Shephard & Chib 1998) or indirect approaches (Gallant et al. 1997, Monfardini 1998, Calzolari et al. 2004).

The Gaussian assumption is often unsatisfactory for applied purposes, as observed series tend to be heavy-tailed and display discontinuities. Therefore, several studies (Chib et al. 2002, Jacquier, Polson & Rossi 2004) have considered the possibility to employ t-distributed innovations in the return equation. Nevertheless, the use of t distributions is arbitrary and in a sense dampens most of the theoretical appeal of SV models: the resulting continuous-time process is not anymore a Brownian motion, which is a requirement of most of the theoretical finance literature.

In order to account for possible discontinuities which may result from surprise events, it has also been proposed to introduce jump components in the observation equation or in the variance equation; the role of jumps in the structure of the model is discussed in detail by Eraker, Johannes & Polson (2003). The decision to include jumps in the observation or in the state equation (or in both) is not neutral and has different effects on the pattern of the volatility. Most of the research up to now has concentrated on heavy tails and jumps in the observation equation. This implies that shocks are transient and, contrary to what would happen in a GARCH framework, have no subsequent impact on volatility. Therefore, using Gaussian innovations in the volatility equation, in order to achieve a rapid increase in volatility similar to that observed in real datasets, one would need an unlikely long sequence of positive innovations. Instead, jumps in volatility have been widely documented (Bates 2000): in order to account for that, one can either introduce jumps or consider distributions with heavier tails.

In this setting, α-stable distributions are peculiar in what they generate processes with discontinuities (McCulloch 1978) that can, in a way, be interpreted as jumps. It is also interesting to remark that the inclusion of a jump component yields an unconditional mixture of normals representation; this parallels with the fact that symmetric α-stable distributions can actually be represented as scale mixtures of normals.

To the best of our knowledge, the first explicit appearance of α-stable distributions in the setting of stochastic volatility model is in Casarin (2004a), in which a stochastic volatility model with symmetric α-stable innovations in the returns equation is estimated via Sequential Monte Carlo methods:

$$\begin{aligned} \ln h_t &= \delta + \varphi \ln h_{t-1} + \sigma_h v_t, & v_t &\sim \mathcal{N}(0,1), \\ r_t &= w_t \sqrt{h_t}, & w_t &\sim \mathcal{S}(\alpha,0). \end{aligned} \qquad (6)$$

Such a model represents a very natural way to incorporate α-stable distributions but it presents a number of inconvenients. In the first place, it does not allow heavy tails to interact with the volatility pattern: contrary to GARCH models, in a SV framework the evolution of the volatility is determined uniquely by its past values. Therefore, the spikes observed in the returns have no impact on the volatility pattern. This can in certain cases be a positive feature, but our feeling is that in the majority of the situations extreme returns are supposed to impact positively on the volatility. A second shortcoming is that allowing for α-stable innovations in the returns equation is bound to deteriorate the quality of the estimates, as heavy tails make more difficult to reconstruct the latent volatility pattern. Another remark is that, whereas returns are conditionally α-stable, their unconditional distribution is a scale mixture of α-stable distributions

5

with mixing weights given by the square root of a log-normal distribution; this yields an unknown distributional form. We will refer to this model with the shorthand notation SVSR (Stochastic Volatility with Stable Returns).

de Vries (1991) examines the relation between GARCH and stable processes and proposes an α-stable quasi-GARCH specification:

$$
\begin{aligned}
h_t &= \varphi h_{t-1} + \sigma_h v_{t-1}, & v_t &\sim \mathcal{D}(0,1), \\
r_t &= w_t \sqrt{h_t}, & w_t &\sim \mathcal{D}(0,1),
\end{aligned}
\tag{7}
$$

where $\mathcal{D}(0,1)$ denotes a generic distribution with zero mean and unit variance. When one sets $v_{t-1} = w_{t-1}$, the above specification is very similar to a GARCH (1, 1). If instead one employs

$$
\begin{aligned}
h_t &= \varphi h_{t-1} + \sigma_h v_{t-1}, & v_t &\sim \mathcal{S}\left(\tfrac{\alpha}{2}, 1\right), \\
r_t &= w_t \sqrt{h_t}, & w_t &\sim \mathcal{N}(0,1),
\end{aligned}
\tag{8}
$$

it is shown (de Vries 1991) that the unconditional distribution of the returns is α-stable with characteristic exponent α. The model is indicated as Stable Subordination with Conditional Scaling (SSCS), but can in actual facts thought of as a particular form of stochastic volatility. This model has very nice theoretical properties, as returns are conditionally Gaussian and unconditionally stable. However, it must be remarked that having perfectly skewed innovations in the volatility equation may be undesirable, as shocks are necessarily positive, and may also yield unrealistic patterns of volatility. A consequence of the positivity of the shocks is that the estimates on real data indicate a very small degree of persistence[1]: since negative innovations are not allowed, in order for the volatility to decrease realistically after a shock it is necessary to have a low φ. We furthermore remark that the specification above does not include a constant in the volatility equation, and the unconditional distribution of h_t is $\mathcal{S}(\tfrac{\alpha}{2}, \gamma, 0)$, where

$$
\gamma = \varphi \left[\frac{1}{1 - \sigma_h^{\frac{\alpha}{2}}} \right]^{\frac{2}{\alpha}}.
$$

It follows that, contrary to what happens in traditional stochastic volatility specifications, in this case the scale of the returns is determined uniquely by φ and σ_h, which is definitely not desirable.

An alternative approach could be to consider a more traditional SV specification allowing for α-stable innovations in the volatility equation, but without necessarily restricting α to be less than one:

$$
\begin{aligned}
\ln h_t &= \delta + \varphi \ln h_{t-1} + \sigma_h v_t, & v_t &\sim \mathcal{S}(\alpha, 1), \\
r_t &= w_t \sqrt{h_t}, & w_t &\sim \mathcal{N}(0,1).
\end{aligned}
\tag{9}
$$

Such a model also allows for negative shocks in the volatility whenever $\alpha > 1$ and yields a smoother pattern of volatility. Of course, positive jumps that tend to be less

