ELSEVIER

Statistics & Probability Letters 53 (2001) 381–390

STATISTICS & PROBABILITY LETTERS

www.elsevier.nl/locate/stapro

A method for fitting stable autoregressive models using the autocovariation function

Colin M. Gallagher

Mathematical Sciences, Clemson University, Clemson, SC 29634, USA

Received June 1999; received in revised form January 2001

Abstract

We use the sample covariations to estimate the parameters in a univariate symmetric stable autoregressive process. Unlike the sample correlation, the sample covariation can be used to estimate the tail decay parameter of the process. The fitted model will be consistent with the dependence as measured by the covariation. The limit distribution of the sample covariation can be used to derive confidence intervals for the autoregressive parameter in a first order process. Simulations show that confidence intervals coming from the covariation have better coverage probabilities than those coming from the sample correlations. The method is demonstrated on a time series of sea surface temperatures.

Keywords: Time series; Autocovariation; Autoregression

1. Introduction

In this paper, we describe a method for fitting a symmetric alpha-stable (SαS) autoregressive (AR) process

$$X_t - \phi_1 X_{t-1} - \cdots - \phi_p X_{t-p} = Z_t, \tag{1.1}$$

where the sequence $\{Z_t\}$ is an iid sequence of SαS random variables with $\alpha \in (1,2]$, to time series data generated from a process with finite absolute mean. We note here that if $\alpha = 2$ we are fitting a traditional Gaussian AR model. For background on stable processes the reader is referred to Samorodnitsky and Taqqu (1994).

When $\alpha < 2$ the process given by (1.1) has infinite variance. Should we ever fit such a model to a data set? Can we sometimes get a better fit using a model with infinite variance? After we develop our methodology in Sections 2 and 3, we will fit a Gaussian AR process and a stable non-Gaussian AR process to the same data set.

E-mail address: cgallag@ces.clemson.edu (C.M. Gallagher).

0167-7152/01/$ - see front matter
PII: S0167-7152(01)00041-4

In a traditional time series model we assume an underlying stationary stochastic process $\{X_t\}$ has a finite second moment (see for example Brockwell and Davis, 1991). For ease of presentation, we will assume throughout that our discrete time process has mean zero (if this were not the case we could subtract the sample mean). Given time series data $x_1, x_2, \ldots, x_n$, we can use the sample autocorrelation function (ACF)

$$\hat{\rho}(h) = \frac{\sum_{t=1}^{n-h} x_t x_{t+h}}{\sum_{t=1}^{n} x_t^2} \quad h = 0, 1, 2, \ldots \tag{1.2}$$

to fit a time series model. In fact, we can mimic the observed sample ACF for any integer number of lags by a Gaussian autoregressive process. The autoregressive parameters of our approximating process come from the Yule–Walker equations which use the empirical ACF to estimate the parameters. If $\{X_t\}$ is stationary and ergodic then $\hat{\rho}(h)$ converges in the almost sure sense to the theoretical correlation function:

$$\rho(h) = \frac{EX_t X_{t+h}}{E|X_t|^2}. \tag{1.3}$$

For the process $\{X_t\}$ with $E|X_t|^2 = \infty$, the theoretical ACF no longer exists. However, the sample ACF is still well defined. For any causal ARMA process with representation:

$$X_t = \sum_{j=0}^{\infty} \psi_j Z_{t-j} \quad t = 1, 2, \ldots, \tag{1.4}$$

where $\{Z_t\}$ is an iid sequence of random variables in the domain of attraction of a stable law, the sample ACF $\hat{\rho}(h)$ converges in probability to the constant limit

$$\rho(h) = \frac{\sum \psi_j \psi_{j+h}}{\sum \psi_j^2}.$$

(Davis and Resnick, 1985) time series modeling techniques, such as the Yule–Walker equations, which depend on the sample ACF can be used to fit autoregressive models to heavy tailed ARMA data.

