Inference for Lévy-Driven Continuous-Time ARMA Processes

Peter J. Brockwell
Richard A. Davis
Yu Yang

Colorado State University

May 23, 2007

Outline

- Background
- Lévy-driven CARMA processes
- Second order properties
- Canonical representation
- Estimation via the uniformly sampled process
- Estimation for non-negative CAR(1)
- The gamma-driven CAR(1) process
- Recovering the Lévy increments
- Estimation for continuously-observed CAR(1)
- Conclusions

Background

Why study continuous-time models?

- For handling irregularly-spaced data.
- For financial applications—option pricing.
- For taking advantage of the now wide-spread availability of high-frequency data.

Background

Why study continuous-time models?

- For handling irregularly-spaced data.
- For financial applications—option pricing.
- For taking advantage of the now wide-spread availability of high-frequency data.

In recent years, various attempts have been made to use continuous-time in order to capture the so-called stylized features of financial time series

- tail heaviness
- dependence without correlation
- volatility clustering

Background (cont)

log-returns for Nikkei (7/97 – 4/99)

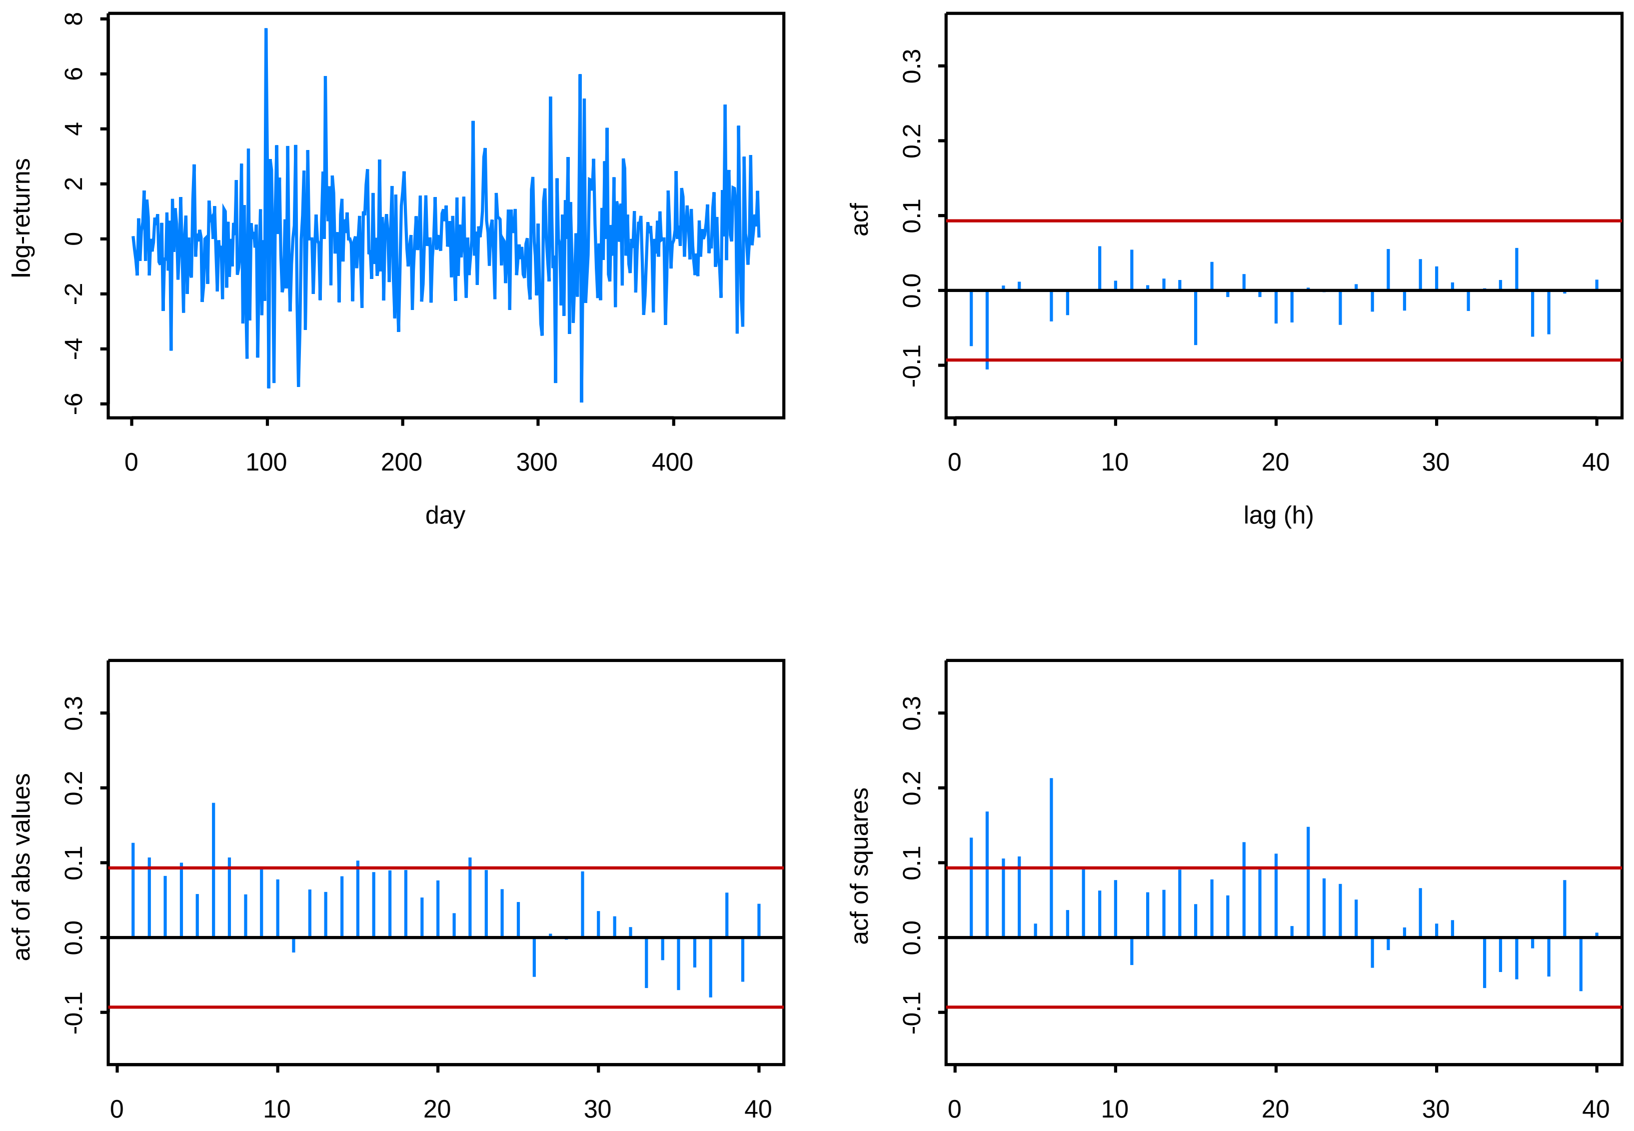

A stochastic volatility model

Barndorff-Nielsen and Shephard (2001) introduced the following SV model for the log-asset price X^*:

$$dX^*(t) = (\mu + \beta V(t))dt + \sqrt{V(t)}dW(t),$$

where $W(t)$ is SBM. The volatility process V is an independent stationary non-negative Lévy-driven Ornstein-Uhlenbeck process satisfying

$$dV(t) + aV(t)dt = \sigma dL(t), \quad a > 0,$$

i.e.,

$$V(t) = \sigma \int_{-\infty}^{t} e^{-a(t-u)} dL(u)$$

with $L(t)$ a Lévy process.

Background (cont)

Daily Volatility Estimates for DM/$ (12/1/86 to 6/30/99) based on 5-minute returns (see Todorov).

