

Accurate and efficient numerical calculation of stable densities via optimized quadrature and asymptotics

Sebastian Ament* and Michael O’Neil†

August 31, 2021

Abstract

Stable distributions are an important class of infinitely-divisible probability distributions, of which two special cases are the Cauchy distribution and the normal distribution. Aside from a few special cases, the density function for stable distributions has no known analytic form, and is expressible only through the variate’s characteristic function or other integral forms. In this paper we present numerical schemes for evaluating the density function for stable distributions, its gradient, and distribution function in various parameter regimes of interest, some of which had no pre-existing efficient method for their computation. The novel evaluation schemes consist of optimized generalized Gaussian quadrature rules for integral representations of the density function, complemented by asymptotic expansions near various values of the shape and argument parameters. We report several numerical examples illustrating the efficiency of our methods. The resulting code has been made available online.

Keywords: Stable distributions, α -stable, generalized Gaussian quadrature, infinitely-divisible distributions, numerical quadrature.

1 Introduction

Continuous random variables that follow *stable laws* arise frequently in physics [6], finance and economics [9, 13], electrical engineering [10], and many other fields of the natural and social sciences. Certain sub-classes of these distributions are also referred to as α -stable, a -stable, stable Paretian distributions, or Lévy alpha-stable distributions. Going forward, we will merely refer to them as stable distributions. The defining characteristic of random variables that follow stable laws is that the sum of two independent copies follows the *same* scaled and translated distribution [14]. For example, if X_1 and X_2 are independent, identically distributed (iid) stable random variables, then *in distribution*

$$X_1 + X_2 \sim aX + b, \tag{1.1}$$

where X has the same distribution as each X_ℓ . Several discrete random variables, such as those following Poisson distributions, also obey this *stability-of-sums* law, but we restrict our attention to continuous distributions. Modeling with stable distribution has several advantages. For example, even though in general they do not have finite variances, they are closed under sums and satisfy a

*Department of Computer Science, Cornell University, New York, NY. E-mail: ament@cs.cornell.edu.

†Department of Mathematics, Courant Institute and Tandon School of Engineering, New York University, New York, NY. E-mail: oneil@cims.nyu.edu. Research partially supported by AIG-NYU Award #A15-0098-001.

generalized-type of Central Limit Theorem [10, 22]. This property is directly related to the fact that they have tails which are heavier than those of normal random variables. For this reason, these distributions are useful in describing many real-world data sets from finance, physics, and chemistry. On the other hand, computing with stable distributions requires more sophistication than does, for example, computing with normal distributions. When modeling with multivariate normal distributions, all of the relevant calculations (in, for example, likelihood evaluation) are linear-algebraic in nature: matrix inversion, determinant calculation, eigenvalue computation, etc. [18]. The analogous operations for stable distributions are highly nonlinear, often requiring technical multivariable optimization and Monte Carlo codes, slowing down the resulting calculation many-fold. In this work, for these reasons, we restrict our attention to one-dimensional stable distributions. Numerical schemes for multivariate stable distributions are an area of current research.

To be more precise, if we denote by $\alpha \in (0, 2]$ the *stability parameter*, $\beta \in [-1, 1]$ the *skewness parameter*, $\gamma \in \mathbb{R}$ the *location parameter*, and $\lambda \in \mathbb{R}^+$ the *scale parameter* of X , then these random variables satisfy the relationship [22]:

$$a_1 X_1 + a_2 X_2 \sim a X + \begin{cases} \lambda \gamma (a_1 + a_2 - a) & \alpha \neq 1 \\ \lambda \beta (2/\pi) (a_1 \log(a_1/a) + a_2 \log(a_2/a)) & \alpha = 1 \end{cases} \quad (1.2)$$

where, as before \sim is used to denote *equality in distribution* and $a = (a_1^\alpha + a_2^\alpha)^{1/\alpha}$. Enforcing the previous stability laws, although partially redundant, places conditions on the characteristic function of X (i.e. the Fourier transform of the probability density function). Since the density of the sum of two iid random variables is obtained via convolution of their individual densities, in the Fourier domain this is equivalent to multiplication of the characteristic functions. It can be shown that, in general, the class of characteristic functions for stable distributions must be of the form:

