

Sharp Gaussian Approximation Bounds for Linear Systems with α -stable Noise

Marina Riabiz, Tohid Ardeshiri, Ioannis Kontoyiannis and Simon Godsill

Department of Engineering, University of Cambridge, Cambridge, CB2 1PZ, UK

mr622@cam.ac.uk, ta417@cam.ac.uk, ik355@cam.ac.uk, sjg@eng.cam.ac.uk

Abstract—We report the results of several theoretical studies into the convergence rate for certain random series representations of α -stable random variables, which are motivated by and find application in modelling heavy-tailed noise in time series analysis, inference, and stochastic processes. The use of α -stable noise distributions generally leads to analytically intractable inference problems. The particular version of the Poisson series representation invoked here implies that the resulting distributions are “conditionally Gaussian,” for which inference is relatively straightforward, although an infinite series is still involved. Our approach is to approximate the residual (or “tail”) part of the series from some point, $c > 0$, say, to ∞ , as a Gaussian random variable. Empirically, this approximation has been found to be very accurate for large c . We study the rate of convergence, as $c \rightarrow \infty$, of this Gaussian approximation. This allows the selection of appropriate truncation parameters, so that a desired level of accuracy for the approximate model can be achieved. Explicit, nonasymptotic bounds are obtained for the Kolmogorov distance between the relevant distribution functions, through the application of probability-theoretic tools. The theoretical results obtained are found to be in very close agreement with numerical results obtained in earlier work.

I. INTRODUCTION

A. Background

Time series arising in the natural sciences, engineering and finance [1], [2] are frequently characterised by high data rates and irregular sampling, a situation well-represented by continuous-time state-space models. Perhaps the simplest, most tractable such model is the linear diffusion model with linear observations at discrete times $\{t_i\}$,

$$d\mathbf{x}_t = \mathbf{A}\mathbf{x}_t dt + \mathbf{h} d\ell_t, \quad y_{t_i} = \mathbf{b}'\mathbf{x}_{t_i} + v_{t_i}$$

where $\mathbf{x}_t = [x_{1,t}, \dots, x_{P,t}]'$ is the state, \mathbf{A} is a $P \times P$ matrix describing the interaction of the components of \mathbf{x}_t , \mathbf{h} is a P -dimensional vector describing the effects of the noise process $\{d\ell_t\}$, \mathbf{b} is a P -dimensional vector and $\{v_{t_i}\}$ is the observation noise process. A wide range of results have been developed in the literature for the case when $\{\ell_t\}$ is a Brownian motion [3], [4]. However, such models cannot account for the heavy tails and the large “jumps” in the state process often observed in applications. In such cases, the use of linear state-space systems driven by (non-Gaussian) Lévy processes is more appropriate, since these models do exhibit heavy-tailed, discontinuous behaviour [5]–[7]. Despite the simple characterization of Lévy processes [8], specialized probabilistic tools are required in their analysis. The main difficulty stems from lack of closed-form expressions for many relevant quantities of interest. We refer to [9]–[11] for an approximate framework

for simulation. The family of α -stable Lévy processes is of special importance, in that the class of α -stable laws are the natural limiting distributions in the generalized central limit theorem (CLT) with heavy-tailed summands [12]; also see [13], [14] for other relevant background. The self-similarity of α -stable Lévy processes [13] implies that transition densities, although still intractable, all come from the same α -stable family, and hence they may be considered to be the natural first candidates in generalising the classical Gaussian process framework to the heavy-tailed case.

In this paper we provide convergence results for the Poisson series representation (PSR) of α -stable random variables. These find application in discrete-time time-series models with α -stable disturbances, and in the above continuous time model when the observation noise $\{v_{t_i}\}$ is α -stable. See [15]–[18] for PSR approaches to the full α -stable Lévy process. Despite extensive earlier work on properties of α -stable systems, there are few results on likelihood or Bayesian parameter inference for such linear models, see e.g. [19]–[21] for some examples.

Our approach to the inference problem involves an auxiliary-variables version of Bayesian Monte Carlo approaches. In our models, for example, where part of the state is conditionally linear and Gaussian, efficient Rao-Blackwellised versions of Sequential Monte Carlo (SMC) can be applied [22], [23], or MCMC in Bayesian parameter inference may be used [16], [24], [25].

We note that [9] presents an alternative Gaussian approximation to ours, based on simulation only of jumps greater than some magnitude ϵ , including a CLT and convergence rates. We will present comparisons between the two approaches in future work.