What if our data is not generated by an ARMA process, but is coming from a process which is stationary and ergodic? Can we use the sample ACF to approximate $\{X_t\}$ with an autoregression as in the finite second moment case? Resnick et al. (1998) give examples of stationary ergodic stable processes for which the sample ACF converges to a random limit and caution that, "the usual time series model fitting and diagnostic tools such as the Akiake Information Criterion or Yule–Walker estimators will be of questionable applicability" (see also Resnick, 1997).

We consider zero mean stationary ergodic processes, $\{X_t\}$ with $E|X_t| < \infty$. For this class of processes the *autocovariation function* (AcovF) at lag k,

$$\lambda(k) = \frac{EX_t S_{t-k}}{E|X_{t-k}|} \quad \text{for } k = 0, \pm 1, \pm 2, \ldots, \tag{1.5}$$

where

$$S_t = \operatorname{sign}(X_t)$$

is well defined. Its sample version

$$\hat{\lambda}(k) = \frac{\sum_{t=l}^{r} x_t s_{t-k}}{\sum_{t=1}^{n} |x_t|}, \tag{1.6}$$

where $l = \max(1, 1+k)$ and $r = \min(n, n+k)$, satisfies

$$\hat{\lambda}(k) \to \lambda(k), \tag{1.7}$$

where the convergence in (1.7) is in the almost sure sense. The limiting behavior of the sample AcovF is investigated in Gallagher (2000).

The *covariation* was introduced in Cambanis and Miller (1981) as a measure of dependence for jointly SαS vectors. For a SαS process $\lambda(k)$ is the normalized covariation of X_t on X_{t-k} and

$$E(X_t|X_{t-k}) = \lambda(k)X_{t-k}. \tag{1.8}$$

Given an empirical autocovariation function we can estimate the parameters of an autoregressive process using a generalization of the Yule–Walker equations (see Nikias and Shao, 1995).

In Section 3, we see that we can use the sample AcovF to estimate the remaining parameters in model (1.1). We use the sample AcovF to derive confidence intervals for ϕ_1 in an AR(1) process in Section 4 and give simulation results indicating that these intervals perform better (in terms of coverage probability) than do intervals based on the ACF. We conclude the paper by fitting model (1.1) to a data set consisting of 217 daily sea surface temperatures recorded at an open sea buoy.

2. Generalized Yule–Walker equations

In this section, we describe how to fit an autoregression to a data set using the sample AcovF. For an AR(p) process, we can use the sample AcovF to estimate the autoregressive parameters.

Let $\{X_t\}$ be an autoregressive process satisfying (1.4). We describe the relationship between the autoregressive parameters and the AcovF.

Using the fact that $\{X_t\}$ is causal we have,

$$\lambda(k) = \phi_1\lambda(k-1) + \phi_2\lambda(k-2) + \cdots + \phi_p\lambda(k-p) \quad \text{for } k = 1, 2, \ldots, p.$$

In matrix form,

$$\lambda = \mathbf{\Lambda}\phi, \tag{2.9}$$

where,

$$\lambda = (\lambda(1), \ldots, \lambda(p))^{\mathrm{t}},$$

$$\phi = (\phi_1, \ldots, \phi_p)^{\mathrm{t}},$$

and $\mathbf{\Lambda}$ is the $p \times p$ matrix with (i,j)th element

$$\mathbf{\Lambda}(i,j) = \lambda(i-j) \quad \text{for } i, j = 1, \ldots, p. \tag{2.10}$$

We use the sample autocovariation to develop estimators of the parameters $\phi_1, \ldots, \phi_p$.

Nikias and Shao (1995) use (2.9) to estimate the autoregressive parameters by replacing $\lambda(k)$ with an estimated value $\hat{\lambda}(k)$, and give simulation evidence indicating that this estimator performs as well as the usual Yule–Walker estimator when $\hat{\lambda}(k)$ is given by (1.6). The generalized Yule–Walker estimate $\hat{\phi}$ is any vector ϕ satisfying

$$\hat{\lambda} = \hat{\mathbf{\Lambda}}\hat{\phi}, \tag{2.11}$$

where

$$\hat{\lambda} = (\hat{\lambda}(1), \ldots, \hat{\lambda}(p))^{\mathrm{t}},$$

and $\hat{\mathbf{\Lambda}}$ is the $p \times p$ matrix with (i,j)th element

$$\hat{\mathbf{\Lambda}}(i,j) = \hat{\lambda}(i-j) \quad \text{for } i,j = 1,\ldots,p.$$