$$\begin{aligned} \mathbb{E} \left[e^{itX} \right] &= \varphi_X(t) \\ &= e^{\lambda(it\gamma - |t|^\alpha + it\omega(t, \alpha, \beta))}, \end{aligned} \quad (1.3)$$

with

$$\omega(t, \alpha, \beta) = \begin{cases} |t|^{\alpha-1} \beta \tan \frac{\pi\alpha}{2} & \text{if } \alpha \neq 1, \\ -\frac{2\beta}{\pi} \log |t| & \text{if } \alpha = 1, \end{cases} \quad (1.4)$$

and as before,

$$\alpha \in (0, 2], \quad \beta \in [-1, 1], \quad \gamma \in (-\infty, \infty), \quad \lambda > 0. \quad (1.5)$$

This particular parameterization of the characteristic function φ_X in terms of α , β , γ , and λ is the canonical one [22], and often referred to as the **A**-parameterization. As discussed in Section 2, we will deal solely with an alternative parameterization, the **M**-parameterization. This parameterization is obtained by merely a change of variables in the x -parameter but, in contrast to (1.3), is jointly continuous in all of its parameters.

Often, the form of the above characteristic function is taken to be the *definition* of stable distributions because of the absence of an analytic form of the inverse transform. Special cases of these

distributions are normal random variables ($\alpha = 2$ and $\beta = 0$), Cauchy distributions ($\alpha = 1$ and $\beta = 0$), and the Lévy distribution ($\alpha = 0.5$ and $\beta = 1$). Each of these distributions has a closed-form expression for its density and cumulative distribution function. However, as mentioned before, in general, the density and distributions functions for stable random variables have no known analytic form, and are expressible only via their Fourier transform or special-case asymptotic series. Because of this, performing inference or developing models based on these distribution laws can be computationally intractable if the density and distribution functions are expensive to compute (i.e. if numerically evaluating the corresponding Fourier integral is expensive). We will focus our attention on the numerical evaluation of the density function for a unit, centered distribution: $\gamma = 0$ and $\lambda = 1$. We will denote this class of unit, centered stable distributions in the \mathbf{A} -parameterization as $\mathcal{S}(\alpha, \beta, \mathbf{A})$, and say that $X \sim \mathcal{S}(\alpha, \beta, \mathbf{A})$ if X has characteristic function (1.3) with $\gamma = 0$ and $\lambda = 1$. In a slight change of notation from [14], we make the particular parameterization explicit in the definition of $\mathcal{S}(\alpha, \beta, \cdot)$.

Most existing numerical methods for the evaluation of the corresponding density function f ,

$$f(x; \alpha, \beta) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \varphi_X(t) e^{-itx} dt, \quad (1.6)$$

rely on some form of numerical integration [11] or, in the symmetric case ($\beta = 0$), asymptotic expansions for extremal values of x , and α [8]. Often, if the shape parameters of the distribution are being inferred or estimated from data, as in the case of maximum likelihood calculations or Bayesian modelling, the density function f must be evaluate at the same x values for *many* values of the parameters α and β . In order to ensure the accuracy of the numerical integration scheme, often adaptive quadrature is used. However, for new values of the parameters, no previous information can be used, nor are these quadratures optimal (as in the sense that the trapezoidal rule is optimal for smooth periodic function or that Gauss-Legendre quadrature is optimal for polynomials on finite intervals [4]).

The main contribution of this work is to develop an efficient means by which to evaluate the density function for stable random variables using various integral formulations, optimized quadrature schemes, and asymptotic expansions. We develop generalized Gaussian quadrature rules [1, 21] which are able to evaluate the density f for various *ranges* of the shape parameters α, β , as well as the argument x . We generate quadratures that consist of a *single* set of weights and nodes which are able to integrate the characteristic function (or deformations thereof) for large regions of the parameters and argument space. Using a small collection of such quadrature rules we are able to cover most of $\alpha\beta x$ -space. This class of quadrature schemes is an extension of the classical Gaussian quadrature schemes for polynomials which are able to *exactly* integrate polynomials of degree $d \leq 2n - 1$ using n nodes and n weights (i.e. using a total of $2n$ degrees of freedom). We discuss these quadrature rules in detail in Section 3, as well as derive new asymptotic expansions in the asymmetric case in the \mathbf{M} -parameterization. For regions with large x , we derive efficient asymptotic expansions which can be used for evaluation.