B. The α -stable Distribution

We adopt standard notation as in the text [13]. Let $X \sim \mathcal{S}_\alpha(\sigma, \beta, \mu)$ denote an α -stable distributed random variable, where μ is the location parameter, $\sigma > 0$ is the scale parameter, $\beta \in [-1, 1]$ is the skewness parameter, and $\alpha \in (0, 2)$ is the tail parameter. Recall that the probability density function (pdf) of X decays like $1/|x|^{1+\alpha}$ as $|x| \rightarrow \infty$, which is a consequence of α -stable version of the CLT. The polynomial tails can be viewed as the result of the presence of extreme values, with more extreme values (and hence heavier tails) appearing for smaller values of α . When $\beta = 0$ the pdf of X is symmetric, while $\beta = \pm 1$ correspond to the fully left or right skewed cases, respectively.

Also we recall that the pdf of α -stable distributions is not available in closed form, except in few special

cases (Gaussian, Cauchy and Lévy distributions), obtained for specific choices of α and β . Although, as mentioned above, this complicates the analysis of questions related to inference, there is a substantial body of earlier work developing practically applicable methods.

On the other hand, an α -stable random variable has characteristic function (CF) $\phi_X(s) := \mathbb{E}[\exp(isX)]$, for $s \in \mathbb{R}$, explicitly given by:

$$\log(\phi_X(s)) = \begin{cases} -\sigma^\alpha |s|^\alpha \left\{ 1 - i\beta \operatorname{sgn}(s) \tan \frac{\pi\alpha}{2} \right\} + i\mu s & \text{if } \alpha \neq 1, \\ -\sigma |s| \left\{ 1 + i\beta \operatorname{sgn}(s) \frac{2}{\pi} \log |s| \right\} + i\mu s & \text{if } \alpha = 1. \end{cases}$$

For the sake of simplicity, we concentrate only on the non-singular cases, i.e., on values of $\alpha \neq 1$.

Thus motivated, in this paper we consider the PSR representation of α -stable random variables, and adopt the series-truncation approach of the earlier work referenced above. Our main goal is to obtain simple, explicit, quantitative bounds on the quality of approximating the tail of the series by a Gaussian, or, equivalently, of treating the PSR as a conditionally Gaussian representation of the α -stable distribution. Our first main contribution is the derivation of closed-form expressions for the relevant CFs, and our second main contribution is the development of non-asymptotic bounds on the distance between the PSR residual and an appropriately defined Gaussian. Only brief outlines of the proofs are given here; complete details can be found in the full paper [26].

II. PSR AND THE CONDITIONALLY GAUSSIAN REPRESENTATION

The PSR of the α -stable random variable $X \sim S_\alpha(\sigma, \beta, 0)$, $\alpha \in (0, 2)$, $\alpha \neq 1$, is the following random series,

$$X \stackrel{\mathcal{D}}{=} \sum_{j=1}^{\infty} \Gamma_j^{-1/\alpha} W_j - \mathbb{E}[W_1] b_j^{(\alpha)},$$

where $\stackrel{\mathcal{D}}{=}$ denotes equality in distribution, $\{\Gamma_j\}_{j=1}^{\infty}$ are the arrival times of a unit rate Poisson process, and $\{W_j\}_{j=1}^{\infty}$ are independent and identically distributed (i.i.d.) random variables independent of $\{\Gamma_j\}_{j=1}^{\infty}$, with $\mathbb{E}[|W_1|^\alpha] < \infty$; see [13, p.28] for a detailed exposition. The coefficients $b_j^{(\alpha)}$ are non-zero only if $\alpha \in [1, 2)$ and for this case they are readily computed and have a telescoping structure. We further refer to [13] for the non-linear transformations that map the moments of W_j and α to the parameters β, σ .

From the PSR it follows that, if we choose the distribution of the $\{W_j\}$ to be i.i.d. with $W_j \sim \mathcal{N}(\mu_W, \sigma_W^2)$, we can write a conditionally Gaussian model for X as,

$$X | \{\Gamma_j\}_{j=1}^{\infty} \sim \mathcal{N}(\mu_W m, \sigma_W^2 S^2), \quad (1)$$

where $m := \sum_{j=1}^{\infty} \Gamma_j^{-1/\alpha} - b_j^{(\alpha)}$ and $S^2 := \sum_{j=1}^{\infty} \Gamma_j^{-2/\alpha}$ are treated as auxiliary random variables. Figure 1 shows the first 100 terms of sampled PSR realizations, for different values of the parameters α, μ_W, σ_W .