---

[1]The estimates reported in de Vries (1991) point to values between 0.15 and 0.45, and our application to real data yields similar results.
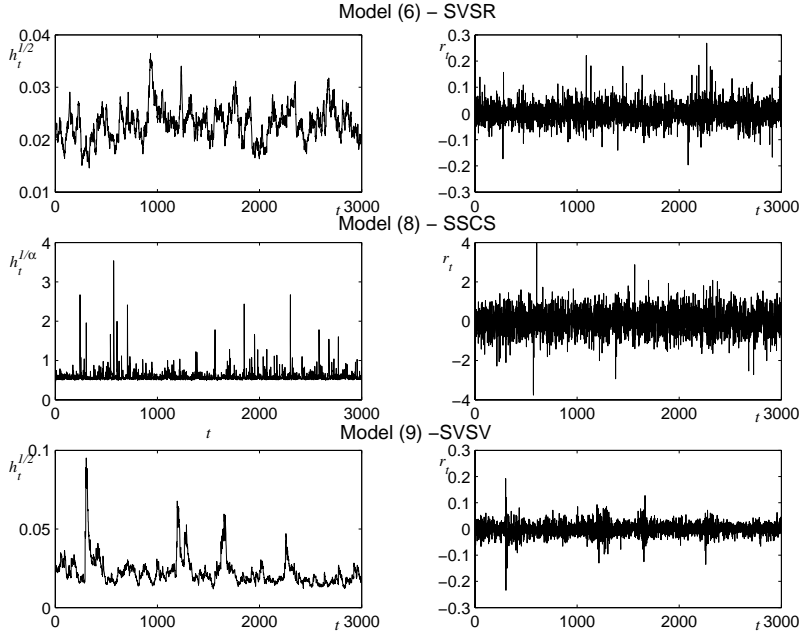
Figure 1: Simulated volatility (left panels) and returns (right panels) for model (6) (top panel, $\alpha = 1.9$), model (8) (central panel, $\alpha = 1.8$) and (9) (bottom panel, $\alpha = 1.7$).

frequent as α approaches 2, since the skewness parameter $\tilde{\beta}$ loses relevance as we move towards a Gaussian distribution. In this case, the returns have conditional Gaussian distribution, but their unconditional distribution is unknown, as it is given by a scale mixture of normals with weights given by a log-stable distribution. This model will be denoted a Stochastic Volatility with Stable Volatility (SVSV).

To illustrate the three different approaches outlined above and convince the reader that α-stable distributions can appropriately model jumps and generate plausible patterns of volatility, we report (Figure 1) the simulated paths of volatility and returns for each of the models under scrutiny. For the SVSR model, α was fixed to 1.9 (with $\tilde{\beta} = 0$), for model SSCS we employ the same α and $\tilde{\beta} = 1$ and finally in model SVSV we employed $\alpha = 1.7$ (with $\tilde{\beta} = 1$). The parameters of the volatility equation for models SVSV and SVSR were fixed to $\delta = -0.15$, $\varphi = 0.98$, $\sigma_h = 0.06$, whereas for model SSCS we use $\varphi = 0.35$ and $\sigma_h = 0.015$, values similar to the estimates reported in de Vries (1991) on exchange rate data. We observe that the inclusion of the heavy-tailed noise in the observation equation generates spikes in the pattern of returns, but is unable to yield a realistic pattern of the volatility. The SSCS model also yields a not very persistent pattern of volatility, with very frequent spikes in returns. On the other hand, the SVSV model yields in our opinion more realistic patterns of volatility, with shocks rapidly increasing the volatility and then fading away smoothly. The volatility clustering pattern in the returns is much more visible, but no transient shocks are allowed in the SVSV case.

7

To sum up, the use of α-stable distributions should be in our opinion preferable over other heavy-tailed distributions and/or jumps for three different reasons:

– The presence of the central limit theorem should justify the arising of such a kind of distributions in the volatility setting: if (log-)volatility is thought of as the stream of news arriving to the market, then it is natural to assume it to be composed of a large number of individual contributions of each event;

– In continuous time, the α-stable generalization of the Brownian motion (the α-Stable Lévy motion) has been studied extensively (Samorodnitsky & Taqqu 1994, Janicki & Weron 1994). Finance theorists have widely exploited the α-stable assumption and several models are available, ranging from asset allocation to option pricing (see, for an excellent survey, McCulloch (1996));

– Unlike Brownian motions, Lévy motions are not almost surely continuous, but instead they are almost surely dense with discontinuities (McCulloch 1978); this means they can actually account for empirically-observed jumps without having to include a separate jump component.

3 The indirect estimation approach

The indirect estimation (Gouriéroux, Monfort & Renault 1993) is an inferential approach which is suitable for situations where the estimation of the statistical model of interest is too difficult to be performed directly while it is straightforward to produce simulated values from the same model. It was first motivated by econometric models with latent variables, but it can be applied in virtually every situation in which the direct maximization of the likelihood function turns out to be difficult.

The principle underlying "indirect inference" (Gouriéroux et al. 1993) is very simple: suppose we have a sample of T observations $\mathbf{y}$ and a model whose likelihood function $\mathcal{L}^\star(\mathbf{y}; \theta)$ is difficult to handle and maximize; the model could also depend on a matrix of explanatory variables $\mathbf{X}$. The maximum likelihood estimate of $\theta \in \Theta$, given by

$$\hat{\theta} = \arg\max_{\theta \in \Theta} \sum_{t=1}^{T} \ln \mathcal{L}^\star(\theta; y_t),$$

is thus unavailable. Let us now take an alternative model, depending on a parameter vector $\zeta \in Z$, which will be indicated as *auxiliary model*, easier to handle, and suppose we decide to use it in the place of the original one. Since the model is misspecified, the quasi-ML estimator

$$\hat{\zeta} = \arg\max_{\zeta \in Z} \sum_{t=1}^{T} \ln \tilde{\mathcal{L}}(\zeta; y_t),$$

is not necessarily consistent: the idea is to exploit simulations performed under the original model to correct for inconsistency.

The first step consists of computing the quasi maximum likelihood estimate of ζ, which will be denoted as $\hat{\zeta}$. Next, one simulates a set of S vectors of size T from

the original model on the basis of an arbitrary parameter vector $\hat{\theta}^{(0)}$. Let us denote the observations of these vectors as $y_t^s(\hat{\theta}^{(0)})$. The simulated values are then estimated using the auxiliary model, yielding

$$\tilde{\zeta}(\hat{\theta}^{(0)}) = \arg\max_{\zeta \in Z} \sum_{s=1}^{S} \sum_{t=1}^{T} \ln \tilde{\mathcal{L}} \left[\zeta; y_t^s(\hat{\theta}^{(0)}) \right]. \tag{10}$$

The idea is to numerically update the initial guess $\hat{\theta}^{(0)}$ in order to minimize the distance

$$\left[\hat{\zeta} - \tilde{\zeta}(\theta) \right]' \Omega \left[\hat{\zeta} - \tilde{\zeta}(\theta) \right], \tag{11}$$

where Ω is a symmetric nonnegative matrix defining the metric.