The paper is organized as follows: in Section 2 we review some standard integral representations and asymptotic expansions for the density functions of stable distributions. In Section 3 we discuss

the procedure for constructing generalized Gaussian quadratures for evaluating the integral representations presented in the previous section. In Section 4, we demonstrate the effectiveness of our schemes for evaluating the density functions via various numerical examples. In Section 5, the conclusion, we discuss some additional areas of research, and point out regimes in which the algorithms of this paper are not applicable.

2 Stable distributions

In this section we review some basic facts regarding stable distributions, and review the integral and asymptotic expansions that we will use to evaluate the density function. As mentioned in the previous section, there are several different parameterizations of stable densities. We now detail the parameterization useful for numerical calculations, most commonly referred to as Zolotarev's **M**-parameterization. Random variables that follow stable distributions with this parameterization will be denoted $X \sim \mathcal{S}(\alpha, \beta, \mathbf{M})$.

2.1 Basic Facts

There are a number of different parameterizations for stable distributions, each of which is useful in a particular regime: integral representations, asymptotic expansions, etc. It was shown in [11] that Zolotarev's **M**-parameterization is particularly useful for numerical computations as it allows for the computation of a *unit* density that can be later scaled and translated. In our numerical scheme, we also use the **M**-parameterization because of its continuity in all underlying parameters. This is a standard procedure among other numerical methods as well.

This parameterization, and therefore the density function, is defined by the characteristic function

$$\begin{aligned} \varphi_X(t) &= e^{\lambda(it\gamma - |t|^\alpha + it\omega_M(t, \alpha, \beta))}, \\ \omega_M(t, \alpha, \beta) &= \begin{cases} (|t|^{\alpha-1} - 1)\beta \tan \frac{\pi\alpha}{2} & \text{if } \alpha \neq 1, \\ -\frac{2\beta}{\pi} \log |t| & \text{if } \alpha = 1. \end{cases} \end{aligned} \quad (2.1)$$

For the rest of the paper, we will work with unit stable laws ($\gamma = 0$, $\lambda = 1$) unless otherwise mentioned. We will refer to the case of $\beta = 0$ as the *symmetric* case, and otherwise for $\beta \neq 0$ the *asymmetric* case. In general, the parameters α and β cannot be interchanged between different parameterizations, and in this particular case the change of variables from the **A**- to the **M**-parameterization is given by:

$$\alpha_A = \alpha_M = \alpha, \quad \beta_A = \beta_M = \beta, \quad \gamma_A = \gamma_M - \beta \tan \frac{\pi\alpha}{2}, \quad \lambda_A = \lambda_M \quad (2.2)$$

where the subscripts are used to denote the parameterization. Under this change of variable, the characteristic function in the **A**-parameterization lacks the term $-it\beta \tan(\pi\alpha/2)$ in the exponent. The existence of this term in the **M**-parameterization makes characteristic function continuous at $\alpha = 1$. This change of variables is mostly done for analytical convenience. However, the mode of the density in the **A**-parameterization approaches infinity as $\alpha \rightarrow 1$ if $\beta \neq 0$. Therefore, neither of the one-sided limits $\alpha \rightarrow 1^\pm$ equals a distribution in the **A**-parameterization [22].

From (2.1), it follows that for unit stable laws,

$$\varphi_X(-t; \alpha, \beta) = \varphi_X(t; \alpha, -\beta) = \overline{\varphi_X(t; \alpha, \beta)}, \quad (2.3)$$

where \bar{z} denotes the complex-conjugate of z , and we have used the Fourier transform convention

$$\begin{aligned} \varphi_X(t) &= \mathbb{E} \left[e^{itX} \right] \\ &= \int_{-\infty}^{\infty} f(x) e^{itx} dx. \end{aligned} \quad (2.4)$$

The density f is therefore given by:

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \varphi_X(t) e^{-itx} dt. \quad (2.5)$$