Remark 2.1. If X_t is either an AR(1) process or an AR(2) process, then the parameter estimates are unique and strongly consistent. In either case both $\hat{\mathbf{\Lambda}}$ and $\mathbf{\Lambda}$ are nonsingular. In both cases $\hat{\mathbf{\Lambda}}$ is almost surely nonsingular if $P(X_t = 0) = 0$. If X_t is an AR(1) process, $\Lambda = \lambda(1)$. If X_t is an AR(2) process, Λ is nonsingular unless $\lambda(1) = \lambda(-1) = \pm 1$. In such a case

$$\lambda(1) = \lambda(0)\phi_1 + \lambda(-1)\phi_2 = \phi_1 + \lambda(1)\phi_2.$$

So that X_t is not causal. Consistency follows from (1.7).

3. Fitting a stable process

In this section, we consider fitting model (1.1) to a data set. All the parameters in the model can be estimated with a set of $p+2$ equations. From a fitted model we obtain residuals which can be used to check the adequacy of the fit. Since the estimators are simple and quickly calculated on a computer, we can fit various models and use residual diagnostics to choose a parsimonious model which provides an adequate fit.

The SαS AR process is completely determined by the autoregressive parameters, the tail decay parameter α and the scale parameter of the iid sequence $\{Z_t\}$. We can use the sample AcovF to estimate ϕ as in the previous section. Once we have an estimate $\hat{\phi}$ we can use this and the sample AcovF to estimate α as described below. The scale parameter σ can be estimated from the residuals obtained from the fitted AR parameters.

We have seen that given an observed autocovariation function on the integers in $[-p, p]$, we can estimate the autoregressive parameters of an AR(p) process and use the process

$$X_t - \hat{\phi}_1 X_{t-1} - \cdots - \hat{\phi}_p X_{t-p} = \hat{Z}_t, \tag{3.12}$$

to estimate our underlying process. Does this process have the same theoretical autocovariation for every iid sequence $\{\hat{Z}_t\}$? The answer is no. The autocovariation is both a function of the autoregressive parameters and α. How do we use the observed AcovF to estimate α?

The following fact, which we state as a proposition, suggests a procedure for estimating α from the sample AcovF.

Proposition 3.1. *Let X_t be a causal* SαS *AR(p) process. The following equation holds*

$$(E|Z_1|/E|X_1|)^\alpha = 1 - \phi_1\lambda(-1) - \cdots - \phi_p\lambda(-p). \tag{3.13}$$

Proof. Using properties of stable random variables (1.8) and (1.1)

$$E|X_t| = E|X_t| \left(\phi_1\lambda(-1) + \cdots + \phi_p\lambda(-p) + \left(\sum |\psi_j|^\alpha \right)^{-1} \right)$$

and

$$(E|Z_1|/E|X_1|)^\alpha = \left(\sum |\psi_j|^\alpha \right)^{-1}. \qquad \square$$

We can estimate α by replacing the autoregressive parameters by their estimates and $\lambda(\cdot)$ by $\hat{\lambda}(\cdot)$ in Eq. (3.13) and solving for α

$$\hat{\alpha} = \frac{\log(1 - \hat{\phi}_1\hat{\lambda}(-1) - \cdots - \hat{\phi}_p\hat{\lambda}(-p))}{\log(\sum_{t=1}^{n} |\hat{Z}_t| / \sum_{t=1}^{n} |X_t|)}. \tag{3.14}$$

Of course α could be estimated from the residuals using any of the estimators for iid data discussed in the literature, but this estimate of α is consistent with the observed AcovF.

