JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ., **13**, 659–670 (1998)

OPTIMAL UNIVARIATE INFLATION FORECASTING WITH SYMMETRIC STABLE SHOCKS

PRASAD V. BIDARKOTA[a]† AND J. HUSTON MCCULLOCH*[b]‡
[a]*School of Business, La Trobe University, Bundoora, Victoria 3083, Australia*
E-mail: p.bidarkota@latrobe.edu.au
[b]*Department of Economics, The Ohio State University, 410 Arps Hall, 1945 North High Street, Columbus, OH 43210, USA*
E-mail: mcculloch.2@osu.edu

SUMMARY

Monthly inflation in the United States indicates non-normality in the form of either occasional big shocks or marked changes in the level of the series. We develop a univariate state space model with symmetric stable shocks for this series. The non-Gaussian model is estimated by the Sorenson–Alspach filtering algorithm. Even after removing conditional heteroscedasticity, normality is rejected in favour of a stable distribution with exponent 1·83. Our model can be used for forecasting future inflation, and to simulate historical inflation forecasts conditional on the history of inflation. Relative to the Gaussian model, the stable model accounts for outliers and level shifts better, provides tighter estimates of trend inflation, and gives more realistic assessment of uncertainty during confusing episodes. © 1998 John Wiley & Sons, Ltd.

1. INTRODUCTION

Inflation forecasts are central to macroeconomic analysis. Sargent (1993) suggests simulating the public's forecasts as if they had been formed using econometric techniques. A Gaussian assumption is commonly motivated by the Central Limit Theorem, since economic disturbances are the sum of a large number of unobserved contributions. However, the Generalized Central Limit Theorem (Zolotarev, 1986) indicates that such a sum could also have a non-Gaussian stable distribution. There is evidence that the empirical distributions of inflation rates have thick tails (e.g. Baillie *et al.*, 1996; Balke and Fomby, 1994; Blanchard and Watson, 1986), making stable distributions a natural extension. Batchelor (1981) notes that if errors are truly stable, models built on a false normality assumption can give misleading simulations of the public's forecasts of inflation and other macroeconomic variables.

In this paper, we develop a univariate state space model for the monthly US inflation series using symmetric stable disturbances. Our model can be used for forecasting future inflation, and for simulating historical inflation forecasts conditional on the observed history of inflation. Because the stable model makes more efficient use of the data, it ordinarily gives more precise forecasts of future inflation.

McCulloch (1996) provides a survey of the financial applications of stable distributions. In addition to their Central Limit attributes, stable distributions also have the desirable characteristic that they simplify the long-run forecast error distributions, as described below.

* Correspondence to: J. Huston McCulloch, 410 Arps Hall, 1945 North High Street, Columbus, OH 43210, USA. E-mail: mcculloch.2@osu.edu
†Lecturer, School of Business, La Trobe University, ‡Professor of Economics and Finance, Department of Economics, The Ohio State University, respectively.

CCC 0883–7252/98/060659–12$17·50

Received 9 April 1996
Revised 27 August 1997

Both univariate time series and multivariate structural models have been developed in the literature. A comparison of their relative forecasting performance is made in Fama and Gibbons (1984) and Hafer and Hein (1990). While the univariate models are commonly simple integrated moving average (IMA) models, both Hamilton (1985) and Burmeister *et al.* (1986) use more elaborate multivariate state space models. The state space models developed by Hamilton (1985) and Burmeister *et al.* (1986) fit bivariate Gaussian processes to nominal interest rates and observed inflation.

The univariate model fit here tentatively ignores interest rates, money supplies, and other variables that might have some forecasting power. However, Nelson (1972) notes that parsimonious univariate models often do as well as more elaborate structural and multivariate models.

Inflation rates are often subject to sudden changes in level (Barsky, 1987). Markov switching models have been advocated for modeling such regime shifts (e.g. Hamilton, 1989; Lam, 1990). Whereas these models limit the number of states to a discrete few, the stable model can account for a continuum of state values.

We use monthly inflation data from the United States. Our measure of inflation is based primarily on the CPI-U. However, the CPI-U is generally recognized to have seriously mismeasured the housing component of the true cost of living prior to its 1983 revision. The Bureau of Labor Statistics (BLS) CPI-X series recomputes the CPI-U on the post-1983 basis, from June 1967 until its adoption as the official CPI-U in 1983. We, therefore, constructed a CPI-UX series by linking together the CPI-U from 1953:10 to 1967:6, the CPI-X from 1967:6 to 1983:1, and again the CPI-U from 1983:1 to 1993:9. The 1967 base year was used for both series, so as to reduce rounding error. During 1967:6 to 1983:1, the CPI-U shows 8·5 percentage points more total inflation than does the CPI-X.

The inflation rate was constructed as the first differences of the natural logarithms of the constructed CPI-UX series, expressed as annualized percentage rates. This gave 479 inflation observations, from 1953:11 (October–November inflation) to 1993:9. The 1953 starting date was chosen because of the extensive improvements that were made to the process of constructing the CPI in the early 1950s (see Cosimano and Jansen, 1988; Barsky, 1987, for a discussion, and evidence for structural change in the inflation process). The inflation rate was then seasonally adjusted by subtracting constant seasonal dummies, estimated jointly with a preliminary Gaussian homoscedastic local level model.