A. Truncation of the PSR and the Approximate Conditionally Gaussian Representation

While the exact representation of the stable law (1) is theoretically very appealing, in practice it is computationally intractable because of the infinite summations involved in the definitions of m and S . Given that the summands of the PSR are stochastically decaying, the approach we adopt is to truncate the series to values of $\Gamma_j \leq c$, where $c \geq 0$ is a truncation parameter, and to approximate the distribution of the residual term of the series by an appropriately chosen Gaussian. Then, the PSR can be split as,

$$X = X_{(0,c)} + R_{(c,\infty)}, \quad (2)$$

where $X_{(0,c)}$ is the truncated PSR,

$$X_{(0,c)} := \sum_{j:\Gamma_j \in [0,c]} W_j \Gamma_j^{-1/\alpha}, \quad (3)$$

and $R_{(c,\infty)}$ is the obvious tail of the series.

The residual $R_{(c,\infty)}$ is not Gaussian. However, it can be shown that it converges to a Gaussian for large c , in that, under mild conditions,

$$Z_{(c,\infty)} := (R_{(c,\infty)} - m_{(c,\infty)}) / S_{(c,\infty)}$$

converges to the standard Gaussian distribution, asymptotically as $c \rightarrow \infty$. A first result [16] studies the deterministic case $\sigma_W^2 = 0$ and our more recent results give the general case [27]. Such a CLT served to justify our adoption of the following Gaussian approximation of the PSR residual, $\hat{R}_{(c,\infty)}$, for practical inference procedures,

$$R_{(c,\infty)} \stackrel{\text{approx}}{\sim} \hat{R}_{(c,\infty)} \sim \mathcal{N}(m_{(c,\infty)}, S_{(c,\infty)}^2),$$

where $\stackrel{\text{approx}}{\sim}$ means that the distribution of $R_{(c,\infty)}$ converges to the Gaussian on the right-hand side, as $c \rightarrow \infty$, and $m_{(c,\infty)}$ and $S_{(c,\infty)}^2$ are, respectively, the mean and the variance of $R_{(c,\infty)}$, with explicit expressions provided in [27]. It follows that, by analogy with (2), we can introduce the random variable,

$$\hat{X} := X_{(0,c)} + \hat{R}_{(c,\infty)},$$

that converges in distribution to X , as $c \rightarrow \infty$ and it is this approximating random variable that is used in our inference procedures.

In fact, our CLT result for the PSR residual does not assume Gaussianity of the variables W_j . Therefore, even in this case we have the following *approximate* conditionally Gaussian representation for the α -stable model,

$$X | \{\Gamma_j \leq c\} \stackrel{\text{approx}}{\sim} \mathcal{N}(m_{(0,c)} + m_{(c,\infty)}, S_{(0,c)}^2 + S_{(c,\infty)}^2),$$

where we view $m_{(0,c)} := \mu_W \sum_{j:\Gamma_j \in [0,c]} \Gamma_j^{-1/\alpha}$, and $S_{(0,c)}^2 := \sigma_W^2 \sum_{j:\Gamma_j \in [0,c]} \Gamma_j^{-2/\alpha}$ as auxiliary random variables, which can now be generated exactly by direct sampling of the truncated Poisson process.

This approximate conditionally Gaussian structure illustrates the power of the proposed approximation: The inference methods that are valid for the exact PSR can be used for such an approximation, with the quality of the approximation controlled directly by the truncation parameter c . The inference schemes [15]–[18], [24], [25]

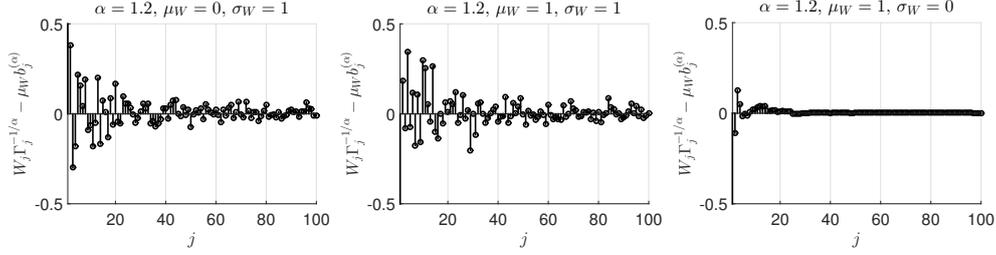


Fig. 1. First 100 terms of PSR realizations, with $W_j \sim \mathcal{N}(\mu_W, \sigma_W^2)$ for $\alpha = 1.2$ and different values for μ_W and σ_W^2 .

referenced in the Introduction were based on this approximation; however, we had no numerical measure of *how good* the approximation was and how it might vary with c, α, β .