An alternative but similar approach, leading to the so-called EMM (Gallant & Tauchen 1996), considers directly the score function of the auxiliary model:

$$\sum_{t=1}^{T} \frac{\partial \ln \tilde{\mathcal{L}}(\zeta; y_t)}{\partial \zeta}, \tag{12}$$

which is clearly zero for the quasi-maximum likelihood estimator $\hat{\zeta}$. The idea is to make as close as possible to zero the score computed on the simulated observations, namely

$$\arg\min_{\theta} \left\{ \sum_{s=1}^{S} \sum_{t=1}^{T} \frac{\partial \ln \tilde{\mathcal{L}} \left[\zeta; y_t^s(\theta) \right]}{\partial \zeta} \Big|_{\zeta=\hat{\zeta}} \right\}' \Xi \left\{ \sum_{s=1}^{S} \sum_{t=1}^{T} \frac{\partial \ln \tilde{\mathcal{L}} \left[\zeta; y_t^s(\theta) \right]}{\partial \zeta} \Big|_{\zeta=\hat{\zeta}} \right\}, \tag{13}$$

where Ξ is a symmetric nonnegative definite matrix. This approach is especially useful when an analytic expression for the gradient of the auxiliary model is available, since it allows us to avoid the numerical optimization routine for the computation of the $\hat{\zeta}(\theta)$s.

Indirect estimators are consistent and asymptotically normal under certain regularity conditions. The most difficult one to establish is that the *binding function*, that is the function that maps the parameter (sub-)space of the auxiliary model onto the parameter space of the true model, is one-to-one. In general, the binding function cannot be expressed analytically and the above condition needs to be verified numerically. It is clear that the choice of the auxiliary model is crucial for the successful implementation of the algorithm: see for instance Heggland & Frigessi (2004) for further details concerning the desirable properties of the auxiliary model for the "just-identified" case and Gallant & Tauchen (1996) for a particular and "over-identified" choice of the auxiliary model producing efficient estimators.

Once one manages to specify an adequate auxiliary model, indirect estimators for the parameters of α-stable distributions can be readily implemented and exploited by relying on the pseudo-random number generator of Chambers et al. (1976).

3.1 Indirect estimation of stochastic volatility models

We will here discuss the indirect estimation of the three stochastic volatility specifications outlined above, namely the SVSR (6), the SSCS (8) and the SVSV (9). The

characteristic exponent α will be constrained inside teh interval $(1, 2)$ for models SVSR and SVSV and inside $(0, 1)$ for model SSCS; the persistence parameter φ will also be constrained inside $(-1, 1)$ to ensure stationarity of the volatility.

Calzolari et al. (2004) employ a GARCH with Student's t innovations as an over-identified auxiliary model for the indirect estimation of a Gaussian stochastic volatility model. This choice could be fruitful also in our case, as the degrees of freedom of the t distribution are naturally linked to the characteristic exponent α in determining the tail thickness.

We will express the t distribution in terms of $\eta = \nu^{-1}$, where ν are the degrees of freedom, and we constrain η on $(0.01, 1)$. On the lower bound, the auxiliary distribution is very close to the Gaussian, while on the upper bound it corresponds to an α-stable distribution with $\alpha = 1$. It is important to remark that here, contrary to Calzolari et al. (2004), we do not employ the standardized version of the t distribution and we allow for values of η that give rise to distributions with infinite variance ($\eta \geq 0.5$). This choice was made with the goal of avoiding the degrees of freedom to clash on their lower bound because of the heavier-tailedness of the true model. The only caveat is that h_t cannot be interpreted anymore as a conditional variance: when $\nu < 2$ the variance is infinite and when $\nu \geq 2$ the conditional variance will be determined as $\frac{h_t}{\nu-2}$. Therefore, we will see h_t as a scale parameter. The fact that the auxiliary model may produce infinite variances is not a limitation in our opinion: in the case of α-stable noise in the observation equation (6), the variance of the returns is infinite in the true model as well and the volatility equation yields only a time-varying scale parameter. When the α-stable noise is employed in the variance equation, on the other hand, the returns do have finite variance – given by h_t – but the variance of the volatility, which is related to the fourth moment of returns, is infinite because of the α-stable innovations. Therefore the parameter σ_h would not be interpreted as the variance of the (log-)volatility, but rather as a simple scale coefficient.

To sum up, the auxiliary model will be:

$$
\begin{aligned}
h_t &= \omega + \psi r_{t-1}^2 + \beta h_{t-1} & (14) \\
r_t &= u_t \sqrt{h_t}, & u_t \sim t_{1/\eta};
\end{aligned}
$$

In order to ensure non-negativity of the conditional variance in the auxiliary model, we will impose $\omega > 0$, $\psi > 0$ and $\beta > 0$ (Nelson & Cao 1992); to safeguard against weird asymptotic behavior of the QML estimator (Lumsdaine 1996) we will also enforce the constraint $\psi + \beta < 1$. In the case $\psi = 0$, it turns out (Andrews 1999) that β gets asymptotically unidentified; we will therefore impose $\psi > \varepsilon$, with ε arbitrarily small.

Sections of the binding functions for models (6), (8) and (9) are plotted in Figure 3, 4 and 5, respectively, and signal that such auxiliary model could be appropriate in all cases. In particular, we observe that the persistence parameter φ is, as expected, linked to both β and ψ. Also the scale of the innovations in the volatility equation depend on the ARCH and GARCH coefficients, albeit in a different fashion: an increase in σ_h causes an increase in β and a decrease in ψ. This effect is caused by the fact that bigger innovations in the volatility in the SV model cannot be captured by the ARCH coefficient, which is related to innovations in the returns, and need to be discounted by

10

the past volatility dynamics which is controlled via the GARCH coefficient.

We also remark that the binding function seems to behave quite similarly for all three models; some differences arise with respect to the sections $\alpha - \nu$. The relation between the tail-thickness parameters is more smooth in model (9) and the degrees of freedom of the auxiliary model do not appear to explode to ∞. This could be caused by the fact that, when allowing for heavy tails in the volatility equation, also very small shocks may, by means of the multiplicative effect, yield a pattern in the returns that can be captured only via tails heavier than the Gaussian.

3.2   Simulation results

In order to evaluate the performance of the proposed estimator, we have conducted a simulation exercise; all the results are based on a set of 1000 replications of the indirect estimation, with $S = 10$. We will use three different sample sizes: 1000, 3000 and 5000 observations.