Section 2 implements a simple local level model with homoscedastic stable shocks. Section 3 refines the model to include conditional heteroscedasticity. Section 4 investigates the implications of the model for inflationary expectancies, conditional on realized inflation. Section 5 concludes and briefly discusses possible extensions.

2. A HOMOSCEDASTIC LOCAL LEVEL MODEL

We begin with a simple local level model for inflation:

Model 1

$$y_t = x_t + \varepsilon_t \quad \varepsilon_t \sim \text{i.i.d. } S_\alpha(0, c_\varepsilon) \tag{1a}$$

$$x_t = x_{t-1} + \eta_t \quad \eta_t \sim \text{i.i.d. } S_\alpha(0, c_\eta) \tag{1b}$$

Here, y_t is the seasonally adjusted observed inflation, and x_t is an unobserved trend inflation. The ε_t shock captures transitory deviations of observed inflation about the trend, and η_t captures permanent shocks that drive the trend inflation. The two shocks are assumed to be serially and mutually independent. Within the framework of state space models, x_t is the state variable, ε_t is the measurement error, and η_t is the signal shock.

A random variable X is said to have a symmetric stable distribution $S_\alpha(\delta, c)$ if its log-characteristic function can be expressed as:

$$\ln \mathrm{E} \exp(iXt) = i\delta t - |ct|^\alpha$$

The parameters $c > 0$ and $\delta \in (-\infty, \infty)$ are measures of scale and location, respectively, and $\alpha \in (0, 2]$ is the characteristic exponent governing the tail behaviour, with a smaller value of α indicating thicker tails. The stable distribution and density may be evaluated by using Zolotarev's (1986, pp. 74, 78) proper integral representations, or by taking the inverse Fourier transform of the characteristic function. McCulloch (1998) has developed a fast numerical approximation to the symmetric stable distribution and density that has an expected relative density precision of 10^{-6} for $\alpha \in [0{\cdot}84, 2]$. We therefore restrict ourselves in this paper to symmetric stable distributions with α in this range for computational convenience. Although asymmetric stable distributions exist and are well defined, there is no empirical evidence to suggest that asymmetries are important in the inflation series, once outliers have been accounted for (Balke and Fomby, 1994).

The normal distribution belongs to the symmetric stable family with $\alpha = 2$, and is the only member with finite variance, equal to $2c^2$. In this case, the local level model above is equivalent to an IMA(1,1) process, with the first-order autocorrelation coefficient restricted between zero and minus one-half. A number of studies, including Barsky (1987), Ball and Cecchetti (1990), and Brunner and Hess (1993), identify such a process as describing the time series of US inflation rates, to at least a first approximation.

For $\alpha < 2$, a subtle distinction emerges between the local level and IMA(1,1) models. The present study adheres to the former model, because of its natural interpretability as a slowly changing regime observed with noise. In doing so, however, we are validly extending the Gaussian IMA(1,1) literature to the non-Gaussian stable cases.

In the Gaussian case, the Kalman filter provides optimal, linearly adaptive estimates of x_t and, therefore, of future y_t in equation (1) above. However, if $\alpha < 2$, the Kalman filter does not provide optimal estimates, and indeed has an infinite variance stable distribution. Nevertheless, the filtering algorithm of Sorenson and Alspach (1971) enables estimation of both the likelihood function as well as the conditional densities of the state vector under general assumptions about the distributions of the disturbances.

Let $Y_t = (y_1, \ldots, y_t)$. The recursive formulae for obtaining one-step-ahead prediction and filtering densities, due to Sorenson and Alspach, are as follows:

$$p(x_t \mid Y_{t-1}) = \int_{-\infty}^{\infty} p(x_t \mid x_{t-1}) p(x_{t-1} \mid Y_{t-1}) \, \mathrm{d}x_{t-1} \tag{2a}$$

$$p(x_t \mid Y_t) = p(y_t \mid x_t) p(x_t \mid Y_{t-1}) / p(y_t \mid Y_{t-1}) \tag{2b}$$

$$p(y_t \mid Y_{t-1}) = \int_{-\infty}^{\infty} p(y_t \mid x_t) p(x_t \mid Y_{t-1}) \, \mathrm{d}x \tag{2c}$$

Finally, the log-likelihood function, conditional on the hyperparameters α, c_ε, and c_η is given by:

$$\log p(y_1, \ldots, y_T) = \sum_{t=1}^{T} \log p(y_t \mid Y_{t-1}) \tag{3}$$

These formulae have been applied to non-Gaussian data, and extended to include a smoother formula, by Kitagawa (1987). (References to the substantial research in non-Gaussian state space modelling provoked by that work can be seen in Kitagawa and Gersch, 1996.) When $\alpha = 2$, the above filter collapses to the Kalman filter. When disturbances are non-normal stable, the integrals in equations (2a) and (2c) may be evaluated numerically. Details are given in the Appendix.

Maximum likelihood (ML) estimates of the hyperparameters α, c_ε, and c_η, using monthly US inflation data for 1953:11 (October–November inflation) to 1993:9 ($T = 479$), are given in the first row of Table I. The estimate of the characteristic exponent α is 1·803, well below the value of 2 for normal shocks. The estimated signal–noise scale ratio of 1:4·9 suggests a relatively smooth trend inflation as compared to the erratic observed series at the monthly frequency. Hessian-based asymptotic standard errors are in parentheses.