In conjunction with (2.3), one can show that

$$\begin{aligned} f(x; \alpha, \beta) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \varphi_X(t; \alpha, \beta) dt \\ &= \frac{1}{2\pi} \left(\int_0^{\infty} e^{itx} \varphi_X(t; \alpha, -\beta) dt + \int_{-\infty}^0 e^{-itx} \varphi_X(t; \alpha, \beta) dt \right) \\ &= \frac{1}{\pi} \operatorname{Re} \int_0^{\infty} e^{itx} \varphi_X(t; \alpha, -\beta) dt. \end{aligned} \quad (2.6)$$

Furthermore, (2.3) and the Fourier inversion formula imply that

$$f(-x; \alpha, \beta) = f(x; \alpha, -\beta). \quad (2.7)$$

This symmetry allows for considerable restriction in the relevant values of x . In particular, for $\alpha \neq 1$, defining $\zeta(\alpha, \beta) = \beta \tan \pi\alpha/2$, we need only address the case $x > \zeta$. Indeed, if $x < \zeta$, then

$$-x > -\zeta(\alpha, \beta) = \zeta(\alpha, -\beta) \quad (2.8)$$

as can be seen from (2.1).

2.2 Integral representations

Inserting φ_X from (2.1) into the inverse Fourier transform (2.6), we see that

$$f(x; \alpha, \beta) = \frac{1}{\pi} \int_0^{\infty} \cos(h(t; x, \alpha, \beta)) e^{-t^\alpha} dt, \quad (2.9)$$

where for $\alpha \neq 1$

$$\begin{aligned} h(t; x, \alpha, \beta) &= (x - \zeta)t + \zeta t^\alpha, \\ \zeta(\alpha, \beta) &= -\beta \tan \frac{\pi\alpha}{2}, \end{aligned} \quad (2.10)$$

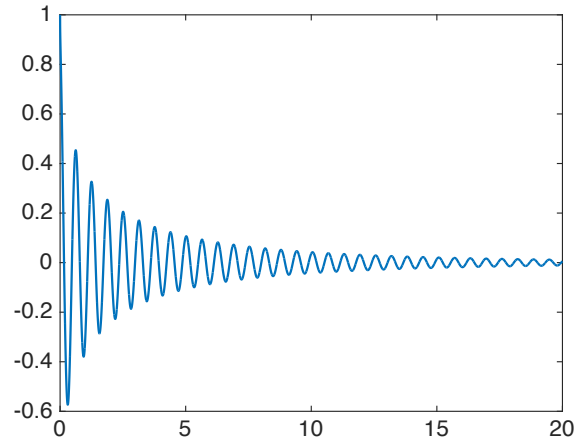


Figure 1: Graph of the Fourier transform of f , i.e. the integrand in (2.9), for parameter values of $x = 10$, $\alpha = .5$ and $\beta = 0$.

and for $\alpha = 1$

$$h(t; x, \alpha, \beta) = xt + \frac{2\beta t}{\pi} \log t, \quad (2.11)$$

$$\zeta(\alpha, \beta) = 0.$$

Because of the linear dependence of h on t , the integrand in (2.9) is oscillatory for modestly sized values of x . See Figure 1 for a plot of this integrand. For numerical calculations, the infinite interval of integration can be truncated based on the decay of the exponential term if x is not too large. However, for small values of α , this region of integration can still be prohibitively large. Furthermore, for large $|x|$, the integrand becomes increasingly oscillatory and standard integration schemes (e.g. trapezoidal rule, Gaussian quadrature) not only become expensive, but lose accuracy due to the oscillation. It is possible that Filon-type quadratures [15] could be applicable, but this has yet to be thoroughly investigated in the literature. Section 3.4 contains a brief discussion of quadrature techniques for highly oscillatory integrands, but these methods, unfortunately, would likely prove to be more computationally expensive than those techniques presented in this work. For these reasons, this representation of the density f in (2.9) cannot be used efficiently in the following parameter ranges: small α and/or large $|x|$.