For models SVSR and SVSV, which are more traditional stochastic volatility specifications, we will use two sets parameters in the stochastic volatility equation with various choices of the characteristic index α, namely one that roughly matches typical values obtained on weekly returns ($\delta = -0.7$, $\varphi = 0.9$, $\sigma_h = 0.35$) and one roughly matching daily returns ($\delta = -0.15$, $\varphi = 0.98$, $\sigma_h = 0.06$); the simulation design is therefore similar to that of Jacquier et al. (1994) and Calzolari et al. (2004). For the model SSCS, we will instead use values similar to the estimates reported by de Vries (1991), namely $\varphi = 0.3$ and $\sigma_h = 0.015$.

As remarked by Jacquier et al. (1994) in the SV setting, a parameter set yielding a low signal-to-noise-ratio (SNR) is bound to give rise to estimation difficulties. In the α-stable case, however, some caveats are necessary. In the Gaussian SV model, the (power) SNR is:

$$\mathrm{SNR}_2 = \frac{\mathrm{Var}(h_t)}{\mathrm{Var}(r_t)} = \frac{\mathrm{Var}(h_t)}{\mathrm{E}(h_t)^2}.$$

Employing the above expression under model (6) yields a SNR of zero, whereas for models (8) and (9) the SNR becomes infinite. An alternative which is often use in signal processing is to consider

$$\mathrm{SNR}_{\mathrm{A}} = \frac{\mathrm{E}|h_t|}{\mathrm{E}|r_t|}.$$

When the α-stable noise is in the observation equation (6), this choice would lead to a nonzero quantity as long as $\alpha > 1$. In the other two cases, however, the numerator would be the expected value of, respectively, a perfectly skewed stable and a log-stable random variable, which are unfortunately infinite.

In Tables 1 and 2 we report, respectively, the outcome of the simulation experiment on the model (6) with the first and the second parameter set and for different values of α. We remind that in this case we only consider symmetric distributions, therefore $\tilde{\beta}$ is fixed to 0. First of all, we remark that in the case with high SNR (Table 1), the estimation went smoothly and the results are very satisfactory. While the parameter α is estimated with strong precision, the parameters in the volatility equation have larger

Table 1: Monte Carlo mean and standard error (in parentheses) for the first set of parameters ($\delta = -0.7$, $\varphi = 0.9$, $\sigma_h = 0.35$) and various values of α, SVSR model (6). The rows "*Time*" report the average time to convergence (in seconds) of one iteration.

	$\alpha = 1.8$			$\alpha = 1.9$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.8279	1.8145	1.8096	1.9104	1.9052	1.9040
	(0.0848)	(0.0551)	(0.0447)	(0.0593)	(0.0384)	(0.0304)
δ	−1.0493	−0.8785	−0.8351	−0.8682	−0.7528	−0.7449
	(1.0936)	(0.7286)	(0.6373)	(0.6778)	(0.3757)	(0.3076)
φ	0.8502	0.8744	0.8807	0.8762	0.8924	0.8936
	(0.1559)	(0.1042)	(0.0909)	(0.0964)	(0.0535)	(0.0439)
σ_h	0.4038	0.3816	0.3720	0.3737	0.3590	0.3577
	(0.2599)	(0.1829)	(0.1569)	(0.1710)	(0.1071)	(0.0891)
Time	6.5613	6.2942	5.7529	3.1315	3.0858	2.7685
	$\alpha = 1.95$			$\alpha = 1.99$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.9515	1.9526	1.9525	1.9798	1.9851	1.9870
	(0.0440)	(0.0304)	(0.0249)	(0.0279)	(0.0177)	(0.0147)
δ	−0.8019	−0.7349	−0.7299	−0.7201	−0.6967	−0.7017
	(0.4879)	(0.2687)	(0.2148)	(0.2898)	(0.1772)	(0.1380)
φ	0.8856	0.8950	0.8958	0.8974	0.9005	0.8998
	(0.0696)	(0.0382)	(0.0306)	(0.0411)	(0.0251)	(0.0196)
σ_h	0.3609	0.3564	0.3555	0.3377	0.3437	0.3465
	(0.1234)	(0.0806)	(0.0661)	(0.0792)	(0.0535)	(0.0426)
Time	4.1063	4.5971	4.7299	4.9718	10.8701	15.9124

standard errors, especially in the presence of heavier tails. This is of course consistent with the impact of heavy tails on the SNR we have discussed above. For what concerns the execution time, we notice a weird effect: for smaller values of α the algorithm takes less time to estimate larger sample sizes. This is related to the fact that, when heavy tails are involved, having more observations implies a clearer picture of the tail behavior: therefore, working with smaller sample sizes may imply convergence difficulties. On the other hand, we remark that, not surprisingly, slow convergence arises also when α is very close to 2. In that case, since no heavy-tailed behavior has to be observed, the sample size constitutes an hinderance and considerably increases the computational time.

When the lower SNR is concerned (Table 2), the performance of the indirect estimator worsens dramatically. While the performance of the estimator for α remains satisfactory, we have found a considerable distortion and inefficiency for the estimators of the parameters in the volatility equation. As we have previously remarked, smaller values of α do impact negatively on the performance of the estimator as they imply an even smaller SNR_A. A control experiment[2] aimed at assessing the asymptotic

---

[2]Results are not reported here for the sake of brevity, but are available upon request.

Table 2: Monte Carlo mean and standard error (in parentheses) for the second set of parameters ($\delta = -0.15$, $\varphi = 0.98$, $\sigma_h = 0.06$) and various values of α, SVSR model (6). The rows "*Time*" report the average time to convergence (in seconds) of one iteration.