For comparison, we also estimated the above model, constraining $\alpha = 2$, i.e. assuming normal distributions for the error terms. In this case the Kalman filter was employed for estimation. The resulting ML hyperparameters are given in the second row of Table I. The estimated scales (which for $\alpha = 2$ equal standard deviation divided by $\sqrt{2}$) imply a first-order autocorrelation coefficient of −0·49, and a signal–noise scale ratio c_η/c_ε of 1:4·38.

The likelihood ratio (LR) statistic $2\Delta \log L$ for the restriction $\alpha = 2$ is reported in Table I to be 24·76 for Model 1. This LR test statistic does not have the usual $\chi^2(1)$ distribution, since the null hypothesis lies on the boundary of admissible values for α and, hence, the regularity conditions are not satisfied. However, McCulloch (1997) has tabulated small-sample critical values for such a test by Monte Carlo simulations in the simple case when α, c, and δ are being estimated. The models in this paper are more complicated, but this is unlikely to greatly affect the critical values

Table I. Estimation results

Homoscedastic models	α	c_ε	c_η		$\log L$	$2\Delta \log L$ for normality ($\alpha = 2$)[a]	$2\Delta \log L$ for homoscedasticity ($\omega_1 = 0$)[b]
Stable	1·803 (0·094)	1·369 (0·077)	0·278 (0·064)		−1100·78	24·76	
Normal	2	1·563 (0·029)	0·357 (0·051)		−1113·16		
Heteroscedastic models		ω_0	ω_1	ρ			
Stable	1·830 (0·053)	1·376 (0·164)	0·091 (0·027)	0·197 (0·031)	−1089·35	25·98	22·86
Normal	2	1·900 (0·187)	0·085 (0·029)	0·232 (0·037)	−1102·34		27·64

Standard errors are in parentheses.
[a]0·005 critical value equals 6·688 for $T = 1000$, and 7·664 for $T = 300$ (McCulloch, 1997, Table 4, panel b).
[b]0·005 critical value for the χ_1^2 equals 7·879.

for testing $\alpha = 2$. Based on these critical values, the LR test easily rejects normality in favour of $\alpha < 2$ at the 0·005 level of significance (critical value 7·664 for $T = 300$, 6·688 for $T = 1000$), under the homoscedastic assumption of Model 1. A Wald statistic computed from the asymptotic standard error on α likewise does not have an asymptotic normal distribution for the null hypothesis $\alpha = 2$, and indeed it greatly underpredicts the actual change in log-likelihood.

A plot of the standardized residuals (not shown) indicates neither any apparent deviation from a zero mean process nor any trending behaviour, and the Box–Ljung test reveals no significant residual serial correlation. The asymptotic distribution of the Box–Ljung test statistic as a χ^2 requires a finite second-moment assumption. Therefore, although the inferences from this test hold asymptotically only in the Gaussian case, these can only be taken as suggestive when shocks are non-normal stable.

3. LOCAL LEVEL MODEL WITH ARCH ERRORS

Although early work by Engle (1983) established that the US inflation rates were not well described by a homoscedastic process, Cosimano and Jansen (1988) attribute Engle's finding of autoregressive conditional heteroscedasticity (ARCH) to a misspecification in his model. Other work by Bollerslev (1986), Brunner and Hess (1993), Evans and Wachtel (1993), and Baillie *et al.* (1996) also provides mixed evidence of conditional heteroscedasticity in this series. The evidence on ARCH seems to be strong when the sample dataset includes the period prior to 1953, but weaker otherwise.

We extend the homoscedastic model analysed in Section 2 by incorporating GARCH disturbances:

Model 2

$$y_t = x_t + \varepsilon_t \quad \varepsilon_t \sim S_\alpha(0, c_t) \tag{4a}$$

$$x_t = x_{t-1} + \eta_t \quad \eta_t \sim S_\alpha(0, \rho c_t) \tag{4b}$$

$$c_t^\alpha = \omega_0 + \omega_1 |y_{t-1} - E(y_{t-1} \mid Y_{t-2})|^\alpha + \omega_2 c_{t-1}^\alpha, \quad \omega_0 > 0, \omega_1, \omega_2 \geqslant 0 \tag{4c}$$

When errors are normal, equation (4c) reduces to the familiar GARCH-normal process. A test of the GARCH(1,1) versus the ARCH(1) model ($\omega_2 = 0$) yields an LR test statistic of 0·16 when shocks are stable, and 0·00 when shocks are normal. This is contrary to Baillie *et al.* (1996), but the difference is presumably due to their inclusion of data from 1947 to 1953. Hence, we settled on a more parsimonious ARCH(1) parameterization. A GARCH(1,1) process, with α powers as in equation (4c), has been estimated for foreign currency returns by Liu and Brorsen (1995). McCulloch (1985) fits a GARCH-stable model to bond returns, but using absolute values in place of α-powers.