In this paper we provide explicit probability-theoretic results that can guide the practical choice of the truncation parameter c . Specifically, we provide non-asymptotic bounds on the approximation error between the corresponding cumulative distribution functions (cdfs) for finite values of c , by making use of Fourier-inversion theorems, summarized in the following.

III. DISTANCES BETWEEN DISTRIBUTIONS: THE SMOOTHING LEMMA

Let $c \geq 0$ be the PSR truncation parameter as in (2), (3). Suppose S_c and T are random variables with CFs $\phi_{S_c}(s)$ and $\phi_T(s)$, $s \in \mathbb{R}$, respectively, let $F_{S_c}(x)$ and $F_T(x)$, $x \in \mathbb{R}$, be the corresponding cdfs, and assume that $\mathbb{E}[S_c] = \mathbb{E}[T] = 0$. Furthermore, assume that $F_T(x)$ has derivative $p_T(x)$ such that $|p_T(x)| \leq m < \infty$, $\forall x \in \mathbb{R}$. Finally, denote by,

$$\Delta(S_c, T) := \sup_{x \in \mathbb{R}} |F_{S_c}(x) - F_T(x)|,$$

the Kolmogorov distance between the distributions of S_c and T , see e.g. [28]. Then, the Esséen smoothing lemma [29, Lemma 2, p.538] states that, for any $\Theta > 0$,

$$\Delta(S_c, T) \leq \frac{1}{\pi} \int_{-\Theta}^{\Theta} \frac{|\phi_{S_c}(s) - \phi_T(s)|}{|s|} ds + \frac{24m}{\pi\Theta} := I(S_c, T) \quad (4)$$

$$\xrightarrow{\Theta \rightarrow \infty} \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{|\phi_{S_c}(s) - \phi_T(s)|}{|s|} ds := \bar{I}(S_c, T), \quad (5)$$

where (5) is meaningful only if the improper integral converges. Note that the assumption $\mathbb{E}[S_c] = \mathbb{E}[T] = 0$ can be relaxed, if either (4) or (5) are finite.

We first use the smoothing lemma to investigate convergence of the PSR standardized residual to the Gaussian distribution, by deriving an upper bound for the distance,

$$\Delta(Z_{(c, \infty)}, Z) := \sup_{x \in \mathbb{R}} |F_{Z_{(c, \infty)}}(x) - F_Z(x)|, \quad (6)$$

where $F_{Z_{(c, \infty)}}(x)$ and $F_Z(x)$ denote the cdf of $Z_{(c, \infty)}$ and the standard normal cdf, respectively.

We then use the resulting bound to further bound the Kolmogorov distance between the approximated stable law with Gaussian approximation of the PSR residual, \hat{X} , and the ‘exact’ stable law, X ,

$$\Delta(X, \hat{X}) := \sup_{x \in \mathbb{R}} |F_X(x) - F_{\hat{X}}(x)|, \quad (7)$$

where $F_X(x)$ and $F_{\hat{X}}(x)$ are the cdfs of X and \hat{X} , respectively.

From this point on, we restrict attention to the case when the $\{W_j\}$ are normally distributed. We provide results for the symmetric stable law (corresponding to $\mu_W = 0$), for which we are able to obtain closed-form expressions for the CF of the residual, as reported in the following lemma, proved in [26].

Lemma 1: Suppose $W_1 \sim \mathcal{N}(0, \sigma_W^2)$, and denote,

$$a := \frac{\alpha}{2}, \quad \eta := \frac{1-a}{a}, \quad w := \frac{\eta s^2}{2c}, \quad u := w S_{(c, \infty)}^2, \quad (8)$$

for $\alpha \in (0, 2), \alpha \neq 1$. Then,

$$\begin{aligned} \phi_{Z_{(c, \infty)}}(s) &= \psi_{Z_{(c, \infty)}}(w) \\ &= \exp(c(1 - e^{-w} - w^a \gamma(1-a, w))), \end{aligned} \quad (9)$$

where $\gamma(s, x) := \int_0^x t^{s-1} e^{-t} dt$, $\text{Re}(s) > 0$, is the lower incomplete gamma function. Moreover,

$$\begin{aligned} \phi_{X_{(0, c)}}(s) &= \omega_{X_{(0, c)}}(u) \\ &= \exp(-c(1 - e^{-u} + u^a \Gamma(1-a, u))), \end{aligned} \quad (10)$$

where $\Gamma(s, x) := \int_x^\infty t^{s-1} e^{-t} dt$, $\text{Re}(s) > 0$, is the upper incomplete gamma function.