The final parameter in our model is the scale parameter. Once we have a fitted model we can estimate the scale of Z_t using the residuals. Under model (1.1)

$$\sigma = \pi E|Z_t|/(2\Gamma(1 - 1/\alpha)).$$

A simple estimate is

$$\hat{\sigma} = \pi/(2\Gamma(1 - 1/\hat{\alpha}))n^{-1}\sum |\hat{Z}_t|. \tag{3.15}$$

Let $\hat{\lambda}(\cdot)$ be the sample autocovariation of an observed time series on the integers in $[-p, p]$. We can approximate the underlying stochastic process from which our data is generated with (3.12), where $Z_t \sim S\alpha S(\sigma)$, α is given by (3.14), σ is given by (3.15) and the autoregressive parameters are given by (2.11).

Diagnostic checking: If our model fits the data well, the residual $\hat{Z}_t$ should be close to Z_t. Since we assume the sequence $\{Z_t\}$ is iid, if the model provides a good fit the residuals should appear to be from an iid sequence. We consider two diagnostic checking procedures, both of which test to see if the residuals behave like an iid sample from a heavy tailed data set. The first is based on the sample AcovF of the residuals. The second is a nonparametric generalization of the Portmanteau test.

For an iid sequence, the AcovF is zero for all $k \neq 0$. Since $\hat{\lambda}(k) \to \lambda(k)$, the sample AcovF should be close to zero as well. The limiting distribution of the sample AcovF for causal ARMA processes is derived in Gallagher (2000) and can be used to decide what is a significant departure from zero. This can be done with a formal hypothesis test or using a graphical procedure. This is discussed in more detail in Gallagher (2000).

The usual Portmanteau test statistic doesn't perform well in the heavy tailed case (see Runde, 1998). Dufour and Roy (1986) derive the limiting distribution of a generalized Portmanteau test statistic which is based on the ranks of the data. The chi-square limiting distribution is valid for iid sequences of continuous random variables.

4. Confidence intervals for AR(1) processes: a simulation

In this section, we consider the AR(1) model

$$X_t = \phi X_{t-1} + Z_t, \tag{4.16}$$

where $Z_t \sim S\alpha S$ with $\alpha > 1$. For this simple model, ϕ can be estimated with the sample ACF at lag 1 (Yule–Walker estimate) or the AcovF at lag 1 (Generalized Yule–Walker). The limiting distributions of the sample ACF and AcovF can be used to derive confidence intervals for ϕ based on their respective estimates.

Let $\hat{\phi}_{\mathrm{ACF}}$ be the estimate of ϕ coming from the sample ACF. The limiting distribution of $\hat{\phi}_{\mathrm{ACF}}$ for $\alpha < 2$ is derived in Davis and Resnick (1986) and is given by

$$(n/\log(n))^{(1/\alpha)}(\hat{\phi}_{\mathrm{ACF}} - \phi) \Rightarrow (1 - \phi^{\alpha})^{1/\alpha}U/V, \tag{4.17}$$

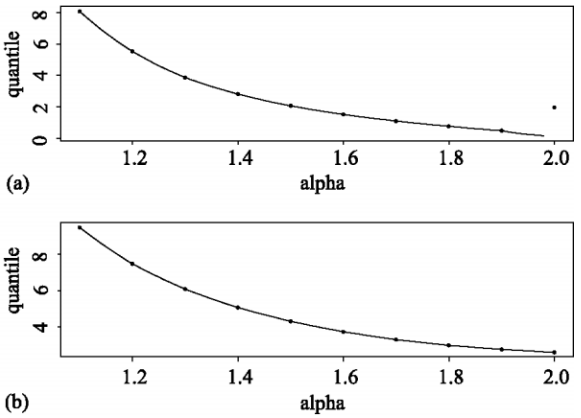

Fig. 1. (a) 0.975 quantiles of U/V in (4.17) and N in (4.18) and (b) 0.975 quantiles of S in (4.19). Quantiles for alpha values other than 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, and 2 were approximated using interpolation.