Results of estimating the ARCH(1) version of equation (4), both when shocks are stable and when they are constrained to be normal, are presented in the lower panel of Table I. The results indicate an estimate of 0·091 for ω_1, which is statistically significant by both the Wald and the LR tests at conventional significance levels. The estimate of α rises to 1·830 with the introduction of ARCH, indicating that some of the leptokurtosis can be attributed to conditional heteroscedasticity. However, an LR test for $\alpha = 2$ rejects normality more strongly now in favour of

$\alpha < 2$, at better than the 0·005 level. Thus, inflation appears to display strong leptokurtosis that is not merely a result of conditional heteroscedasticity. The estimated signal–noise scale ratio ρ is now 1:5·08.

Much of the empirical work on modelling uncertainty in inflation is motivated by the Friedman hypothesis that higher rates of trend inflation lead to higher uncertainty in future inflation. The evidence largely seems to corroborate this hypothesis (e.g. Brunner and Hess, 1993; Evans and Wachtel, 1993). We tested this hypothesis by alternatively specifying the conditional scales as dependent on the trend level of inflation x_t, instead of as an ARCH(1) process. Our estimates confirmed earlier evidence that high rates of trend inflation do lead to greater uncertainty, particularly in the variability of permanent shocks than in the variability of temporary shocks to inflation, as suggested by Ball and Cecchetti (1990). However, such a specification for conditional heteroscedasticity was easily dominated by the GARCH(1,1)/ARCH(1) process, in terms of the Akaike Information Criterion (AIC).

4. EFFECTS OF NON-NORMALITY ON INFLATION FORECASTS

It can be shown that so long as $\alpha > 1/2$, the filter density has finite variance for $t \geqslant 2$, despite the infinite variances of the stable shocks. Figure 1 plots the observed inflation series along with the filter mean from the stable ARCH model, $E(x_t \mid Y_t) = E(y_{t+j} \mid Y_t)$, $j \geqslant 1$, and a two-standard error band. The mean and the standard error are computed using the first two moments of the posterior densities of the state variable that were evaluated at a grid of nodes by numerical integration. Divergences between the forecasts from the stable and Gaussian models are particularly important during two episodes, 1973–4 and 1986, which we will examine in detail below.

In Figure 2 we plot the standard errors computed from the filter densities from the stable ARCH model, and for comparison, for the Gaussian ARCH model, computed as if its normal assumption were valid. Without any heteroscedasticity, the standard errors from the homoscedastic normal model would be constant, except for the start-up of the filter. The variation in the computed normal standard errors that does appear after the start-up is entirely due to conditional heteroscedasticity. The normal model standard errors, like those of a standard Kalman filter, do not respond at all to the magnitude of the local disturbances, except indirectly through ARCH effects. The stable model standard errors are usually smaller than those estimated from the normal model. Thus, it ordinarily shows more precise forecasts of future inflation, reacting to forecast errors in a similar way but at a lower level.

During unsettled periods when large outliers or regime shifts appear to have occurred, however, the stable filter density spreads out, and, as we shall see, may even become multimodal. During these periods, the stable forecasts appropriately have higher standard errors. Because the normal filter density does not respond directly to the local disturbances, it must overcompensate during quiet periods with an unnecessarily high standard error.

Figure 3 illustrates the response of the two models to one very big shock that occurred in August 1973. The posterior densities of trend inflation from the two models are plotted from 1973:6 to 1974:1, and the dashed vertical line indicates the actual observed inflation during a specific month. The first two panels show that both models estimated trend inflation to be around 5–7% prior to this period. The observed seasonally adjusted inflation for August was nearly 20% — a jump of about 14 percentage points. The panels show that the Gaussian model reacts strongly with a large shift in location (from about 6% to about 9%), whereas the stable model