Alternatively, the integral in (2.9) can be rewritten using the method of stationary phase. This calculation was done in [11]. To this end, we begin by rewriting (2.9) as

$$f(x; \alpha, \beta) = \frac{1}{\pi} \operatorname{Re} \int_0^\infty e^{ih(z) - z^\alpha} dz, \quad (2.12)$$

where $\operatorname{Re} w$ denotes the real part of the complex number w and we have suppressed the explicit dependence on α and β for simplicity. Deforming along the contour with zero phase, we have

$$f(x; \alpha, \beta) = \frac{\alpha}{\pi|\alpha - 1|} \frac{1}{(x - \zeta)} \int_{-\theta_0}^{\frac{\pi}{2}} g(\theta; x, \alpha, \beta) e^{-g(\theta; x, \alpha, \beta)} d\theta, \quad (2.13)$$

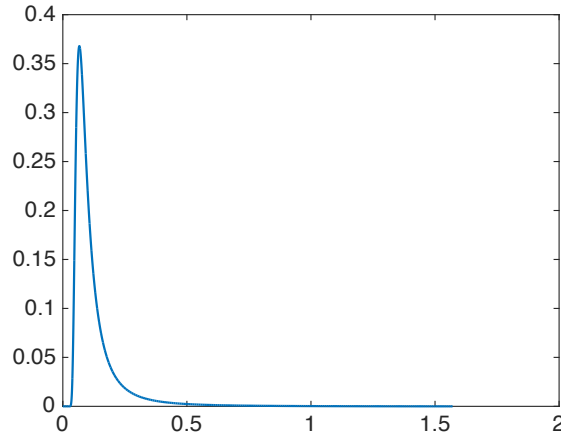


Figure 2: Graph of the stationary phase integrand in (2.13) for parameter values of $x = 10^{-1}$, $\alpha = 1.5$ and $\beta = 0$. Note the large derivatives and spiked behavior.

with

$$g(\theta; x, \alpha, \beta) = (x - \zeta)^{\frac{\alpha}{\alpha-1}} V(\theta; \alpha, \beta), \quad (2.14)$$

where for $\alpha \neq 1$

$$\begin{aligned} \zeta(\alpha, \beta) &= -\beta \tan \frac{\pi\alpha}{2}, \\ \theta_0(\alpha, \beta) &= \frac{1}{\alpha} \arctan \left(\beta \tan \frac{\pi\alpha}{2} \right), \\ V(\theta; \alpha, \beta) &= (\cos \alpha\theta_0)^{\frac{1}{\alpha-1}} \left(\frac{\cos \theta}{\sin \alpha(\theta_0 + \theta)} \right)^{\frac{\alpha}{\alpha-1}} \frac{\cos(\alpha\theta_0 + (\alpha-1)\theta)}{\cos \theta}, \end{aligned} \quad (2.15)$$

and where for $\alpha = 1$

$$\begin{aligned} \zeta(\alpha, \beta) &= 0, \\ \theta_0(\alpha, \beta) &= \frac{\pi}{2}, \\ V(\theta; \alpha, \beta) &= \frac{2}{\pi} \left(\frac{\frac{\pi}{2} + \beta\theta}{\cos \theta} \right) \exp \left(\frac{1}{\beta} \left(\frac{\pi}{2} + \beta\theta \right) \tan \theta \right). \end{aligned} \quad (2.16)$$

While seemingly more complicated than that in (2.9), the integrand in (2.13) is strictly positive, has no oscillations and the interval of integration is finite. Unfortunately, this is not a fail-safe transformation. In particular, for very small and very large x , and α close to 1 and 2, the integrand has large derivatives (e.g. very spiked) and is hard to efficiently integrate, see Figure 2. For this reason, previous schemes [8, 11] have used zero-finding methods to locate the integrand's unique extremum point θ_{max} , where $g(\theta_{max}) = 1$. Subsequently, adaptive quadrature schemes were applied to the two subintervals created by splitting the original interval of integration at θ_{max} . This procedure is often computationally expensive.