where U and V are independent stable random variables. When $\alpha = 2$, the limiting distribution and normalization change

$$n^{1/2}(\hat{\phi}_{\mathrm{ACF}} - \phi) \Rightarrow (1 - \phi^2)^{1/2} N, \tag{4.18}$$

where N has a normal distribution. For a fixed sample size n the normalizing constant has a discontinuity at $\alpha = 2$. Likewise, the limiting distribution has a discontinuity at $\alpha = 2$ (e.g. the limit of the 0.975 quantile of U/V as $\alpha \to 2$ is not the appropriate quantile of the normal distribution). In practice this can present a problem. If for example the data is normally distributed ($\alpha = 2$), and α is estimated to be around 1.98, the quantile used in the confidence interval corresponding to $\alpha = 1.98$ is vastly different from the appropriate quantile of the normal distribution. This can be seen from Fig. 1a which graphs the 0.975 quantile of U/V in (4.17) and N in (4.18) versus α. Let $\hat{\phi}_{\mathrm{COV}}$ be the estimate of ϕ coming from the sample AcovF. The limiting distribution of $\hat{\phi}_{\mathrm{COV}}$ is given in Gallagher (1998) and is given by

$$n^{1-1/\alpha}(\hat{\phi}_{\mathrm{COV}} - \phi) \Rightarrow (1 - \phi^{\alpha})^{1/\alpha} S, \tag{4.19}$$

where $S \sim S\alpha S$ for $\alpha \in (1, 2]$. Note that there is no discontinuity in either the normalization or the limiting distribution.

To use the above limiting results to make confidence intervals we need appropriate quantiles of the limiting distributions. Quantiles of the distribution of U/V are given in Adler et al. (1998) for various equally spaced α values in $(0, 2)$. Quantiles of S can be found in Samorodnitsky and Taqqu (1994).

To compare confidence intervals using the two competing estimation procedures mentioned above, we simulate from model (4.16) with $\phi = 0.9$ for various α values and calculate confidence intervals counting the number of intervals which cover ϕ. The results are shown in Table 1. For each α considered we simulated $10,000$ samples with a sample size of $n = 1000$. From each simulated sample, five confidence intervals were calculated

(i) Using $\hat{\phi}_{\mathrm{ACF}}$ and assuming α is known.
(ii) Using $\hat{\phi}_{\mathrm{COV}}$ and assuming α is known.
(iii) Using $\hat{\phi}_{\mathrm{ACF}}$ and estimating α using McCulloch's method (McCulloch, 1986).
(iv) Using $\hat{\phi}_{\mathrm{COV}}$ and estimating α using McCulloch's method.
(v) Using $\hat{\phi}_{\mathrm{COV}}$ and estimating α using (3.14).

Table 1
Simulated coverage percentages for 95% confidence intervals for ϕ in model (4.16). When α was assumed to be unknown it was estimated using $\hat{\alpha}$ coming from (3.14) or using $\hat{\alpha}_1$ coming from McCulloch's (1986) quantile estimation method

	Alpha unknown			Alpha known	
α	$\hat{\phi}_{COV}$		$\hat{\phi}_{ACF}$	$\hat{\phi}_{COV}$	$\hat{\phi}_{ACF}$
	$\hat{\alpha}$	$\hat{\alpha}_1$	$\hat{\alpha}_1$	NA	NA
1.1	93.77	96.31	85.48	96.04	81.48
1.2	93.84	96.92	86.83	97.49	85.62
1.3	93.66	97.48	87.68	97.84	87.86
1.4	94.52	97.87	88.39	98.07	89.89
1.5	94.35	97.47	88.13	97.75	90.18
1.6	94.89	96.91	86.70	97.60	90.10
1.7	95.44	96.71	84.76	97.46	88.15
1.8	95.62	96.16	81.72	96.81	82.20
1.9	96.10	95.98	81.31	95.91	70.00
2.0	96.32	96.62	81.19	94.93	94.71

For estimated α values for which the quantile of the limiting distribution was unknown, interpolation was used to estimate the appropriate quantile (see Fig. 1). Also, any estimated value of α (coming from either method) which was outside the appropriate range was set to be the closest value in $[1.1, 2]$. For example if $\hat{\alpha} = 1$, $\hat{\alpha}$ was set to be 1.1 while if $\hat{\alpha} = 2.1$, $\hat{\alpha}$ was taken to be 2.