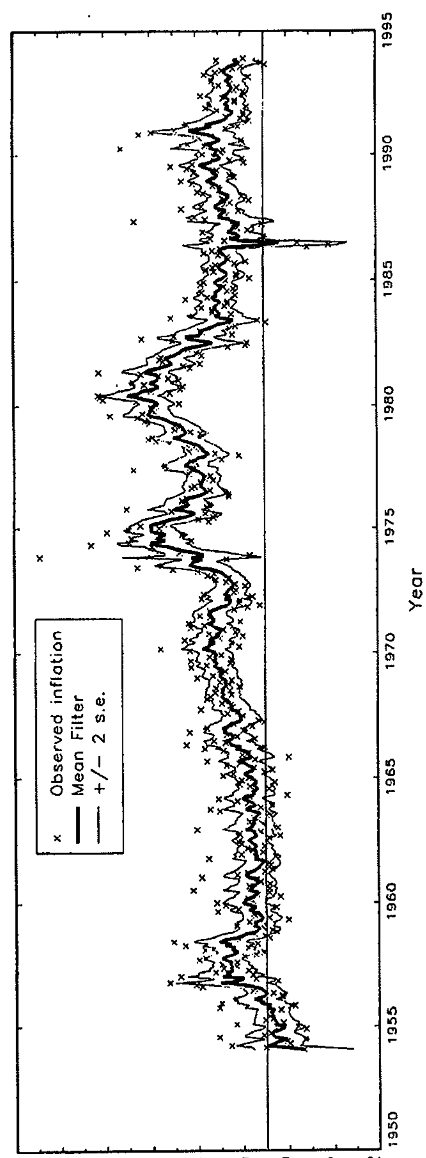

Figure 1. Stable filter

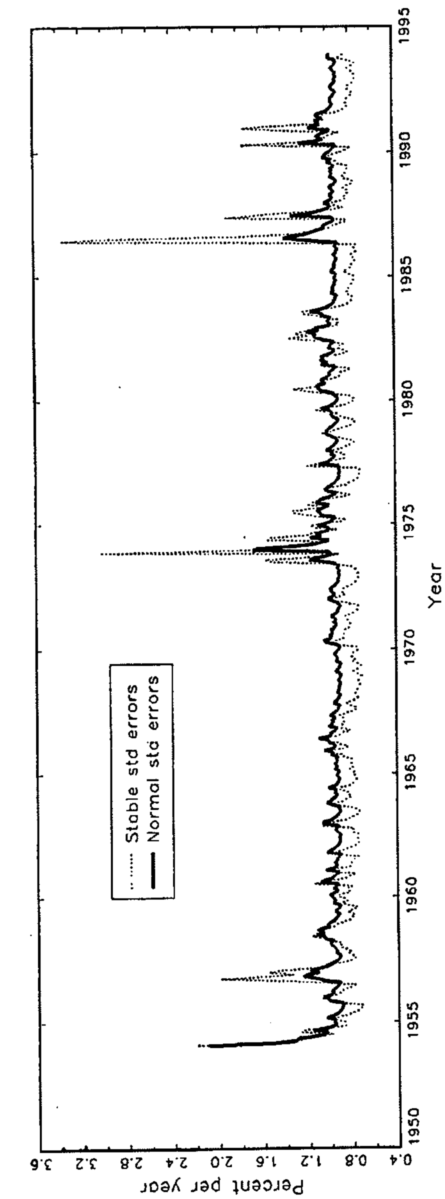

Figure 2. Standard errors from the filter densities

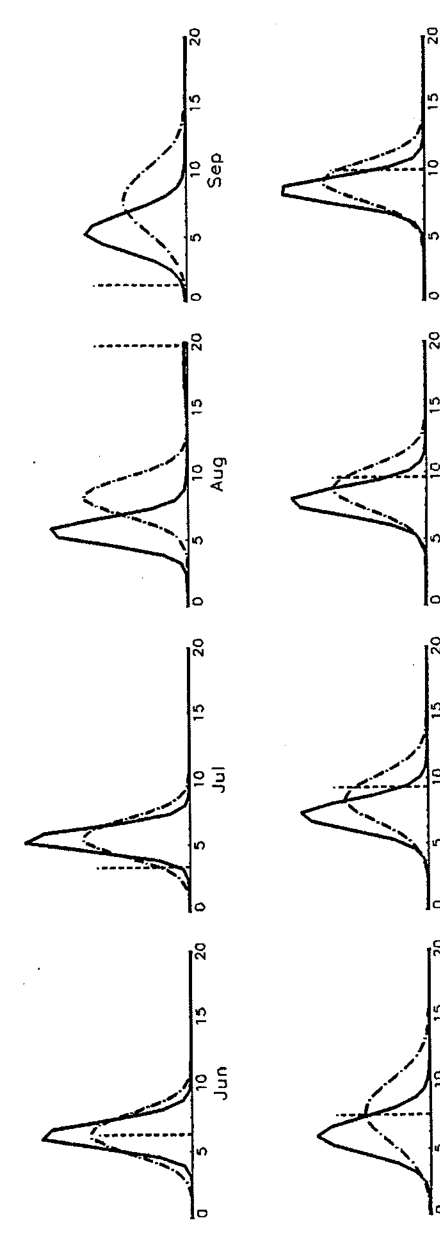

Figure 3. Filter densities, $p(x_t \mid Y_t)$, for 1973:6 to 1974:1. —— stable, —·— Gaussian, - - - newly observed inflation

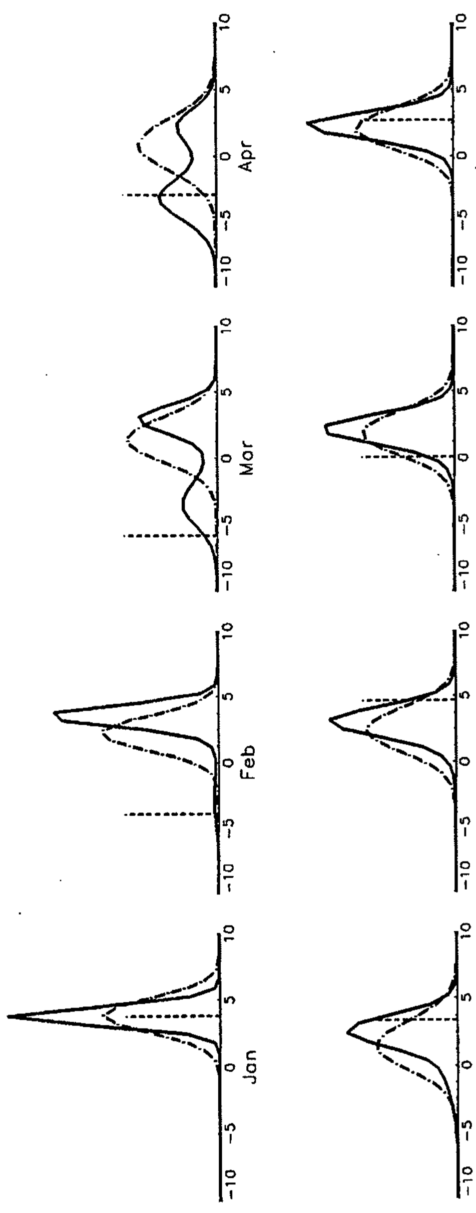

Figure 4. Filter densities, $p(x_t \mid Y_t)$, for 1986:1 to 1986:8. —— stable, —·— Gaussian, - - - newly observed inflation