Even though an expression for the cumulative distribution function (CDF) can be derived by straightforward integration of (2.13), this is not ideal for various numerical considerations described

later. We obtain an expression for F , the CDF for f , by using an inversion theorem found in [19]:

$$\begin{aligned} F(x) &= \frac{1}{2} - \frac{1}{2\pi} \int_0^\infty \left(\varphi(t)e^{-ixt} - \varphi(-t)e^{ixt} \right) \frac{dt}{it} \\ &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \sin(h(t; x, \alpha, \beta)) e^{-t^\alpha} \frac{dt}{t}, \end{aligned} \quad (2.17)$$

The theorem in [19] assumes the existence of a mean of the random variable which is associated with the characteristic function. However, stable variates with $\alpha < 1$ do not have a mean. Fortunately, one can relax the assumptions of the theorem to integrability of the integrand in (2.17). This means that the expression is valid for all parameter combinations. The integrand of (2.17) has similar behavior as the integrand in (2.9), and therefore we expect that similar numerical schemes for evaluating the integral will be applicable. In the asymmetric case however, the integrand has a singularity at the origin when $\alpha < 1$ and $\beta \neq 1$. The advantages of this representation are explained in detail in Section 4.

2.3 Series and asymptotics

Fortunately, there are series and asymptotic expansions for Zolotarev's \mathbf{M} -parameterization which nicely complement the integral representations above. Specifically, they yield accurate results for very small and large x . Zolotarev derived series and asymptotic expansions for the \mathbf{B} -parameterization in [22], and in the following, we will derive similar representations for the \mathbf{M} -parameterization valid in the general case¹.

Lemma 1. Let $\alpha \neq 1$, $\beta \in [-1, 1]$, $\zeta = -\beta \tan(\pi\alpha/2)$, and

$$S_n^0(x; \alpha, \beta) := \frac{1}{\alpha\pi} \sum_{k=0}^n \frac{\Gamma(\frac{k+1}{\alpha})}{\Gamma(k+1)} (1 + \zeta^2)^{-\frac{k+1}{2\alpha}} \sin([\pi/2 + (\arctan \zeta)/\alpha] (k+1)) (x - \zeta)^k. \quad (2.18)$$

Then for any $n \in \mathbb{N}$,

$$|f(x; \alpha, \beta) - S_{n-1}^0(x; \alpha, \beta)| \leq \frac{1}{\alpha\pi} \frac{\Gamma(\frac{n+1}{\alpha})}{\Gamma(n+1)} (1 + \zeta^2)^{-\frac{n+1}{2\alpha}} |x - \zeta|^n. \quad (2.19)$$

Proof. To obtain a series representation centered at $x = \zeta$, we follow the derivation in [22], but for the \mathbf{M} -parameterization instead of the \mathbf{B} -parameterization:

$$\begin{aligned} f(x; \alpha, \beta) &= \frac{1}{\pi} \operatorname{Re} \int_0^\infty e^{itx} \varphi_X(t; \alpha, -\beta) dt \\ &= \frac{1}{\pi} \operatorname{Re} \int_0^\infty \sum_{k=0}^\infty \frac{(it[x - \zeta])^k}{k!} \exp(-(1 - i\zeta)t^\alpha) dt \\ &= \frac{1}{\pi} \sum_{k=0}^{n-1} \frac{(x - \zeta)^k}{k!} \operatorname{Re} \int_0^\infty (it)^k \exp(-(1 - i\zeta)t^\alpha) dt + R_n, \end{aligned} \quad (2.20)$$

¹After the publication of this article, we were made aware of related prior work on asymptotic expansions in [20].

where

$$R_n = \frac{1}{\pi} \operatorname{Re} \int_0^\infty \left[e^{it(x-\zeta)} - \sum_{k=0}^{n-1} \frac{(it[x-\zeta])^k}{k!} \right] \exp(-(1-i\zeta)t^\alpha) dt \quad (2.21)$$

Applying the change of variables $s = (1-i\zeta)^{1/\alpha}t$ to the last integral in (2.20) and subsequently rotating the contour of integration to the real axis yields

$$f(x; \alpha, \beta) = S_{n-1}(x; \alpha, \beta) + R_n^0. \quad (2.22)$$

The change of contour can be justified with Lemma 2.2.2 in [22]. It remains to show that R_n is bounded in magnitude by (2.19). Indeed,