	$\alpha = 1.8$			$\alpha = 1.9$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.8238	1.8171	1.8152	1.9105	1.9087	1.9079
	(0.0697)	(0.0474)	(0.0433)	(0.0527)	(0.0340)	(0.0280)
δ	–0.5966	–0.5100	–0.4879	–0.6110	–0.4651	–0.4154
	(1.5701)	(1.4438)	(1.4170)	(1.3140)	(1.0989)	(1.0787)
φ	0.9206	0.9320	0.9351	0.9187	0.9381	0.9448
	(0.2086)	(0.1924)	(0.1884)	(0.1748)	(0.1459)	(0.1431)
σ_h	0.1006	0.0880	0.0831	0.1001	0.0871	0.0807
	(0.1625)	(0.1415)	(0.1411)	(0.1390)	(0.1215)	(0.1125)
Time	30.401	51.324	55.424	26.958	41.921	49.806
	$\alpha = 1.95$			$\alpha = 1.99$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.9557	1.9546	1.9537	1.9894	1.9907	1.9911
	(0.0388)	(0.0240)	(0.0190)	(0.0168)	(0.0111)	(0.0083)
δ	–0.6050	–0.4250	–0.3442	–0.4967	–0.3449	–0.2721
	(1.0634)	(0.8120)	(0.7156)	(0.7239)	(0.4161)	(0.3416)
φ	0.9195	0.9434	0.9542	0.9340	0.9541	0.9638
	(0.1403)	(0.1077)	(0.0948)	(0.0961)	(0.0552)	(0.0453)
σ_h	0.1047	0.0878	0.0787	0.0900	0.0862	0.0774
	(0.1212)	(0.0917)	(0.0774)	(0.0826)	(0.0507)	(0.0412)
Time	22.538	38.572	44.821	21.043	28.319	32.605

properties of the estimator confirmed indeed its consistency, but it took over 100000 observations to achieve reasonable accuracy. In this case, it is interesting to remark that the weird effect of sample size on computational time highlighted above is not present: smaller values of α – implying a lower SNR – and larger sample sizes always imply an increase in the computational burden.

When one moves to consider model (9), the situation is completely changed. We report only the "difficult" case (low SNR, parameter values $\delta = -0.15$, $\varphi = 0.98$, $\sigma_h = 0.06$, $\beta = 1$ and various values of α). In this case, as we have pointed out, an heavier-tailed noise yields a clearer signal and therefore facilitates the estimation. In actual facts, we observe that in this case, the estimators of the volatility parameters are much more precise, whereas the additional uncertainty carried by the heavy-tailed noise is discounted in larger standard errors for the estimator of α.

Finally, we will consider the estimation of the SSCS model (8). In this case we will only consider $\alpha = 1.6$ and $\alpha = 1.8$, as values close to 2 have no specific interest in this specification. The estimation of this model appears quite straightforward and very quick, albeit it has to be remarked that this model has one parameter less than the

Table 3: Monte Carlo mean and standard error (in parentheses) for the second set of parameters ($\delta = -0.15$, $\varphi = 0.98$, $\sigma_h = 0.06$) and various values of α, SVSV model (9). The rows "*Time*" report the average time to convergence (in seconds) of one iteration.

	$\alpha = 1.5$			$\alpha = 1.7$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.5454	1.5119	1.5041	1.7083	1.6826	1.6992
	(0.3568)	(0.3216)	(0.2984)	(0.3546)	(0.3111)	(0.2743)
δ	–0.2747	–0.2252	–0.1961	–0.2235	–0.1614	–0.1604
	(0.5756)	(0.3133)	(0.2757)	(0.3240)	(0.1163)	(0.1046)
φ	0.9616	0.9626	0.9663	0.9691	0.9779	0.9783
	(0.0753)	(0.0695)	(0.0662)	(0.0437)	(0.0145)	(0.0132)
σ_h	0.0950	0.0806	0.0712	0.0754	0.0667	0.0667
	(0.1181)	(0.0935)	(0.0854)	(0.0570)	(0.0466)	(0.0437)
Time	13.969	46.351	61.503	11.819	36.843	53.529
	$\alpha = 1.9$			$\alpha = 1.95$		
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.7783	1.8053	1.8468	1.8020	1.8373	1.8762
	(0.3204)	(0.2632)	(0.2134)	(0.3060)	(0.2418)	(0.1877)
δ	–0.2761	–0.1715	–0.1614	–0.2961	–0.1782	–0.1631
	(0.3913)	(0.0925)	(0.0670)	(0.3532)	(0.0999)	(0.0674)
φ	0.9628	0.9770	0.9784	0.9603	0.9761	0.9782
	(0.0537)	(0.0121)	(0.0087)	(0.0461)	(0.0132)	(0.0089)
σ_h	0.0639	0.0574	0.0586	0.0617	0.0556	0.0560
	(0.0448)	(0.0305)	(0.0248)	(0.0426)	(0.0273)	(0.0220)
Time	12.482	19.632	28.461	13.713	21.043	31.899

Table 4: Monte Carlo mean and standard error (in parentheses) for parameters $\varphi = 0.3$, $\sigma_h = 0.015$ and various values of α, SSCS model (8). The rows "*Time*" report the average time to convergence (in seconds) of one iteration.

| | $\alpha = 1.6$ | | | $\alpha = 1.8$ | | |
	$T = 1000$	$T = 3000$	$T = 5000$	$T = 1000$	$T = 3000$	$T = 5000$
α	1.6125	1.6034	1.6033	1.8172	1.8048	1.8046
	(0.0486)	(0.0252)	(0.0201)	(0.0544)	(0.0298)	(0.0225)
φ	0.2965	0.2992	0.2993	0.2860	0.2956	0.2968
	(0.0679)	(0.0313)	(0.0248)	(0.1307)	(0.0668)	(0.0456)
σ_h	0.0145	0.0149	0.0149	0.0139	0.0147	0.0148
	(0.0025)	(0.0014)	(0.0011)	(0.0043)	(0.0024)	(0.0018)
Time	3.091	7.393	11.334	4.659	9.711	13.398

others. It is also interesting to remark that smaller values of α imply a clearer signal and, therefore, yield smaller standard errors.

3.3 An empirical application

In this subsection, we will apply the models we have introduced to the analysis of two currency crises. This is a rather new type of application for α-stable distributions, but we see it as very illustrative. A recent paper by Hartmann, Straetmans & de Vries (2004) actually points out that, when the exchange rate fundamentals are heavy-tailed, currency crises tend to spread across different countries. Although their analysis does not involve any kind of nonlinearity in returns, it shall nevertheless be considered as a good argument supporting the use of heavy-tailed distributions, as most of the currency crises that targeted a specific country eventually spread to its neighbors.

From a practical perspective, the patterns of volatility and returns generated during currency crises are very interesting: monetary authorities who let the exchange rate float inside a certain band are sometimes forced to switch to a free floating regime, and it may take quite a while before the monetary authority can attempt to enforce a new floating band. In this case, abandoning a managed floating regime has a strong impact of volatility, and such a shock will be absorbed as soon as the new floating band is adopted.