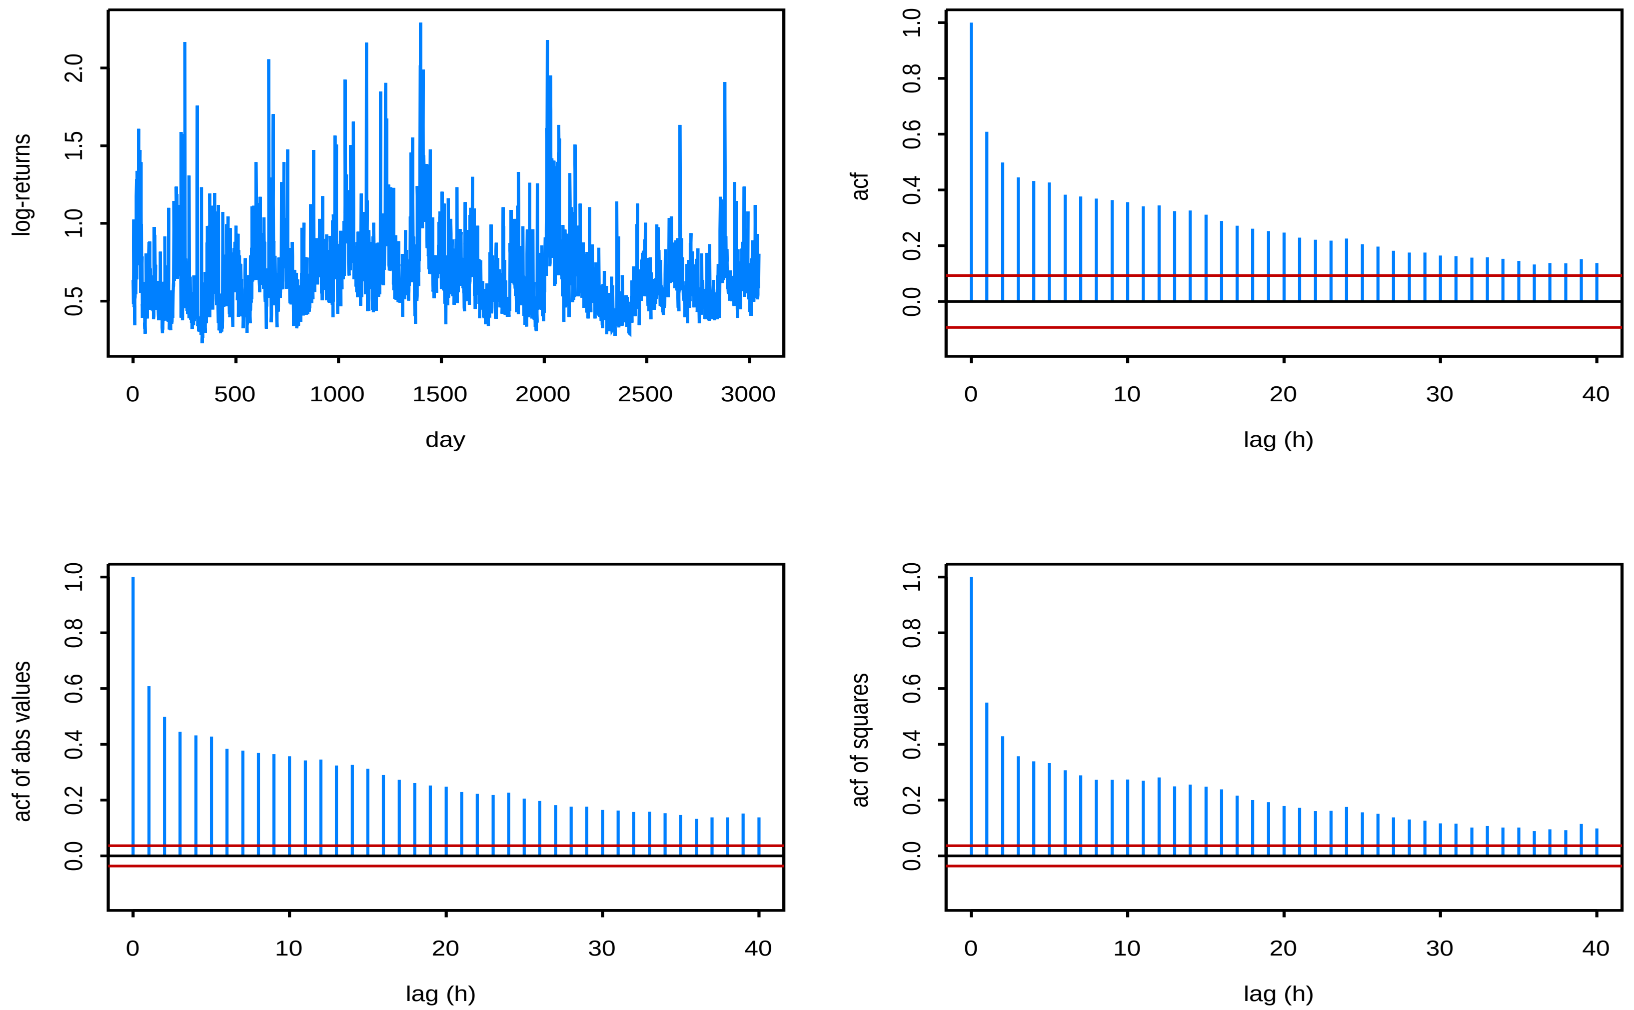

Lévy-driven CARMA processes

The covariance function of Barndorff-Nielsen Shephard model has limited behavior; namely covariance function must decrease exponentially. Instead, we consider the case that V is a subordinator-driven **non-negative continuous-time ARMA (CARMA)**.

Lévy-driven CARMA processes

The covariance function of Barndorff-Nielsen Shephard model has limited behavior; namely covariance function must decrease exponentially. Instead, we consider the case that V is a subordinator-driven **non-negative continuous-time ARMA (CARMA)**.

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \le t \le \infty}, P)$ be a filtered probability space, where $\mathcal{F}_0$ contains all the P-null sets of $\mathcal{F}$ and $(\mathcal{F}_t)$ is right-continuous.

Definition (Lévy Process). $\{L(t), t \ge 0\}$ is an $(\mathcal{F}_t)$-adapted Lévy process if $L(t) \in \mathcal{F}_t$ for all $t \ge 0$ and

- $L(0) = 0$ a.s.,
- $L(t_0), L(t_1) - L(t_0), \ldots, L(t_n) - L(t_{n-1})$ are independent for $0 \le t_0 < t_1 < \cdots < t_n$,
- the distribution of $\{L(s+t) - L(s) : t \ge 0\}$ does not depend on s,
- $L(t)$ is continuous in probability.

Lévy Process