simply develops a small mode far out in the tail. Observed inflation in the next period was about 1%. The stable model now loses its small new mode in the right tail completely, treating the big shock as a clear outlier, while the Gaussian model continues to forecast over 8% inflation. Kitagawa (1987) obtains an even more dramatic demonstration of the relative insensitivity of non-Gaussian models to outliers, but using artificial data.

Figure 4 shows the behaviour of the two models during 1986:1 to 1986:8 when, starting from a period of moderate inflation of about 3·5–4·0%, three fairly large negative shocks were observed in succession. Seasonally adjusted inflation for February, March, and April 1986 was −3·9%, −5·8%, and −3·0%, respectively. The Gaussian model reacts sluggishly with a shift in its entire distribution, and greatly overpredicts inflation during these three months. On the other hand, the stable model develops a very small mode at the first shock, which builds into a bigger mode at the second shock (while retaining the old mode). With the third shock, however, the stable model has adapted (and more than the normal model) to what it perceives as a level shift, with most of the probability mass now concentrated close to the newly observed inflation. However, it still has a smaller mode back near the old mode, just in case these are just three unusual draws of big observation errors in succession. The following realizations suggest that this, in fact, was true, and the stable model is quick to revert back to the old mode, whereas the normal model is slower to react. By August of 1986, both models settle down on an estimate of about 2% trend inflation.

Thus, these two sets of figures illustrate two important differences in the behaviour of thick-tailed (stable) and Gaussian processes. First, when faced with occasional big shocks the Gaussian model reacts strongly with a large jump in the mean, whereas the stable model largely ignores these, treating them instead as probable outliers. Second, when consecutive shocks of similar magnitude are encountered the normal model is very sluggish in fully adapting, whereas the stable model infers that the regime has changed more quickly. A further fact that emerges is that, during such episodes, the stable model more realistically gives higher estimates of uncertainty associated with trend inflation. This is understandable since the stable model usually develops bimodal densities reflecting greater uncertainty during these confusing periods, whereas the normal model does not.

Abstracting from heteroscedasticity and start-up considerations, the Gaussian Kalman filter density is always normal, with a constant estimated standard error, and adapts linearly to new information. Stuck (1978) and Rutkowski (1994) find the optimal stable filter within the linear class for discrete time processes. Similarly, Le Breton and Musiela (1993) generalize the Kalman filter to continuous time processes with and without infinite variance errors. However, the mean of the globally optimal filter does not revise linearly, unless shocks are truly Gaussian.

The level of future inflation, as implied by the local level model, is:

$$y_{T+j} = x_T + \sum_{i=1}^{j} \eta_{T+i} + \varepsilon_{T+j} \tag{5a}$$

In the homoscedastic Model 1, the stability property of the η and ε shocks implies that $(y_{T+j} - x_T) \sim S_\alpha(0, c_j)$ where $c_j = (jc_\eta^\alpha + c_\varepsilon^\alpha)^{1/\alpha}$. Since x_T is unobserved and has to be estimated, the forecast density of future inflation is obtained by convoluting $S_\alpha(0, c_j)$ with the filter density of x_T. Thus, c_j provides a lower bound on the uncertainty associated with forecasting future inflation. For large values of j, c_j is approximately $j^{1/\alpha}c_\eta$, and the contribution of the filter density is small.

Similarly, the future price level, $\ln(p_{T+j})$, is given by

$$\ln(p_{T+j}) = \ln(p_T) + jx_T + \sum_{i=1}^{j}\sum_{k=1}^{i}\eta_{T+k} + \sum_{i=1}^{j}\varepsilon_{T+i} \tag{5b}$$

$$= \ln(p_T) + jx_T + \sum_{k=1}^{j}(j-k+1)\eta_{T+k} + \sum_{i=1}^{j}\varepsilon_{T+i} \tag{5c}$$

The stability property of the errors again implies that

$$\{\ln(p_{T+j}) - \ln(p_T) - jx_T\} \sim S_\alpha(0, C_j)$$

where $C_j = \{[j^\alpha + (j-1)^\alpha + \cdots + 1]c_\eta^\alpha + jc_\varepsilon^\alpha\}^{1/\alpha}$. Once again, this density needs to be convoluted with the filter density of jx_T in order to obtain the forecast density of the future price level. For large values of j, C_j behaves approximately like $\{j^{(\alpha+1)}/(\alpha+1)\}^{1/\alpha}c_\eta$, while the scale of jx_T grows only in proportion to j.