As a consequence of Lemma 1, and through the independence of the random variables $X_{(0, c)}$ and $\hat{R}_{(c, \infty)}$, it follows that, when $\mu_W = 0$ the CF of \hat{X} , the approximated stable distribution is $\phi_{\hat{X}}(s) = \omega_{\hat{X}}(u)$, with:

$$\log \omega_{\hat{X}}(u) = -c(1 - e^{-u} + u^a \Gamma(1-a, u) + u/\eta).$$

IV. NONASYMPTOTIC BOUNDS ON THE CONVERGENCE OF THE PSR RESIDUAL

In this section we apply the smoothing lemma to derive explicit bounds on the distance $\Delta(Z_{(c, \infty)}, Z)$, defined in (6). When $\mu_W = 0$, the closed-form expression in (9) for $\phi_{Z_{(c, \infty)}}(s)$ can be used to further bound above the term $\bar{I}(Z_{(c, \infty)}, Z)$ in (5). This results in the following result, whose proof, together with those of the subsequent theorems, can be found in [26].

Theorem 1: Let $W_j \sim \mathcal{N}(0, \sigma_W^2)$ and let $\Delta(Z_{(c, \infty)}, Z)$ be the Kolmogorov distance between $Z_{(c, \infty)}$ and Z , as in (6). Let a and η be as in (8), and write $g(w) := 1 - e^{-w} - w^a \gamma(1-a, w) < 0$.

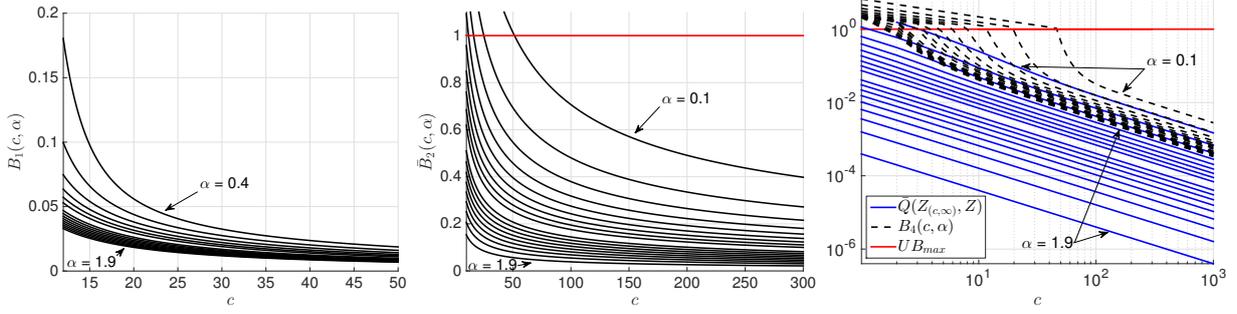


Fig. 2. Bounds on $\Delta(Z_{(c,\infty)}, Z)$. Left: each curve represents the values of the bound $B_1(c, \alpha)$ for $\alpha = 0.4, 0.9, \dots, 1.9$, $\alpha \neq 1$, plotted for $12 \leq c \leq 50$. Centre: each curve represents the values of the bound $\bar{B}_2(c, \alpha)$ for $\alpha = 0.1, 0.2, \dots, 1.9$, $\alpha \neq 1$, plotted for $10 \leq c \leq 300$. Right: comparison between $\bar{Q}(Z_{(c,\infty)}, Z)$ and $B_4(c, \alpha)$, for $\alpha = 0.1, 0.2, \dots, 1.9$, $\alpha \neq 1$, plotted for $3 \leq c \leq 300$. The red horizontal line is simply equal to 1, the maximum possible value of the Kolmogorov distance.

Then, for any $c > 1$, $\Delta(Z_{(c,\infty)}, Z)$ is bounded above by,

$$B_1(c, \alpha) := \frac{1}{\pi} \left(\frac{c}{c-1} \right) \left(\frac{a}{2(2-a)} + \frac{1}{\eta^2} \right) \left[\frac{1}{(c-1)g^2(1)} + \left(\frac{1}{g(1)} - \frac{1}{(c-1)g^2(1)} \right) \exp((c-1)g(1)) + \frac{(c-1) \exp\left\{ (c-1)(1-e^{-1}) \right\}}{a[(c-1)\bar{\gamma}(a)]^{2/a}} \Gamma\left(\frac{2}{a}, (c-1)\bar{\gamma}(a) \right) \right],$$

where $\gamma(s, x)$ and $\Gamma(s, x)$ are lower and upper incomplete gamma functions, respectively, and $\bar{\gamma}(a) := \gamma(1-a, 1)$.