The characteristic function of $L(t)$, $\phi_t(\theta) := E(\exp(i\theta L(t)))$, has the Lévy-Khinchin representation,

$$\phi_t(\theta) = \exp(t\xi(\theta)), \ \theta \in \mathbb{R},$$

where

$$\xi(\theta) = i\theta m - \frac{1}{2}\theta^2\sigma^2 + \int_{\mathbb{R}_0} \left(e^{i\theta x} - 1 - \frac{ix\theta}{1+x^2} \right) \nu(dx),$$

for some $m \in \mathbb{R}$, $\sigma > 0$, and the measure ν is on the Borel subsets of $\mathbb{R}_0 = \mathbb{R} \setminus \{0\}$, known as the *Lévy measure* of the process L, satisfying

$$\int_{\mathbb{R}_0} \frac{u^2}{1+u^2} \nu(du) < \infty.$$

Some examples

$$\xi(\theta) = i\theta m - \frac{1}{2}\theta^2\sigma^2 + \int_{\mathbb{R}_0} \left(e^{i\theta x} - 1 - \frac{ix\theta}{1+x^2} \right) \nu(dx),$$

- $\nu = 0 \Rightarrow$ Brownian motion.
- $m = \sigma^2 = 0$, $\int_{\mathbb{R}_0} \frac{|u|}{1+u^2}\nu(du) < \infty \Rightarrow$ compound Poisson with drift.
- $\nu(du) = \alpha u^{-1} e^{-\beta u} du \Rightarrow$ a gamma process with

$$\xi(\theta) = \int (e^{i\theta x - 1} - 1)\nu(dx) = (1 - i\theta/\beta)^{-\alpha t},$$

- $\nu(du) = \frac{1}{2}\alpha |u|^{-1} e^{-\beta|u|} du \Rightarrow$ a symmetrized gamma process $(L_1 - L_2)$.
- $\xi(\theta) = \exp(-c|\theta|^\alpha)$, $0 < \alpha \leq 2$, $\Rightarrow$ symmetric stable process.