The exchange rate pattern of the British Pound against the Deutsche Mark is very interesting, and in our opinion is a very good example on which the models we have discussed can be illustrated. Britain did not enter the ERM as it was launched in 1979 and concentrated instead on controlling inflation via a tight monetary policy with free floating exchange rate. As inflation stabilized, between 1987 and 1988, the monetary authorities decided to enforce a managed floating regime shadowing the Deutsche Mark. Finally, in October 1990, the Pound entered the ERM with a floating band of $\pm 6\%$, instead of the usual $\pm 2.25\%$ adopted by the other currencies. During September 1992, a speculative attack targeted the British Pound and some other European currencies, eventually forcing the monetary authority to abandon the ERM.
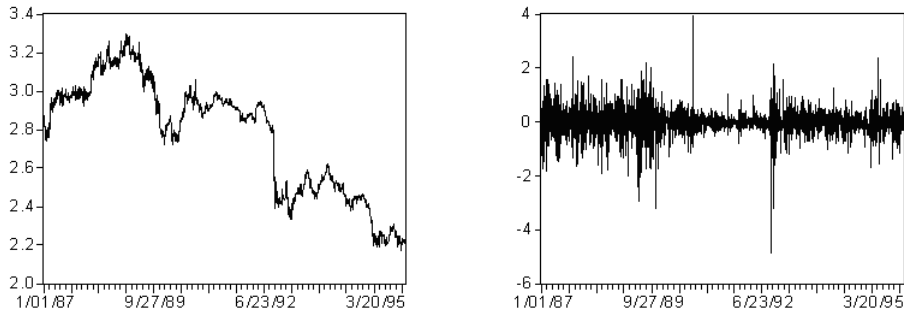
Figure 2: Exchange rate of the British pound against the Deutsche Mark (left) and returns (right), January 1, 1987 to December 31, 1995.

In a sense, the British monetary authorities have tried, for a period of about five years, to manage the exchange rate against the Deutsche Mark first by adopting an informal shadowing, and later by committing themselves to the ERM. Of course, this has exposed the Pound to strong pressure that culminated in the successful speculative attack on September 16, 1992. The evolution of the exchange rate of the British Pound against the Deutsche Mark between 1987 and 1995 (3287 daily observations) is depicted in figure 2.

First of all, we filtered the returns on the exchange rate by means of an AR(1) models in order to remove correlation in the levels. Then, the three SV specifications outlined above were estimated using the AR(1) residuals. It is very illustrative to compare the results obtained on the parameters φ and α. For what concerns φ, using a Gaussian distribution in the volatility equation, as in model SVSR, yields the higher persistence, whereas model SSCS, as already noted by de Vries (1991), has the smallest persistence of the volatility. Model SVSV, as it could be envisaged, is somehow in between. Turning our attention to α, we remark that, as expected according to the theoretical considerations made above, it has nearly the same value for models SSCS and SVSR; for model SVSV, instead, the degree of heavy-tailedness appears smaller.

4 Conclusions

In this paper, we have considered three different stochastic volatility models that allow for α-stable innovations in the returns equation and the volatility equation. An indirect estimation approach has been proposed and its properties have been examined in a simulation study. The models under consideration imply very different patterns of volatility, and we have considered their application to currency crises. Of course, the example we report should not be considered as evidence in favor of the very simple stochastic mechanism implied by the SV model with respect to more complex and somehow structural approaches. We just point out that, allowing for α-stable innovations in the returns or in the volatility equation may yield realistic patterns. A similar

16

Table 5: Estimate and standard error (in parentheses) of the parameters of models (9) and (6) for the returns of the exchange rate of the British Pound versus the Deutsche Mark; January 1, 1987 to December 31, 1995.

α	δ	φ	σ_h
Model (6), SVSR			
1.7963	-0.0076	0.9938	0.0940
(0.)	(0.)	(0.)	(0.)
Model (8), SSCS			
α		φ	σ_h
1.7940		0.2769	0.0404
(0.)		(0.)	(0.)
Model (9), SVSV			
α	δ	φ	σ_h
1.9198	-0.5877	0.5920	0.4893
(0.)	(0.)	(0.)	(0.)
Auxiliary Model			
η	ω	ψ	β
0.2862	0.0002	0.0367	0.9278
(0.)	(0.)	(0.)	(0.)

issue is in a sense addressed by Davidson (2004), who shows that patterns very similar to those observed during currency crises can be obtained by allowing a sequence of shocks in a long-memory HYGARCH model.

In our analysis, we have considered jumps either in the state or in the observation equation. A natural extension will be, as suggested also by Eraker et al. (2003), to allow for jumps in both; this could be accomplished by using bivariate α-stable distributions. In actual facts, since the normal distribution is a particular case of α-stable distribution, such a model would nest both a standard Gaussian stochastic volatility framework and the models we have considered in this paper. Representing the vector $[v_t, w_t]$ by means of a bivariate α-stable distribution would also allow to include correlation among the two innovations, as proposed for example by Jacquier et al. (2004) in the setting of t-distributed innovations. Simulating random numbers from a bivariate α-stable distribution is straightforward (Modarres & Nolan 1994), therefore an indirect estimation approach could be fruitful also in this case. One possible shortcoming of this approach could be that the tail-thickness parameter α has to be the same for both terms of the noise: this could yield unrealistic results as jumps in the volatility are usually more frequent than in returns. Another possible approach could be therefore to employ independent noise terms and induce correlation by means of copula functions. In both cases, however, the specification of an appropriate auxiliary model might be less immediate.

References

Andrews, D. W. K. (1999), 'Estimation when a parameter is on the boundary', *Econometrica* **67**, 1341–1383.

Bates, D. (2000), 'Post-'87 crash fears in the S&P 500 futures options', *Journal of Econometrics* **94**, 181–238.

Buckle, D. J. (1995), 'Bayesian inference for stable distributions', *Journal of the American Statistical Association* **90**, 605–613.

Calzolari, G., Fiorentini, G. & Sentana, E. (2004), 'Constrained indirect estimation', *Review of Economic Studies* **71**, 945–973.

Carr, P. & Wu, L. (2003), 'The finite moment log stable process and option pricing', *The Journal of Finance* **58**, 753–777.

Cartea, A. & Howison, S. (2006), Option pricing with Lévy-stable processes generated by Lévy-stable integrated variance. Birbeck Working Papers in Economics and Finance 0602, University of London.

Casarin, R. (2004*a*), Bayesian inference for generalised Markov switching stochastic volatility models. Cahier du CEREMADE 0414.

Casarin, R. (2004*b*), Bayesian inference for mixture of stable distributions. Cahier du CEREMADE 0428.

Chambers, J. M., Mallows, C. & Stuck, B. W. (1976), 'A method for simulating stable random variables', *Journal of the American Statistical Association* **71**, 340–344.