We can see from the first three columns of Table 1 that the confidence intervals coming from the AcovF outperform those from the ACF. For $\alpha \in [1.1, 1.8]$ the confidence intervals using $\hat{\alpha}$ coming from (3.14) have actual coverage probabilities closest to 95%. When α is closer to 2, it appears that the confidence interval using $\hat{\alpha}_1$ coming from McCulloch's estimator works the best.

Not surprisingly when $\alpha = 2$ (unknown) the confidence interval using $\hat{\phi}_{ACF}$ performs poorly. The 95% confidence interval based on the sample ACF only covered the parameter about 81% of the time. Whereas the confidence interval based on the AcovF covered the parameter about 96% of the time.

We can also see from Table 1 that for heavy tailed data ($\alpha < 2$) the large sample distribution of the sample ACF provides a poor approximation. Even when α is known we can see from the last two columns in the table that the intervals based on the AcovF have true coverage probabilities much closer to 95%. Surprisingly this is true even when $\alpha = 2$. As α approaches 2 the true coverage probabilities based on the AcovF seem to approach 95%. For the intervals coming from $\hat{\phi}_{ACF}$ we observe the opposite of this phenomenon. As α approaches 2 the coverage probabilities become farther from the nominal 95% level. Both the poor approximation of the limiting distribution of the sample ACF and the decrease in accuracy as α approaches 2 have been observed before (For example, see Adler et al., 1998; Runde, 1998.).

Table 1 demonstrates a common problem with large sample approximations in the heavy tailed case. Looking at the last two columns we see that even with a sample size of 1000 and α known, we get a poor distributional approximation. The AcovF gives intervals that are systematically closer to the nominal 95% level than the ACF intervals, but we would only conclude that the true coverage is 95% when $\alpha = 2$. In all other cases the simulated coverage is more than two standard deviations (coming from the binomial distribution) from the nominal level.

5. An example

Our example data set consists of sea surface temperatures for 217 days beginning on October 1st 1988. We see from the time series plot in Fig. 2 that the temperature seems to be decreasing for about the first 120

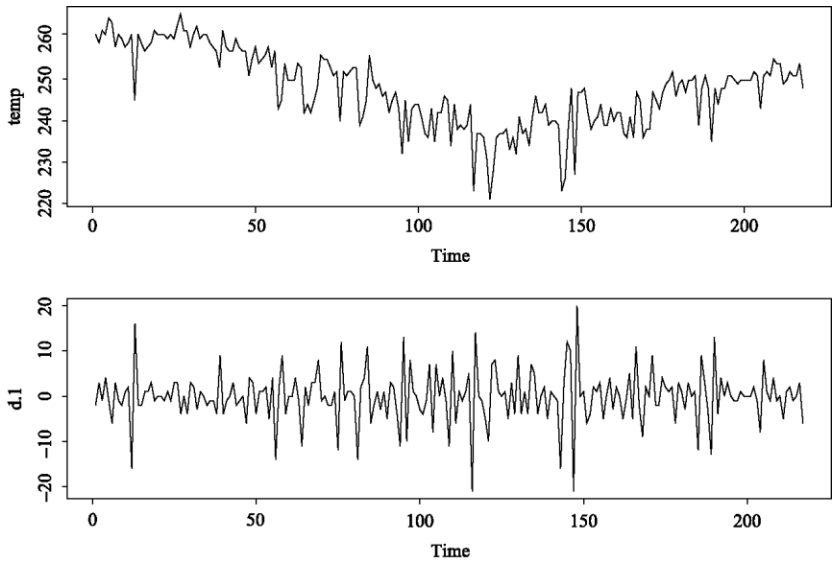

Fig. 2. Plots of temperature time series and time series of differences.

days and then begins to increase. Clearly we cannot assume the data is generated by a stationary stochastic process. We take first differences to remove the trend. The data set of differences is plotted in Fig. 2 as well. A trend is no longer apparent. Notice that there are quite a few jumps in the time series which may be an indication of heavy tails.