2nd order Lévy-driven CARMA Process

Formally, a CARMA process driven by a Lévy process is a stationary solution of the pth order linear differential equation

$$a(D)Y(t) = \sigma b(D)DL(t), \tag{1}$$

where D denotes differentiation with respect to t,

$$\begin{aligned} a(z) &= z^p + a_1 z^{p-1} + \cdots + a_p, \\ b(z) &= b_0 + b_1 z^1 + \cdots + b_{p-1} z^{p-1}, \end{aligned}$$

$b_q = 1,\ b_j := 0$ for $j > q$, and $\{L(t)\}$ is a second-order Lévy process with $\text{Var}(L(1)) = 1$.

Lévy-driven CARMA Process (cont)

The defining SDE (1) is interpreted through the state-space formulation given by the *observation* and *state* equations,

$$Y(t) = \sigma \mathbf{b}'\mathbf{X}(t), t \geq 0, \quad (2)$$

$$d\mathbf{X}(t) = A\mathbf{X}(t)dt + \mathbf{e}\, dL(t), \quad (3)$$

where

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ -a_p & -a_{p-1} & -a_{p-2} & \cdots & -a_1 \end{bmatrix},$$

$$\mathbf{e}' = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \end{bmatrix}, \quad \text{and}$$

$$\mathbf{b}' = \begin{bmatrix} b_0 & b_1 & \cdots & b_q & 0 & \cdots & 0 & 1 \end{bmatrix}.$$

Lévy-driven CARMA Process (cont)

The solution to (3) satisfies

(4) $$\mathbf{X}(t) = e^{At}\mathbf{X}(0) + \int_0^t e^{A(t-u)}\mathbf{e}\,dL(u).$$

Proposition. If $\mathbf{X}(0)$ is independent of $\{L(t)\}$, then $\{\mathbf{X}(t)\}$ given by (4) is strictly (and weakly) stationary if and only if the eigenvalues of the matrix A all have strictly negative real parts and $\mathbf{X}(0) \sim \int_0^\infty e^{Au}\mathbf{e}\,dL(u)$.

Remark. It is easy to check that the eigenvalues of the matrix A are the zeroes of the autoregressive polynomial $a(z)$.

Lévy-driven CARMA Process (cont)

Sometimes it is convenient to define the CARMA process for all real values of t.

Extension to all t. Let $\{M(t),\, 0 \le t < \infty\}$ be an independent copy of L and set

$$L^*(t) = L(t) I_{[0,\infty)}(t) - M(-t-) I_{(-\infty,0]}(t)\,.$$

If the eigenvalues of A have negative real parts, then

$$\mathbf{X}(t) = \int_{-\infty}^{t} e^{A(t-u)} \mathbf{e}\, dL^*(u)\,. \tag{5}$$

is a strictly stationary process satisfying

$$\mathbf{X}(t) = e^{A(t-s)} \mathbf{X}(s) + \int_{s}^{t} e^{A(t-u)} \mathbf{e}\, dL^*(u)\,.$$

Lévy-driven CARMA Process (cont)

Definition (Causal CARMA Process). If the eigenvalues of A have negative real parts, then the CARMA process Y is the strictly stationary process

$$Y(t) = \sigma \mathbf{b}' \mathbf{X}(t)$$

where

$$\mathbf{X}(t) = \int_{-\infty}^{t} e^{A(t-u)} \mathbf{e} \, dL(u),$$

i.e.,

$$Y(t) = \sigma \int_{-\infty}^{t} \mathbf{b}' e^{A(t-u)} \mathbf{e} \, dL(u) \, .$$

That is, $\{Y(t)\}$ is a **causal** function of $\{L(t)\}$,

$$Y(t) = \sigma \int_{-\infty}^{\infty} g(t-u) \, dL(u) \, , \quad \text{where} \quad g(t) = \begin{cases} \sigma \mathbf{b}' e^{At} \mathbf{e}, & t > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Second-order Properties

Using the causal representation of $Y(t)$, the ACVF of Y is

$$\gamma(h) = \sigma^2 \int_{-\infty}^{\infty} \tilde{g}(h-u) g(u) \, du \,,$$

where $\tilde{g}(u) = g(-x)$. After some calculation, one can show that

$$\gamma(h) = \frac{\sigma^2}{2\pi} \int_{-\infty}^{\infty} e^{i\omega h} \left| \frac{b(i\omega)}{a(i\omega)} \right|^2 \, d\omega,$$

i.e., Y has *rational* spectral density $f(\omega) = \frac{\sigma^2}{2\pi} \left| \frac{b(i\omega)}{a(i\omega)} \right|^2$.

Remark. Gaussian processes with rational spectral density have been of interest for many years. (See extensive study by Doob (1944) and a nice paper by Pham Din Duan (1977).) The SDE approach to such processes can be found in the engineering literature and was employed by Jones (1978) for modeling irregularly-spaced data.

Canonical Representation of a CARMA

When the zeroes $\lambda_1, \ldots, \lambda_p$ of the *causal* AR polynomial $a(z)$ are distinct, then the *kernel* function g and ACVF γ have the special form

$$g(h) = \sigma \sum_{r=1}^{p} \frac{b(\lambda_r)}{a'(\lambda_r)} e^{\lambda_r h} I_{[0,\infty)}(h) \text{ and } \gamma(h) = \sigma^2 \sum_{r=1}^{p} \frac{b(\lambda_r)b(-\lambda_r)}{a'(\lambda_r)a(-\lambda_r)} e^{\lambda_r |h|}.$$

Now defining $\alpha_r = \sigma b(\lambda_r)/a'(\lambda_r), \quad r = 1, \ldots, p$, we can write

$$Y(t) = \sum_{r=1}^{p} Y_r(t)\,,$$

where Y_r is the CAR(1) process,

$$Y_r(t) = \int_{-\infty}^{t} \alpha_r e^{\lambda_r (t-u)}\, dL(u)\,.$$

Canonical Representation (cont)

Equivalently,

$$Y(t) = [1, \ldots, 1]\mathbf{Y}(t), \quad t \geq 0,$$

where $\mathbf{Y}$ is the solution of

$$d\mathbf{Y}(t) = \text{diag}[\lambda_i]_{i=1}^{p} \mathbf{Y}\, dt + \sigma B R^{-1} \mathbf{e}\, dL$$

with $\mathbf{Y}(0) = \sigma B R^{-1} \mathbf{X}(0)$, $B = \text{diag}[b(\lambda_i)]$, and $R = [\lambda_j^{i-1}]$.