For Model 2, the exact forecasting uncertainty can only be found by a series of convolutions, taking the heteroscedasticity into account. Nevertheless, it may be expected that the above formulae, using the Model 1 estimated parameters, give a rough approximation of the true forecasting uncertainty.

5. CONCLUSIONS AND EXTENSIONS

The present paper simulates historical US inflation forecasts, conditional on the observed inflation series, using a state space model with symmetric stable disturbances. The model is estimated using the non-linear filtering algorithm of Sorenson and Alspach. Even after adjusting for conditional heteroscedasticity, normality is strongly rejected in favour of stable distributions with characteristic exponent α equal to 1·83. The mean of the filter density for trend inflation is the optimal univariate forecast of future inflation, under the postulated local level model.

The methods used in this paper can readily be extended to other stationary or non-stationary time series that exhibit leptokurtosis, such as stock returns, real income growth, and real interest rates (see Bidarkota, 1996; Bidarkota and McCulloch, 1996; Balke and Fomby, 1994; Blanchard and Watson, 1986). Kitagawa (1987) develops a recursive formula for the smoother density $p(x_t \mid Y_T)$, which may be useful in these contexts. The present study focuses on simulated inflation forecasts, and therefore does not compute the Kitagawa smoother.

Numerical integration of the Sorenson–Alspach filtering equations works best when the dimension of the state vector is small. However, for estimating models with additional explanatory variables, such as money supplies, interest rates, unemployment, etc., and therefore more dimensions, Monte Carlo integration and other procedures proposed by Kitagawa (1996) (see Kitagawa and Gersch, 1996) are promising.

APPENDIX: NUMERICAL INTEGRATION OF THE RECURSIVE FILTERING ALGORITHM

The Sorenson–Alspach filter and predictive densities were evaluated at a grid of 100 points equally spaced on a truncated portion of the real line. The left truncation point was chosen to lie four standard deviations (of the ε shock as measured by a preliminary Kalman filter) below the minimum observed inflation rate (i.e. at −18·0%) and the right truncation point four standard deviations above the maximum observed inflation (i.e. at 32·0%). The likelihood and the predictive density integrals (equations (2c) and (2a) respectively) were evaluated numerically by a piecewise cubic quadrature technique, as follows. Integration between any two interior nodes was performed by fitting a piecewise cubic function through the four nearest nodes, and approximating the required area under the integrand between those nodes by the area under the cubic. The outermost intervals employ the same cubics as the adjacent intervals. For equispaced nodes, eight or more in number, this quadrature procedure yields the weights 8/24, 31/24, 20/24, 25/24, 1, 1, ..., 1, 25/24, 20/24, 31/24, 8/24 for the ordinates. The numerically computed predictive density was normalized at each step in order to ensure that it integrated to unity. The piecewise linear interpolation and the trapezoidal rule for integration suggested by Kitagawa (1987) was not employed. Hodges and Hale (1993) propose an integration by parts procedure to speed up the Kitagawa technique, but this was not used either.

The filter is initialized by the diffuse prior, since we assume that the process is non-stationary, i.e. $p(x_1 \mid Y_1) = s_\alpha(y_1; x_1, c_{\varepsilon 1}(x_1))$, where $s_\alpha(x; \delta, c)$ is the symmetric stable density. Starting points for the hyperparameter estimation are obtained from the Kalman filter under normality.

The accuracy of our numerical quadrature can be gauged by a comparison of the log-likelihood value obtained from our numerical integration with α restricted to be 2, with that obtained from the Kalman filter, for given values of the other hyperparameters. We verified that with 100 nodes our numerical approximation gives log-likelihood values accurate to within 0·001 at the estimated hyperparameters of the Gaussian Model 1. In the light of this, our numerical integration appears to be sufficiently accurate for drawing valid inferences from data. Figures 3 and 4 are constructed by linearly connecting the estimated densities at the computed nodes.

ACKNOWLEDGEMENTS

The authors are grateful to participants at the third annual meeting of the Society for Computational Economics, at the OSU Econometrics Seminar, at the UCSB Statistics Department Symposium on Stable Processes, to John Geweke (editor) and three anonymous referees, for helpful comments and suggestions. All remaining errors are our own responsibility.

REFERENCES

Baillie, R. T., C.-F. Chung, and M. A. Tieslau (1996), 'Analysing inflation by the fractionally integrated ARFIMA-GARCH model', *Journal of Applied Econometrics*, **11**, 23–40.

Balke, N. S. and T. B. Fomby (1994), 'Large shocks, small shocks and economic fluctuations: Outliers in macroeconomic time series', *Journal of Applied Econometrics*, **9**, 181–200.

Ball, L. and S. G. Cecchetti (1990), 'Inflation and uncertainty at short and long horizons', *Brookings Papers on Economic Activity*, 15–54.

Barsky, R. B. (1987), 'The Fisher hypothesis and the forecastability and persistence of inflation', *Journal of Monetary Economics*, **19**, 3–24.

Batchelor, R. A. (1981), 'Aggregate expectations under the stable laws', *Journal of Econometrics*, **19**, 199–210.

Bidarkota, P. V. (1996), 'On modeling real GNP: non-normality and non-linearity, but no long memory', The Ohio State University Department of Economics working paper No. 96-36.