Chambers, J. M., Mallows, C. & Stuck, B. W. (1987), 'Correction to: A method for simulating stable random variables', *Journal of the American Statistical Association* **82**, 704.

Chib, S., Nardari, F. & Shephard, N. (2002), 'Markov chain Monte Carlo methods for stochastic volatility models', *Journal of Econometrics* **108**, 281–316.

Davidson, J. E. (2004), 'Moment and memory properties of linear conditional heteroscedasticity models', *Journal of Business and Economics Statistics* **22**, 16–29.

de Vries, C. G. (1991), 'On the relation between GARCH and stable processes', *Journal of Econometrics* **48**, 313–324.

Engle, R. F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation', *Econometrica* **50**, 987–1008.

Eraker, B., Johannes, M. & Polson, N. G. (2003), 'The impact of jumps in volatility and returns', *The Journal of Finance* **58**, 1269–1300.

Fama, E. F. (1965), 'The behavior of stock market prices', *Journal of Business* **38**, 34–105.

Fiorentini, G., Sentana, E. & Calzolari, G. (2003), 'Maximum likelihood estimation and inference in multivariate conditionally heteroscedastic dynamic regression models with Student t innovations', *Journal of Business and Economic Statistics* **21**, 532–546.

Gallant, A. R., Hsieh, D. & Tauchen, G. (1997), 'Estimation of stochastic volatility models with diagnostics', *Journal of Econometrics* **81**, 159–192.

Gallant, A. R. & Tauchen, G. E. (1996), 'Which moments to match?', *Econometric Theory* **12**, 657–681.

Garcia, R., Renault, E. & Veredas, A. (2004), Estimation of stable distributions by indirect inference. CIRANO working paper.

Ghose, D. & Kroner, K. F. (1995), 'The relationship between GARCH and symmetric stable processes: Finding the source of fat tails in financial data', *Journal of Empirical Finance* **2**, 225–251.

Gouriéroux, C., Monfort, A. & Renault, E. (1993), 'Indirect inference', *Journal of Applied Econometrics* **8**, S85–S118.

Hartmann, P., Straetmans, S. & de Vries, C. G. (2004), Fundamentals and joint currency crises. ECB Working paper 324, European Central Bank.

Harvey, A. C., Ruiz, E. & Shephard, N. (1994), 'Multivariate stochastic variance models', *Review of Economic Studies* **61**, 247–264.

Heggland, K. & Frigessi, A. (2004), 'Estimating functions in indirect inference', *Journal of the Royal Statistical Society B* **66**, 447–462.

Hurst, S. R., Platen, E. & Rachev, S. T. (1999), 'Option pricing for a logstable asset price model', *Mathematical and Computer Modelling* **29**, 105–119.

Jacquier, E., Polson, N. G. & Rossi, P. E. (1994), 'Bayesian analysis of stochastic volatility models (with discussion)', *Journal of Business and Economic Statistics* **12**, 371–417.

Jacquier, E., Polson, N. G. & Rossi, P. E. (2004), 'Bayesian analysis of stochastic volatility models with fat-tails and correlated errors', *Journal of Econometrics* **122**, 185–212.

Janicki, A. & Weron, A. (1994), *Simulation and Chaotic Behavior of α-Stable Stochastic Processes*, Marcel Dekker, New York.

Khindanova, I., Rachev, S. T. & Schwartz, E. (2001), 'Stable modeling of value at risk', *Mathematical and Computer Modelling* **34**, 1223–1259.

Kim, S., Shephard, N. & Chib, S. (1998), 'Stochastic volaitlity: Likelihood inference and comaprison with ARCH models', *Review of Economic Studies* **65**, 361–393.

Kogon, S. M. & Williams, D. B. (1998), *Characteristic Function Based Estimation of Stable Distribution Parameters*, in "A parctical guide to heavy tails", Adler, R. J. and Feldman, R. E. and Taqqu, M. S. eds., chapter 3, pp. 283–310.

Koutrouvelis, I. A. (1980), 'Regression-type estimation of the parameters of stable laws', *Journal of the American Statistical Association* **75**, 919–928.

Liu, S. M. & Brorsen, B. W. (1995), 'Maximum likelihood estimation of a GARCH-stable model', *Journal of Applied Econometrics* **10**, 273–285.

Lombardi, M. J. (2006), 'Bayesian inference for α-stable distributions: A random walk MCMC approach', *Computational Statistics and Data Analysis, in press* .

Lombardi, M. J. & Calzolari, G. (2004), Indirect estimation of α-stable distributions and processes. Working paper 2004/07, Università degli studi di Firenze, Dipartimento di Statistica "G. Parenti".

Lumsdaine, R. L. (1996), 'Consistency and asymptotic normality of the quasi-maximum likelihood estimator in IGARCH(1,1) and covariance stationary GARCH(1,1) models', *Econometrica* **64**, 575–596.

Mandelbrot, B. B. (1963), 'The variation of certain speculative prices', *Journal of Business of the University of Chicago* **36**, 394–419.

McCulloch, J. H. (1978), 'Continuous time processes with stable increments', *Journal of Business* **51**, 601–619.

McCulloch, J. H. (1985), 'Interest-risk sensitive deposit insurance premia: Stable ACH estimates', *Journal of Banking and Finance* **9**, 137–156.

McCulloch, J. H. (1986), 'Simple consistent estimators of stable distribution parameters', *Communications in Statistics – Simulation and Computation* **15**, 1109–1136.

McCulloch, J. H. (1996), *Handbook of Statistics*, Elsevier Science, Amsterdam, chapter Financial Applications of Stable Distributions.

McCulloch, J. H. (2003), The risk-neutral measure and option pricing under log-stable uncertainty. Unpublished working paper, Ohio State University.

Mittnik, S., Doganoglu, T. & Chenyao, D. (1999), 'Computing the probability density function of the stable paretian distribution', *Mathematical and Computer Modelling* **29**, 235–240.

Modarres, R. & Nolan, J. P. (1994), 'A method for simulating stable random vectors', *Computational Statistics* **9**, 11–19.