Remark. Simulation of a CARMA(p, q) process with distinct AR roots can be achieved by the much simpler problem of simulating component CAR(1) processes and adding them together.

Joint Distribution

From the representation, $Y(t) = \int_{-\infty}^{t} g(t-u)\, dL(u)$, the marginal distribution of Y has cumulant generating function

$$\log E\left(\exp\left(i\theta Y(t)\right)\right) = \int_{-\infty}^{\infty} \xi(\theta g(u))\, du\,.$$

Using independence of the increments, one can easily calculate the joint cgf of the fidis.

- $\nu = 0 \Rightarrow$ Gaussian CARMA(p, q).
- $\{L(t)\}$ compound Poisson with bilateral exponential jumps $\Rightarrow$ (in the CAR(1) case) that $\{Y(t)\}$ has marginal cfg, $\kappa(\theta) = -\frac{\lambda}{2a_1} \ln\left(1 + \frac{\theta^2}{\beta^2}\right)$, i.e., $Y(t)$ has a symmetrized gamma distribution (bilateral exponential if $\lambda = 2a_1$).
- For CAR(1) with non-negative Lévy input, see Barndorff-Nielsen and Shephard (2001) and storage theory literature.

Inference for CARMA processes

1. For linear Gaussian CARMA processes, MLE based on observations $Y(t_1), \ldots, Y(t_n)$ can be easily carried out using the state-space representation (see Jones (1981)).

2. For both the linear Lévy-driven CARMA and Gaussian CTAR processes, the likelihood can be computed using the state-space represntation and then optimized.

Estimation via the Sampled Process

- If Y is a Gaussian CARMA process, then it is well known (e.g., Doob (1944), Phillips (1959), Brockwell (1995)), that the sampled process $Y(nh),\ n = 0, \pm 1, \ldots,$ for fixed spacing h is a *strict* Gaussian ARMA(r, s) process with $0 \leq s < r \leq p$.

- If L is non-Gaussian, the sampled process will have the same spectral density (and hence ACVF) as the analogous Gaussian CARMA process. So from a second order perspective, the two sampled processes are identical. However, (except in the CAR(1) case), the non-Gaussian CARMA will not generally be a *strict* ARMA process.

Estimation (cont)

CAR(1) Example. If Y is the CAR(1) process, then the sampled process is the strict AR(1) process

$$Y_n^{(h)} = \phi Y_{n-1}^{(h)} + Z_n, \; n = 0, 1, \ldots,$$

where $\phi = \exp(-ah)$ and

$$Z_n = \sigma \int_{(n-1)h}^{nh} e^{-a(nh-u)} \, dL(u) \, .$$

The noise sequence $\{Z_n\}$ is iid and Z_n has the infinitely divisible distribution with cgf

$$\int_0^h \xi(\sigma\theta e^{-au}) \, du \, ,$$

where $\xi(\theta)$ is the log-characterstic function of $L(1)$.

Estimation (cont)

For the CARMA(p,q) process with $p > 1$, the situation is more complicated. If the AR roots $\lambda_1, \ldots, \lambda_p$ are all distinct then, from the canonical representation, the sampled process is

$$Y(nh) = \sum_{r=1}^{p} Y_r(nh)\,,$$

where $\{Y_r(nh)\}$ is the strict AR(1) process

$$Y_r(nh) = e^{\lambda_r h} Y_r((n-1)h) + Z_r(n), \quad n = 0, \pm 1, \ldots,$$

with

$$Z_r(n) = \alpha_r \int_{(n-1)h}^{nh} e^{\lambda_r (nh-u)}\, dL(u)\,.$$

and

$$\alpha_r = \sigma \frac{b(\lambda_r)}{a'(\lambda_r)}.$$

Estimation for non-negative CAR(1)

Let Y be the CAR(1) process driven by the Lévy process $\{L(t), t \geq 0\}$ with non-negative increments, i.e., Y is the stationary solution of the stochastic differential equation,

$$dY(t) + aY(t)dt = \sigma dL(t).$$

For any $h > 0$, the sampled process $\{Y_n^{(h)} := Y(nh), n = 0, 1, \ldots\}$ is a discrete-time AR(1) process satisfying

$$Y_n^{(h)} = \phi Y_{n-1}^{(h)} + Z_n, \; n = 0, 1, \ldots,$$

where $\phi = \exp(-ah)$ (obviously $0 < \phi < 1$), and

$$Z_n = \sigma \int_{(n-1)h}^{nh} e^{-a(nh-u)} \, dL(u) \, .$$

the noise sequence $\{Z_n\}$ is iid and positive since L has stationary, independent, and positive increments.