From Theorem 1 it is easy to deduce the following.

Corollary 2: Under the same assumptions and notation of Theorem 1, $\Delta(Z_{(c,\infty)}, Z) = O(1/c)$.

For values of α greater than 0.4, $B_1(c, \alpha)$ gives very good bounds, as shown on the left-hand side of Figure 2. But for α below 0.4, the results deteriorate significantly; for example, for $\alpha = 0.2$, $B_1(c, \alpha)$ is below 1 only for $c > 115$. The following result, obtained by bounding $I(Z_{(c,\infty)}, Z)$, gives an $O(1/\sqrt{c})$ bound which is, of course, asymptotically inferior to that in Theorem 1, but which gives sharper results for small c and $\alpha < 0.4$.

Theorem 3: Under the same assumptions and notation as in Theorem 1, for any $\delta \in (0, 2)$, $\Delta(Z_{(c,\infty)}, Z)$ is bounded above by,

$$B_2(c, \alpha, \delta) := \frac{9.6\sqrt{\eta}}{\pi\sqrt{2(2-\delta)}c} + B_3(c, \alpha, \delta),$$

where $B_3(c, \alpha, \delta)$ is the following $O(1/c)$ term:

$$B_3(c, \alpha, \delta) := \frac{1}{\pi c} \left[\frac{a}{2(2-a)} + \frac{1}{\eta^2} \right] \left(\frac{c(2-\delta)}{(c-1)g(2-\delta)} \right)^2 \times \left\{ 1 - [1 - g(2-\delta)(c-1)] \exp(g(2-\delta)(c-1)) \right\}.$$

Numerically minimizing the bound $B_2(c, \alpha, \delta)$ over δ yields $\bar{B}_2(c, \alpha)$, shown in the central part of Figure 2.

Finally, we combine the results of Theorem 1 and Theorem 3, to obtain useful bounds essentially for all values of $\alpha \in (0, 2)$, $\alpha \neq 1$, and $c > 1$ as,

$$B_4(c, \alpha) := \min \{ B_1(c, \alpha), \bar{B}_2(c, \alpha) \}.$$

Figure 2 shows a comparison between the theoretical bound $B_4(c, \alpha)$ and the numerical estimates $\bar{Q}(Z_{(c,\infty)}, Z)$

of $\bar{I}(Z_{(c,\infty)}, Z)$ reported in [30]. The numerical values are produced through the Matlab routine `quadgk`; we do not show here the numerical error intervals because they are negligibly small for $c \geq 3$. Observe that $B_4(c, \alpha)$ has the same asymptotic rate as $\bar{Q}(Z_{(c,\infty)}, Z)$.

V. CONVERGENCE RATE OF THE APPROXIMATED α -STABLE DISTRIBUTION

Finally, we examine the distance $\Delta(X, \hat{X})$; note that in terms of inference, ultimately, it is $\Delta(X, \hat{X})$ that we wish to make small, by appropriately choosing the value of the truncation parameter c . Using the smoothing lemma (5), and the bound in Theorem 1, we can establish the following.

Theorem 4: Let $\Delta(X, \hat{X})$ be the Kolmogorov distance between X and \hat{X} , as in (7), under the same assumptions and notation as in Theorem 1. Let $N \geq 1$, and introduce arbitrary abscissae $0 =: u_0 < u_1 < \dots < u_N =: 1$ together with the corresponding ordinates $f_0 =: 0$ and $f_i =: \log(\omega_{X_{(0,c)}}(u_i))$, for $i = 1, 2, \dots, N$. Also let, $m_i =: (f_{i+1} - f_i)/(u_{i+1} - u_i)$, and $q_i =: -m_i u_i + f_i$, for $i = 0, 1, \dots, N-1$. Then, for any $c > 1$, $\Delta(X, \hat{X})$ is bounded above by:

$$B_5(c, \alpha, N) := \frac{1}{\pi} c \left(\frac{a}{2(2-a)} + \frac{1}{\eta^2} \right) \times \left\{ \sum_{i=0}^{N-1} \frac{e^{q_i}}{\tilde{m}_i} \left[e^{\tilde{m}_i u_{i+1}} \left(u_{i+1} - \frac{1}{\tilde{m}_i} \right) - e^{\tilde{m}_i u_i} \left(u_i - \frac{1}{\tilde{m}_i} \right) \right] + \frac{e^{\tilde{k}_{(1,\infty)}}}{a(\tilde{l}_{(1,\infty)})^{2/a}} \Gamma\left(\frac{2}{a}, \tilde{l}_{(1,\infty)} \right) \right\},$$

where,

$$\begin{aligned} \tilde{m}_i &:= m_i + (c-1)g(1), \\ k_{(1,\infty)} &:= -c((1 - \exp(-1)) + \Gamma(1-a, 1)) < 0, \\ \tilde{k}_{(1,\infty)} &:= k_{(1,\infty)} - (c-1)(e^{-1} - 1), \\ \tilde{l}_{(1,\infty)} &:= (c-1)\gamma(1-a, 1). \end{aligned}$$