Bidarkota, P. V. and J. H. McCulloch (1996), 'Real stock returns: non-normality, seasonality, and volatility persistence, but no predictability', The Ohio State University Department of Economics working paper No. 96-32.

Blanchard, O. J. and M. W. Watson (1986), 'Are business cycles all alike?' in R. J. Gordon (ed.), *The American Business Cycle: Continuity and Change*, University of Chicago Press, pp. 123–156.

Bollerslev, T. (1986), 'Generalised autoregressive conditional heteroskedasticity', *Journal of Econometrics*, **31**, 307–327.

Brunner, A. D. and G. D. Hess (1993), 'Are higher levels of inflation less predictable? A state-dependent conditional heteroskedasticity approach', *Journal of Business and Economic Statistics*, **11**, No. 2, 187–197.

Burmeister, E., K. D. Wall, and J. D. Hamilton (1986), 'Estimation of unobserved monthly inflation using Kalman filtering', *Journal of Business and Economic Statistics*, **4**, No. 2, 147–160.

Cosimano, T. F. and D. W. Jansen (1988), 'Estimates of the variance of U.S. inflation based upon the ARCH model', *Journal of Money, Credit and Banking*, **20**, No. 3, 409–21.

Engle, R. F. (1983), 'Estimates of the variance of U.S. inflation based upon the ARCH model', *Journal of Money, Credit and Banking*, **15**, No. 3, 286–301.

Evans, M. and P. Wachtel (1993), 'Inflation regimes and the sources of inflation uncertainty', *Journal of Money, Credit and Banking*, **25**, No. 3, 475–520.

Fama, E. F. and M. R. Gibbons (1984), 'A comparison of inflation forecasts', *Journal of Monetary Economics*, **13**, 327–48.

Hafer, R. W. and S. E. Hein (1990), 'Forecasting inflation using interest-rate and time-series models: some international evidence', *Journal of Business*, **63**, No. 1, 1–17.

Hamilton, J. D. (1985), 'Uncovering financial market expectations of inflation', *Journal of Political Economy*, **93**, No. 6, 1224–1241.

Hamilton, J. D. (1989), 'A new approach to the economic analysis of nonstationary time series and the business cycle', *Econometrica*, **57**, No. 2, 357–84.

Hodges, P. E. and D. F. Hale (1993), 'A computational method for estimating densities of non-Gaussian nonstationary univariate time series', *Journal of Time Series Analysis*, **14**, No. 2, 163–178.

Kitagawa, G. (1987), 'Non-Gaussian state space modeling of nonstationary time series', *Journal of the American Statistical Association*, **82**, No. 400, 1032–63.

Kitagawa, G. (1996), 'Monte Carlo filter and smoother for non-Gaussian nonlinear state space models', *Journal of Computational and Graphical Statistics*, **5**, No. 1, 1–25.

Kitagawa, G. and W. Gersch (1996), 'Smoothness priors analysis of time series', *Lecture Notes in Statistics* 116, Springer-Verlag, New York.

Lam, P.-S. (1990), 'The Hamilton model with a general autoregressive component: Estimation and comparison with other models of economic time series', *Journal of Monetary Economics*, **26**, 409–32.

Le Breton, A. and M. Musiela (1993), 'A generalisation of the Kalman filter to models with infinite variance', *Stochastic Processes and their Applications*, **47**, 75–94.

Liu, S. M. and B. W. Brorsen (1995), 'Maximum likelihood estimation of a GARCH-stable model', *Journal of Applied Econometrics*, **10**, 273–85.

McCulloch, J. H. (1985), 'Interest-risk sensitive deposit insurance premia: stable ACH estimates', *Journal of Banking and Finance*, **9**, 137–56.

McCulloch, J. H. (1996), 'Financial applications of stable distributions', in G. S. Maddala and C. R. Rao (eds), *Statistical Methods in Finance* (Handbook of Statistics, 14), Elsevier, Amsterdam, pp. 393–425.

McCulloch, J. H. (1997), 'Measuring tail thickness in order to estimate the stable index α: A critique', *Journal of Business and Economic Statistics*, **15**, No. 1, 74–81.

McCulloch, J. H. (1998), 'Numerical approximation of the symmetric stable distribution and density', in R. Adler, R. Feldman, and M. S. Taqqu (eds), *A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tailed Distributions*, Birkhauser, Boston.

Nelson, C. R. (1972), 'The prediction performance of the FRB-MIT-PENN model of the U.S. economy', *The American Economic Review*, **LXII**, No. 5, 902–17.

Rutkowski, M. (1994), 'Optimal linear filtering and smoothing for a discrete-time stable linear model', *Journal of Multivariate Analysis*, **50**, 68–92.

Sargent, T. J. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford.

Sorenson, H. W. and D. L. Alspach (1971), 'Recursive Bayesian estimation using Gaussian sums', *Automatica*, **7**, 465–79.

Stuck, B. W. (1978), 'Minimum error dispersion linear filtering of scalar symmetric stable processes', *IEEE Trans. Aut. Control*, **AC-23**, 507–509.

Zolotarev, Z. M. (1986), *One Dimensional Stable Laws*, American Mathematical Society. (Translation of *Odnomernye Ustoichivye Raspredeleniia*, Nauka, Moscow, 1983.)