Estimation for non-negative CAR(1)

If the process $\{Y(t), 0 \le t \le T\}$ is observed at times $0, h, \dots, Nh$, where $N = [T/h]$, i.e., N is the integer part of T/h, then, since the innovations Z_n of the process $\{Y_n^{(h)}\}$ are non-negative and $0 < \phi < 1$, we can use the highly efficient Davis-McCormick estimator of ϕ,

$$\hat{\phi}_N^{(h)} = \min_{1 \le n \le N} Y_n^{(h)} / Y_{n-1}^{(h)} .$$

To obtain the asymptotic distribution of $\hat{\phi}_N^{(h)}$ with h fixed, we need to suppose the distr F of Z_n satisfies $F(0) = 0$ and that F is regularly varying at zero with exponent α, i.e.,

$$\lim_{t \downarrow 0} \frac{F(tx)}{F(t)} = x^\alpha, \quad \text{for all } x > 0.$$

(These conditions are satisfied by the gamma-driven CAR(1) process as we shall show later.)

Estimation for non-negative CAR(1)

Under these conditions on F, the results of Davis and McCormick (1989) imply that $\hat{\phi}_N^{(h)} \to \phi$ a.s. as $N \to \infty$ with h fixed and that

$$\lim_{N\to\infty} P\left[k_N^{-1}(\hat{\phi}_N^{(h)} - \phi)c_\alpha \le x\right] = G_\alpha(x)$$

where $k_N = F^{-1}(N^{-1}), c_\alpha = (EY_1^{(h)\alpha})^{1/\alpha}$ and G_α is the Weibull distribution function,

$$G_\alpha(x) = \begin{cases} 1 - \exp\{-x^\alpha\}, & \text{if } x \ge 0, \\ 0, & \text{if } x < 0. \end{cases}$$

Estimation for non-negative CAR(1)

From the observations $\{Y_n^{(h)},\ n = 0, 1, \ldots, N\}$, we thus obtain the estimator $\hat{\phi}_N^{(h)}$, and hence the estimator of a is

$$\hat{a}_N^{(h)} = -h^{-1} \log \hat{\phi}_N^{(h)} .$$

Using a Taylor series approximation, we find that

$$\lim_{N \to \infty} P\left[(-h)e^{-ah} k_N^{-1}(\hat{a}_N^{(h)} - a)c_\alpha \leq x\right] = G_\alpha(x),$$

where G_α is the Weibull distribution specified above.
Since $\mathrm{Var}(Y^{(h)}) = \sigma^2/(2a)$, we use the estimator

$$\hat{\sigma}_N^2 = \frac{2\hat{a}_N^{(h)}}{N} \sum_{i=0}^{N} (Y_i^{(h)} - \overline{Y}_N^{(h)})^2$$

to estimate σ^2,

Gamma-driven CAR(1)

Suppose L is a standardized gamma process, i.e., $L(t)$ has the gamma density $f_{L(t)}$ with exponent γt, and scale-parameter $\gamma^{-1/2}$, mean $\gamma^{1/2} t$ and variance t. The Laplace transform of $L(t)$ is

$$\begin{aligned} \tilde{f}_{L(t)}(s) &= E\exp(-sL(t)) \\ &= \exp\{-t\Phi(s)\}, \quad \mathcal{R}(s) \geq 0, \end{aligned}$$

where $\Phi(s) = \gamma \log(1 + \beta s)$, $\beta = \gamma^{-1/2}$, and $\gamma > 0$.

Theorem 1. For the gamma-driven CAR(1) process, we have $\hat{a}_N^{(h)} \to a$ a.s. and

$$\lim_{N\to\infty} P\left[(-h)e^{-ah}k_N^{-1}(\hat{a}_N^{(h)} - a)c_\alpha \leq x\right] = G_\alpha(x)$$

where $\alpha = \gamma h$,

$$k_N^{-1} \sim (\sigma\beta)^{-1}[\Gamma(\gamma h + 1)]^{-1/(\gamma h)} e^{.5ah} N^{1/(\gamma h)},$$

and c_α is computed numerically using a result of Brockwell and Brown '78.

Gamma-driven CAR(1)

Examining the normalizing constant k_N^{-1}, we find that

$$\lim_{h \to 0} \lim_{N \to \infty} \frac{k_N^{-1}}{N^{1/(\gamma h)}} = (\sigma\beta)^{-1} e^{\gamma_E}$$

where γ_E is the Euler-Mascheroni constant.

Convergence is thus extremely fast for large N and small h.

Gamma-driven CAR(1) Process

Example. We now illustrate the estimation procedure with a simulated example. The gamma-driven CAR(1) process defined by,

$$DY(t) + 0.6Y(t) = DL(t), \qquad t \in [0, 5000], \tag{7}$$

was simulated at times $0, 0.001, 0.002, \ldots, 5000$, using an Euler approximation. The parameter γ of the standardized gamma process was 2. The process was then sampled at intervals $h = 0.01$, $h = 0.1$ and $h = 1$ by selecting every 10^{th}, 100^{th} and 1000^{th} observation respectively. We generated 100 such realizations of the process and applied the above estimation procedure to generate 100 independent estimates, for each h, of the parameters a and σ. The sample means and standard deviations of these estimators are shown in Table 1, which illustrates the remarkable accuracy of the estimators.

Gamma-driven CAR(1) Process (Cont.)

Table 1. Estimated parameters based on 100 replicates on $[0, 5000]$ of the gamma-driven CAR(1) process (7) with $\gamma = 2$, observed at times $nh, n = 0, \ldots, [T/h]$.