Monfardini, C. (1998), 'Estimating stochastic volatility models through indirect inference', *The Econometrics Journal* **1**, c113–c128.
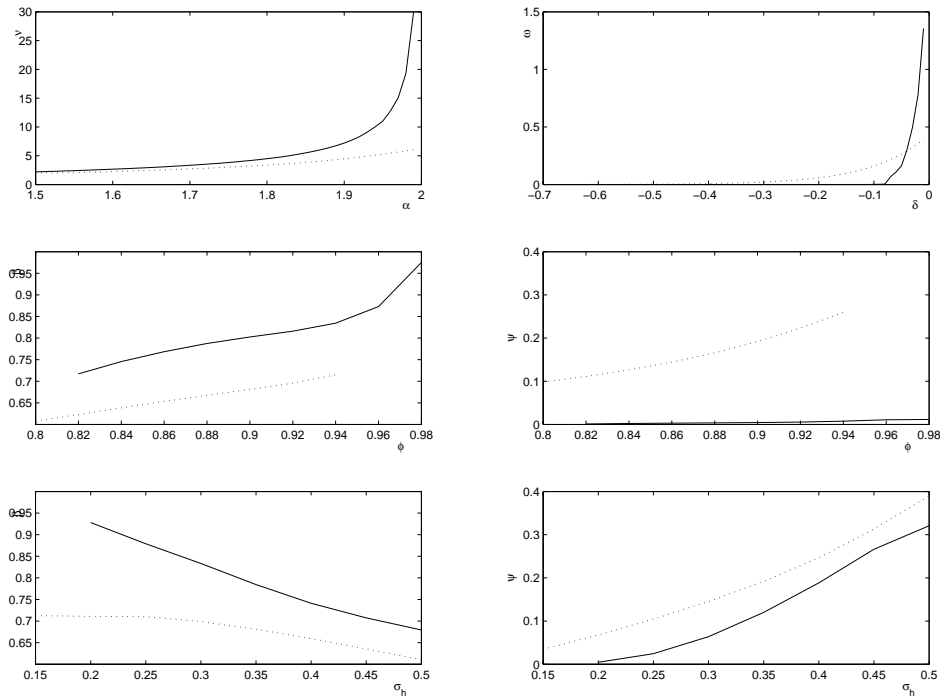
Figure 3: Profiles of the binding function for model (6) various parameter values. The solid line corresponds to "daily parameters" and the dotted line to "weekly parameters".

Nelson, D. B. & Cao, C. Q. (1992), 'Inequality constraints in the univariate GARCH model', *Journal of Business and Economic Statistics* **10**, 229–235.

Nolan, J. P. (1997), 'Numerical computation of stable densities and distribution functions', *Communications in Statistics – Stochastic Models* **13**, 759–774.

Ortobelli Lozza, S., Huber, I. & Schwartz, E. (2002), 'Portfolio selection with stable distributed returns', *Mathematical Methods of Operations Research* **55**, 265–300.

Panorska, A., Mittnik, S., & Rachev, S. T. (1995), 'Stable GARCH models for financial time series', *Applied Mathematics Letters* **8**, 33–37.

Qiou, Z. & Ravishanker, N. (1998), 'Bayesian inference for time series with stable innovations', *Journal of Time Series Analysis* **19**, 235–249.

Samorodnitsky, G. & Taqqu, M. S. (1994), *Stable Non-Gaussian Random Processes*, Chapman & Hall, Boca Raton.

Taylor, S. J. (1986), *Modelling financial time series*, John Wiley and Sons, Chichester.

Zolotarev, V. M. (1986), *One-dimensional Stable Distributions*, American Mathematical Society, Providence.
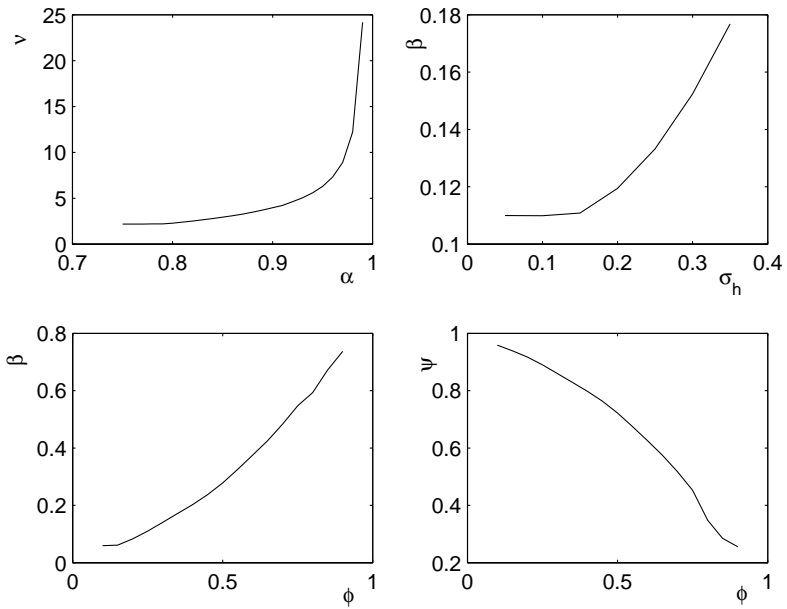
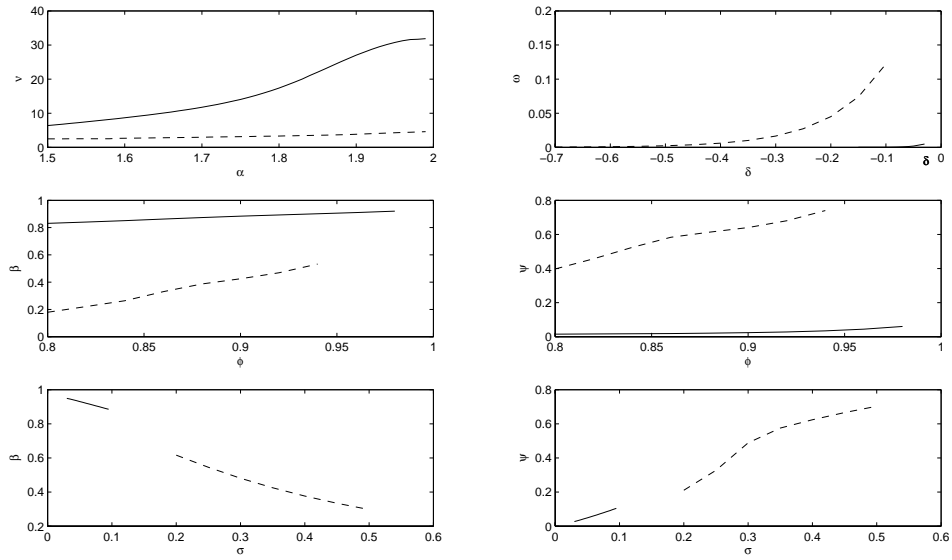Figure 4: Profiles of the binding function for model (8) for various parameter values.

Figure 5: Profiles of the binding function for model (9) for various parameter values. The solid line corresponds to "daily parameters" and the dotted line to "weekly parameters".