The N abscissae u_i and ordinates f_i serve to define a piecewise linear envelope on $\omega_{X_{(0,c)}}(u)$ for $u \in [0, 1]$, that is used in the proof. Increasing N improves (i.e., decreases) $B_5(c, \alpha, N)$, but the changes are minimal for $N \geq 10$ and logarithmically spaced points.

In Figure 3 we compare the numerical estimates of $\bar{I}(X, \hat{X})$ obtained in [30] (denoted $\bar{Q}(X, \hat{X})$) with the

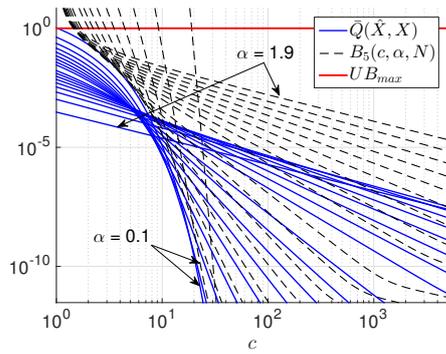


Fig. 3. Bounds on $\Delta(X, \hat{X})$. The blue solid lines represent $\hat{Q}(X, \hat{X})$; the black dashed lines represent $B_5(c, \alpha, N)$ with $N = 10$ and logarithmically spaced points. The values are plotted for $\alpha = 0.1, 0.2, \dots, 1.9$, $\alpha \neq 1$, and for $1 \leq c \leq 5000$. The red horizontal line is simply equal to 1, the maximum possible value of the Kolmogorov distance.

bound $B_5(c, \alpha, N)$ for $N = 10$. Note that the bound of Theorem 4 correctly captures the dependence on α . The approximation error is lower for smaller values of α , a reversal of the trend observed in Figure 2. We believe this is because, as α decreases, the relative significance of the residual term is much smaller, when compared with the more heavy-tailed initial terms in the PSR. We also observe that the rate of convergence is dramatically better for smaller α , again in contrast with the analysis of the residual approximation in Figure 2.

VI. CONCLUSIONS

In this paper we have provided explicit, nonasymptotic bounds for different approximations of the PSR of symmetric α -stable random variables. Our theoretical results are in agreement with our previous numerically computed convergence rates. These results form a collection of tools that can be used to select appropriate values for the PSR truncation parameter, in order to control the quality of the resulting approximation. We expect that the present bounds, as well as our current work extending them to the case of continuous-time stochastic systems, will contribute to providing a foundation for quantifying the accuracy of results obtained by inference algorithms employing the PSR.

REFERENCES

- [1] Brockwell, P. J. and Davis, R. A., *Introduction to Time Series and Forecasting*, ser. Introduction to Time Series and Forecasting. Springer, 2002, no. v. 1.
- [2] Tankov, P. and Cont, R., *Financial Modelling with Jump Processes, Second Edition*, ser. Chapman and Hall/CRC Financial Mathematics Series. Taylor & Francis, 2015.
- [3] Øksendal, B., *Stochastic Differential Equations: An Introduction with Applications*, ser. Hochschultext / Universitext. Springer, 2003.
- [4] Harvey, A. C., *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, 1990.
- [5] Brockwell, P. J., “Lévy-driven CARMA processes,” *Annals of the Institute of Statistical Mathematics*, vol. 53, no. 1, pp. 113–124, 2001.
- [6] —, “Representations of continuous-time ARMA processes,” *Journal of Applied Probability*, vol. 41A, pp. 375–382, 02 2004.
- [7] Brockwell, P. J. and Lindner, A., “Existence and uniqueness of stationary Lévy-driven CARMA processes,” *Stochastic Processes and their Applications*, vol. 119, no. 8, pp. 2660–2681, August 2009.