Spacing	Parameter	Gamma increments	
		Sample mean of estimators	Sample std deviation of estimators
h=1	a	0.59269	0.00381
	σ	0.99796	0.01587
h=0.1	a	0.59999	0.00000
	σ	1.00011	0.01281
h=0.01	a	0.60000	0.00000
	σ	0.99990	0.01175

Recovering the Lévy Increments

In order to suggest an appropriate parametric model for L and to estimate the parameters, it is important to recover an approximation to L form the observed data. If the CAR(1) process is observed continuously on $[0, T]$, we have

$$L(t) = \sigma^{-1} \left[Y(t) - Y(0) + a \int_0^t Y(s) ds \right].$$

Recovering the Lévy Increments

In order to suggest an appropriate parametric model for L and to estimate the parameters, it is important to recover an approximation to L form the observed data. If the CAR(1) process is observed continuously on $[0,T]$, we have

$$L(t) = \sigma^{-1}\left[Y(t) - Y(0) + a\int_0^t Y(s)ds\right].$$

Replacing the CAR(1) parameters by their estimators and the integral by a trapezoidal approximation, we obtain the estimator for the Lévy increments $\Delta L_n^{(h)} := L(nh) - L((n-1)h)$ on the interval $((n-1)h, nh]$, given by

$$\Delta \hat{L}_n^{(h)} = \hat{\sigma}_N^{-1}\left[Y_n^{(h)} - Y_{n-1}^{(h)} + \hat{a}_N^{(h)} h (Y_n^{(h)} + Y_{n-1}^{(h)})/2\right]. \qquad (6)$$

Gamma-driven CAR(1) Process (Cont.)

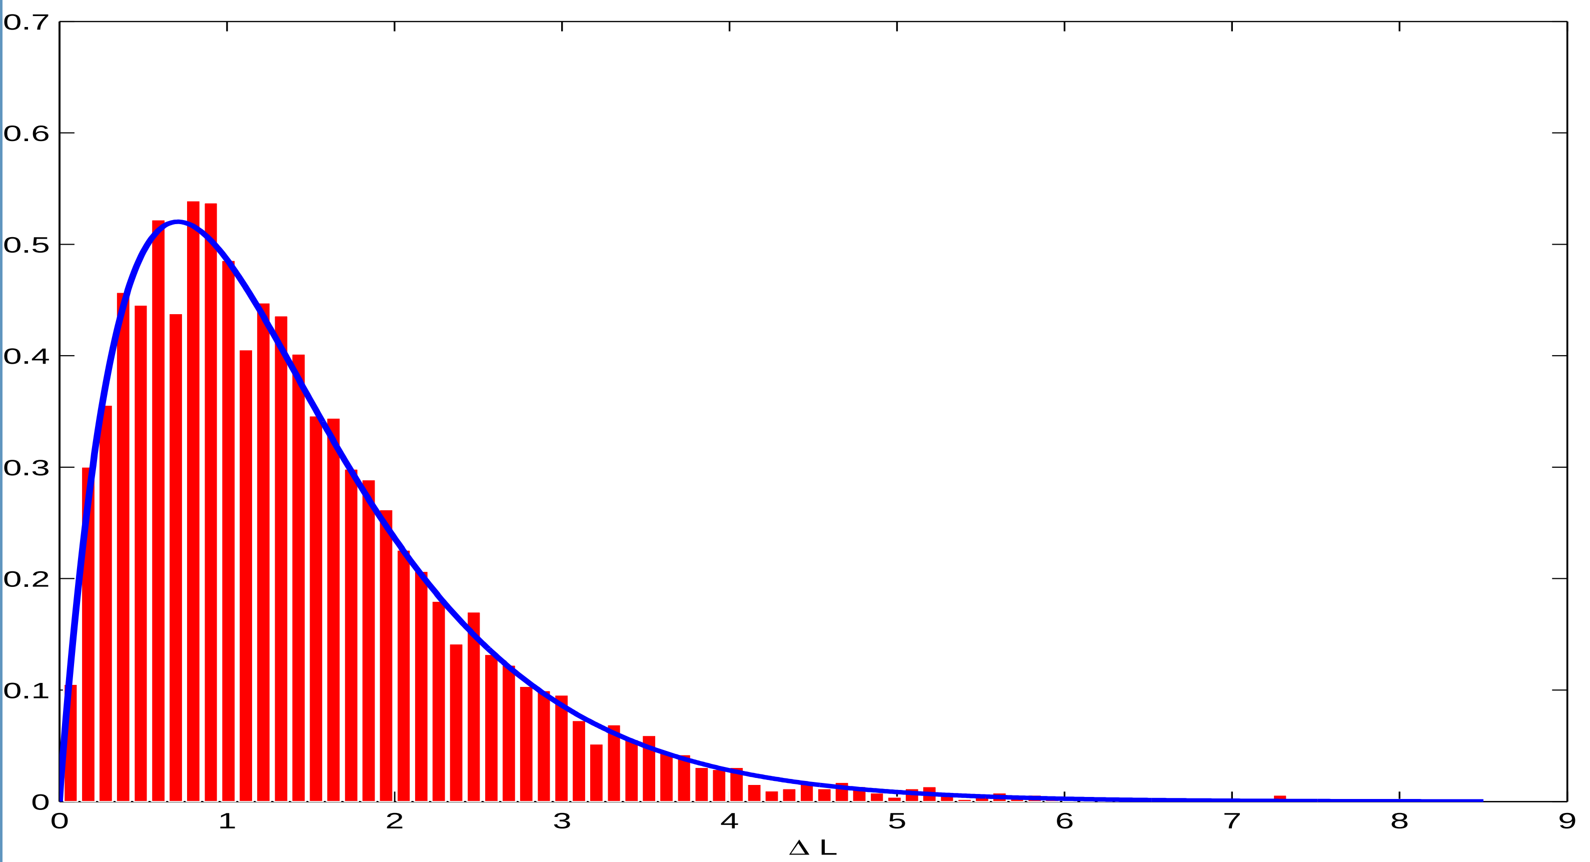

Figure 1: The probability density of the increments per unit time of the standardized Lévy process and the histogram of the estimated increments from a realization of the CAR(1) process (7).