Using the methodology from Section 4, we fit model (1.1) to the data set of differences. We start by fitting an AR(1) model. Using the two diagnostic procedures described above we conclude the residuals cannot be assumed to be iid. We add one parameter at a time fitting models of higher order and examining the residuals until we find a model with residuals which appear to be iid. The chosen model has $p = 6$, $\hat{\sigma} = 2.7$, and $\hat{\alpha} = 1.76$.

We also fit a Gaussian AR model to the data set. The Yule–Walker equations were used to estimate the autoregressive parameters. The order p was decided based on parsimony and Akiake's information criterion as well as residual diagnostics. The fitted Gaussian model has $p = 6$ and $\hat{\sigma} = 4.8$, where $\hat{\sigma}$ is the estimated standard deviation of Z_1.

For the Gaussian model the sequence $\{\hat{Z}_t\}$ is a sequence of iid normal random variables with $\sigma = 4.8$. Fig. 3 is a plot of this normal density function along with a nonparametric density estimate from the residuals of the Gaussian model. Under the stable (non-Gaussian) model the innovation sequence $\{\hat{Z}_t\}$ is an iid sequence of SαS random variables with $\alpha = 1.76$ and $\sigma = 2.7$. Fig. 4 shows a plot of this density function along with a nonparametric density estimate from the residuals of the fitted model. Both nonparametric density estimates were done on S-plus using a normal kernel. The density estimate for the stable model is very close to the fitted stable density.

We note here that the estimated parameters from both the Yule–Walker and generalized Yule–Walker equations are quite similar in this case. However, the sample ACF cannot be used to estimate the tail decay of the stable distribution. The advantage of the method described in this paper is that it provides a complete probability model for the process. The moment estimators considered can be used as preliminary estimates in a numerical maximization of the stable likelihood.

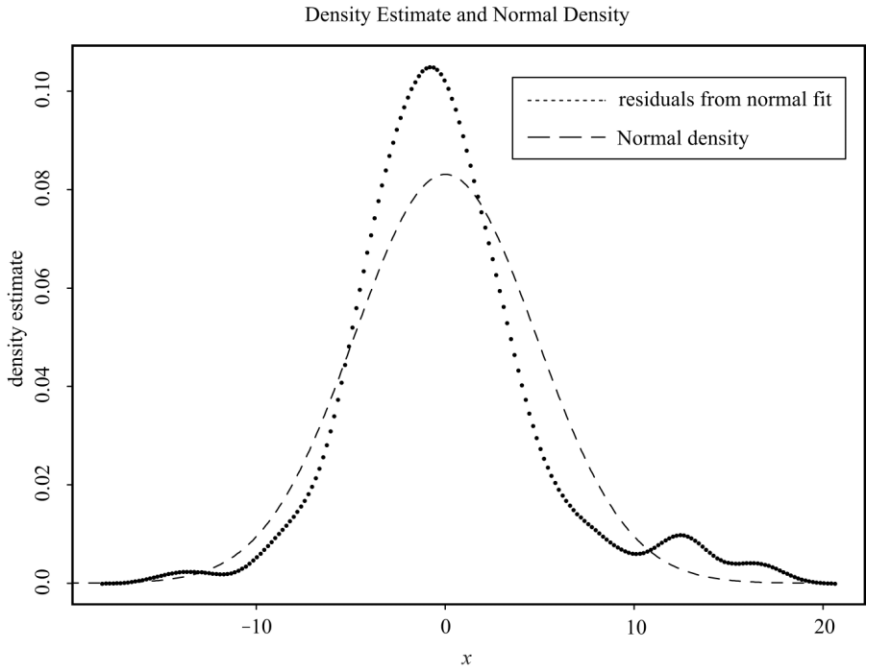

Fig. 3. Normal density function and nonparametric density estimate of residuals from Gaussian model.