- [8] Applebaum, D., *Lévy Processes and Stochastic Calculus*, ser. Cambridge studies in advanced mathematics. Cambridge University Press, 2004.
- [9] Asmussen, S. and Rosinski, J., “Approximations of small jumps of Lévy processes with a view towards simulation,” *Journal of Applied Probability*, vol. 38, no. 2, pp. 482–493, 2001.
- [10] Rosinski, J., “Series representations of Lévy processes from the perspective of point processes,” in *Lévy Processes*, O. Barndorff-Nielsen, S. Resnick, and T. Mikosch, Eds. Birkhauser Boston, 2001, pp. 401–415.
- [11] —, *Simulation of Lévy Processes*. John Wiley & Sons, Ltd, 2014.
- [12] Gnedenko, B. V. and Kolmogorov, A. N., *Limit Distributions for Sums of Independent Random Variables*, ser. Addison-Wesley series in statistics. Addison-Wesley, 1968.
- [13] Samoradnitsky, G. and Taqqu, S., *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*, ser. Stochastic Modeling Series. Taylor & Francis, 1994.
- [14] Janicki A. and Weron A., *Simulation and Chaotic Behavior of α -stable Stochastic Processes*, ser. HSC Books. Hugo Steinhaus Center, Wroclaw University of Technology, 1994.
- [15] Lemke, T. and Godsill, S.J., “Enhanced Poisson sum representation for alpha-stable processes,” in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 4100–4103.
- [16] Lemke, T., “Poisson Series Approaches to Bayesian Monte Carlo Inference for Skewed α -Stable Distributions and Stochastic Processes,” Ph.D. dissertation, Kaiserslautern University of Technology, 2014.
- [17] Lemke, T., and Godsill, S. J., “Inference for models with asymmetric α -stable noise processes,” in *Unobserved Components and Time Series Econometrics*, Koopman S. J. and Shephard N., Ed. Oxford: Oxford University Press, 2015, ch. 9.
- [18] Riabiz, M. and Godsill, S.J., “Approximate simulation of linear continuous time models driven by asymmetric stable Lévy processes,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, March 2017.
- [19] Buckle, D. J., “Bayesian inference for Stable distributions,” *Journal of the American Statistical Association*, vol. 90, no. 430, pp. 605–613, 1995.
- [20] Lombardi, M. J., “Bayesian inference for α -stable distributions: A random walk MCMC approach,” *Computational Statistics & Data Analysis*, vol. 51, no. 5, pp. 2688 – 2700, 2007.
- [21] Riabiz, M. and Lindsten F., and Godsill S., “Pseudo-Marginal MCMC for Parameter Estimation in α -Stable Distributions,” *IFAC-PapersOnLine*, vol. 48, no. 28, pp. 472 – 477, 2015, 17th IFAC Symposium on System Identification SYSID 2015.
- [22] Doucet, A. and Godsill, S. J. and Andrieu, C., “On sequential Monte Carlo sampling methods for Bayesian filtering,” *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [23] Schön, T. and Gustafsson, F. and Nordlund, P. J., “Marginalized particle filters for mixed linear/nonlinear state-space models,” *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2279–2289, July 2005.
- [24] Lemke, T. and Godsill, S. J., “A Poisson series approach to Bayesian Monte Carlo inference for skewed alpha-stable distributions,” in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 8023–8027.
- [25] Lemke, T. and Riabiz, M. and Godsill, S. J., “Fully Bayesian inference for α -stable distributions using a Poisson series representation,” *Digital Signal Processing*, vol. 47, pp. 96 – 115, 2015, special Issue in Honour of William J. (Bill) Fitzgerald.
- [26] Riabiz, M. and Ardeshiri, T., Kontoyiannis, I., and Godsill, S., *arXiv preprint*, 2018.
- [27] Riabiz, M. and Ardeshiri, T. and Godsill, S., “A central limit theorem with application to inference in α -stable regression models,” in *Proceedings of the Time Series Workshop at NIPS 2016*, ser. Proceedings of Machine Learning Research, vol. 55. Barcelona, Spain: PMLR, 09 Dec 2017, pp. 70–82.
- [28] Rachev, S. T. and Klebanov, L. B. and Stoyanov, S. V. and Fabozzi, F. J., *Probability Distances and Probability Metrics: Definitions*. New York, NY: Springer New York, 2013, pp. 11–31.
- [29] Feller, W., *An introduction to probability theory and its applications*, ser. Wiley mathematical statistics series. Wiley, 1966, no. v. 2.
- [30] M. Riabiz, T. Ardeshiri, I. Kontoyiannis, and S. Godsill, “Simulated convergence rates with application to an intractable α -stable inference problem,” in *2017 IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (IEEE CAMSAP 2017)*, Curacao, Netherlands Antilles, Dec. 2017.