Gamma-driven CAR(1) Process (Cont.)

Table 2. Estimated parameter of the standardized driving Lévy process.

Spacing	Parameter	Sample mean of estimators	Sample std deviation of estimators
$h = 1$	γ	1.99598	0.05416
$h = 0.1$	γ	2.00529	0.03226
$h = 0.01$	γ	2.00547	0.02762

Estimation for Cont-Observed CAR(1)

For a continuously observed realization on $[0, T]$ of a CAR(1) process driven by a non-decreasing Lévy process with drift $m = 0$, the value of a can be identified exactly with probability 1. This contrast strongly with the case of a Gaussian CAR(1) process. This results is a corollary of the following theorem.

Theorem 2. If the CAR(1) process $\{Y(t), t \geq 0\}$ is driven by a non-decreasing Lévy process L with drift m and Lévy measure ν, then for each fixed t,

$$\frac{Y(t+h) - Y(t)}{h} + aY(t) \rightarrow m \quad a.s. \quad \text{as } h \downarrow 0.$$

.

Estimation for Cont-Observed CAR(1)

Corollary. If $m = 0$ in the Theorem 2 (this is the case if the point zero belongs to the closure of the support of $L(1)$), then with probability 1,

$$a = \sup_{0 \le s < t \le T} \frac{\log Y(s) - \log Y(t)}{t - s}.$$

For observations available at times $\{nh : n = 0, 1, 2, \ldots, [T/h]\}$, our estimator can be expressed as

$$\hat{a}_T^{(h)} = \sup_{0 \le n < [T/h]} \frac{\log Y(nh) - \log Y((n+1)h)}{h}.$$

The analogous estimator, based on closely but irregularly spaced observations at times $t_1, t_2, \ldots, t_N$ such that $0 \le t_1 < t_2 < \cdots < t_N \le T$, is

$$\hat{a}_T = \sup_n \frac{\log Y(t_n) - \log Y(t_{n+1})}{t_{n+1} - t_n}.$$

Estimation for Cont-Observed CAR(1)

Remark. We have shown that, if the drift of the driving Lévy process is zero and T is any finite positive time, both estimators,

$$\hat{a}_T^{(h)} = \sup_{0 \le n < [T/h]} \frac{\log Y(nh) - \log Y((n+1)h)}{h}.$$

and

$$\hat{a}_T = \sup_n \frac{\log Y(t_n) - \log Y(t_{n+1})}{t_{n+1} - t_n}.$$

converge almost surely to a as the maximum spacing between successive observations converges to zero.

Conclusions

- We found a highly efficient method, based on observations at times $0, h, 2h, \ldots, Nh$, for estimating the parameters of a stationary Ornstein-Uhlenbeck process $\{Y(t)\}$ driven by a non-decreasing Lévy process.

Conclusions

- We found a highly efficient method, based on observations at times $0, h, 2h, \ldots, Nh$, for estimating the parameters of a stationary Ornstein-Uhlenbeck process $\{Y(t)\}$ driven by a non-decreasing Lévy process.
- Under specific conditions on the driving Lévy process, the asymptotic behavior of the estimators can be determined.

Conclusions

- We found a highly efficient method, based on observations at times $0, h, 2h, \ldots, Nh$, for estimating the parameters of a stationary Ornstein-Uhlenbeck process $\{Y(t)\}$ driven by a non-decreasing Lévy process.
- Under specific conditions on the driving Lévy process, the asymptotic behavior of the estimators can be determined.
- If the sample spacing h is small, we used a discrete approximation to the exact integral representation of $L(t)$ in terms of $\{Y(s), s \leq t\}$ to estimate the increments of the driving Lévy process, and hence to estimate the parameters of the Lévy process.

Conclusions

- We found a highly efficient method, based on observations at times $0, h, 2h, \ldots, Nh$, for estimating the parameters of a stationary Ornstein-Uhlenbeck process $\{Y(t)\}$ driven by a non-decreasing Lévy process.
- Under specific conditions on the driving Lévy process, the asymptotic behavior of the estimators can be determined.
- If the sample spacing h is small, we used a discrete approximation to the exact integral representation of $L(t)$ in terms of $\{Y(s), s \leq t\}$ to estimate the increments of the driving Lévy process, and hence to estimate the parameters of the Lévy process.
- Examples suggest extremely good performance of the estimates.

Conclusions (cont)

- If the driving Levý process has no drift, then CAR(1) coefficient a is determined almost surely by a continuously observed realization of Y on any interval $[0, T]$.

Conclusions (cont)

- If the driving Lévy process has no drift, then CAR(1) coefficient a is determined almost surely by a continuously observed realization of Y on any interval $[0, T]$.
- The expression for a suggests an estimator based on discrete observations of Y which, for uniformly spaced observations, is the same as the estimator developed above and establishes the almost sure convergence of our estimator for any fixed T as $h \to 0$.

Conclusions (cont)

- If the driving Levý process has no drift, then CAR(1) coefficient a is determined almost surely by a continuously observed realization of Y on any interval $[0, T]$.
- The expression for a suggests an estimator based on discrete observations of Y which, for uniformly spaced observations, is the same as the estimator developed above and establishes the almost sure convergence of our estimator for any fixed T as $h \to 0$.
- Analogous procedures for non-negative Lévy-driven continuous-time ARMA processes are currently being investigated.