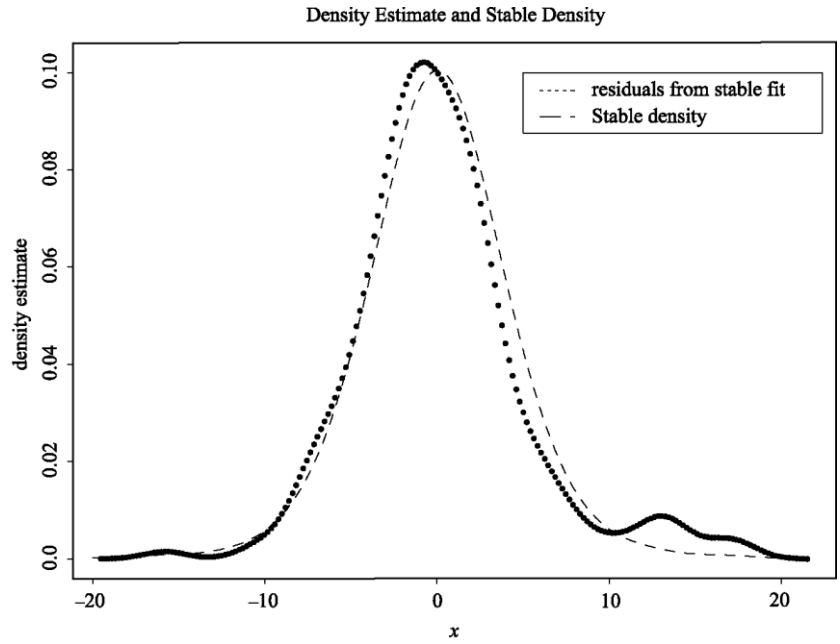

Fig. 4. Stable density function and nonparametric density estimate of residuals from non-Gaussian model.

Acknowledgements

I would like to thank Dr. Roman Glazman and Dr. Benny Chang of JPL for providing the buoy data and the referee for his remarks.

References

Adler, R., Feldman, R., Gallagher, C., 1998. Analysing stable time series. In: Adler, R.J., Feldman, R.E., Taqqu, M.S. (Eds.), A Practical Guide to Heavy Tails. Birkhäuser, Boston.

Brockwell, P.J., Davis, R.A., 1991. Time Series: Theory and Methods, 2nd Edition. Springer, New York.

Cambanis, S., Miller, G., 1981. Linear problems in pth order and stable processes. SIAM J. Appl. Math. 41, 43–69.

Davis, R., Resnick, S., 1985. Limit theory for moving averages of random variables with regularly varying tail probabilities. Ann. Probab. 13, 179–195.

Davis, R., Resnick, S., 1986. Limit theory for the sample covariance and correlation functions of moving averages. Ann. Statist. 14, 533–558.

Dufour, J., Roy, R., 1986. Generalized Portmanteau statistics and tests of randomness. Comm. Statist. Theory Methods 10, 2953–2972.

Gallagher, C., 1998. Fitting ARMA models to heavy-tailed data, PhD dissertation University of California, Santa Barbara (UMI microform 9840773).

Gallagher, C., 2000. Estimating the autocovariation from stationary heavy-tailed data, with applications to time series modeling, Department of Mathematical Sciences Technical Report, #Gal2. Clemson University.

McCulloch, J.H., 1986. Simple consistent estimators of stable distribution parameters. Comm. Statist. Comput. Simulation 15, 1109–1136.

Nikias, C.L., Shao, M., 1995. Signal Processing with Alpha-Stable Distributions and Applications. Wiley, New York.

Resnick, S., 1997. Heavy tail modeling and teletraffic data. Ann. Statist. 25, 1805–1869.

Resnick, S., Samorodnitsky, G., Xue, F., 1998. How misleading can sample acf's of stable MA's be? (very!). Technical Report 1209, Department of Operations Research and Industrial Engineering, Cornel University. Available at http://www.orie.cornell.edu/trlist/trlist.html.

Runde, R., 1998. The asymptotic null distribution of the Box–Pierce Q-statistic for random variables with infinite variance. J. Econometrics 78, 205–216.

Samorodnitsky, G., Taqqu, M., 1994. Stable Non-Gaussian Random Processes. Chapman-Hall, New York.