$$\begin{aligned} |R_n| &= \frac{1}{\pi} \left| \operatorname{Re} \int_0^\infty \left(e^{it(x-\zeta)} - \sum_{k=0}^{n-1} \frac{[it(x-\zeta)]^k}{k!} \right) e^{-(1-i\zeta)t^\alpha} dt \right| \\ &= \frac{1}{\pi} \left| \operatorname{Re} \int_0^\infty \left(\sum_{k=n}^\infty \frac{(i[1-i\zeta]^{-1/\alpha}s[x-\zeta])^k}{k!} \right) e^{-s^\alpha} \frac{ds}{(1-i\zeta)^{1/\alpha}} \right| \\ &\leq \frac{|x-\zeta|^n}{\pi n!} (1+\zeta^2)^{-\frac{n+1}{2\alpha}} \int_0^\infty s^n e^{-s^\alpha} ds \\ &= \frac{1}{\alpha\pi} \frac{\Gamma(\frac{n+1}{\alpha})}{\Gamma(n+1)} (1+\zeta^2)^{-\frac{n+1}{2\alpha}} |x-\zeta|^n. \end{aligned} \quad (2.23)$$

The second equality comes from the change of variables $s = (1-i\zeta)^{1/\alpha}t$, and a rotation of the contour of integration to the real axis. The difference between the exponential and the sum of the first $n-1$ terms of its power series can be bounded by the n^{th} term with the mean value theorem. \square

As a consequence, the series (2.18) converges to the density as $n \rightarrow \infty$ for $\alpha > 1$. For $\alpha < 1$, (2.18) is not convergent, but can be used as an asymptotic expansion as $x \rightarrow \zeta$ if $\beta \neq 1$. While the truncation error bound (2.19) holds regardless of the parameters, (2.18) does not capture the asymptotic behavior of the density as $x \rightarrow \zeta^+$ if $\alpha < 1$ and $\beta = 1$. Indeed, (2.18) is identically zero for this parameter choice. On the other hand, the density falls off exponentially as $x \rightarrow \zeta^+$, which an asymptotic expansion in [22] reveals. Unfortunately, this expansion cannot be efficiently evaluated numerically because its coefficients do not have a closed form.

Still, by rearranging (2.19), we find that truncating (2.18) after the $n-1$ term is accurate to within ϵ of the true value for all x satisfying

$$\begin{aligned} |x-\zeta| &\leq \left[\epsilon \alpha \pi (1+\zeta^2)^{\frac{n+1}{2\alpha}} \frac{\Gamma(n+1)}{\Gamma(\frac{n+1}{\alpha})} \right]^{1/n} \\ &:= B_n^0(\alpha, \beta). \end{aligned} \quad (2.24)$$

After the discussion in the paragraph above, it is important to stress that this bound guarantees *absolute accuracy* of the truncated series, rather than *relative accuracy*.

Lemma 2. Let $\alpha \neq 1$, $\beta \in [-1, 1]$, $\zeta = -\beta \tan(\pi\alpha/2)$, and

$$S_n^\infty(x; \alpha, \beta) := \frac{\alpha}{\pi} \sum_{k=1}^n (-1)^{k+1} \frac{\Gamma(\alpha k)}{\Gamma(k)} (1 + \zeta^2)^{k/2} \sin([\pi\alpha/2 - \arctan \zeta]k) (x - \zeta)^{-\alpha k - 1}. \quad (2.25)$$

Then for any $n \in \mathbb{N}$,

$$|f - S_{n-1}^\infty| \leq \frac{\alpha}{\pi} \frac{\Gamma(\alpha n)}{\Gamma(n)} (1 + \zeta^2)^{\frac{n}{2}} |x - \zeta|^{-\alpha n - 1}. \quad (2.26)$$

Proof. For simplicity, we first derive a series expansion for the **A**-parameterization and convert it to the **M**-parameterization via the shift $x_A = x_M - \zeta$ afterwards. To do this, we extend φ_X to the complex plane, and integrate along the contour $z = iux^{-1/\alpha}$. Again, this is justified by Lemma 2.2.2 in [22].

$$\begin{aligned} x^{-1/\alpha} f_A(x^{-1/\alpha}; \alpha, \beta) &= \frac{1}{\pi x^{1/\alpha}} \operatorname{Re} \int_0^\infty e^{izx^{-1/\alpha}} \varphi_X(z; \alpha, -\beta) dz \\ &= -\frac{1}{\pi} \operatorname{Im} \int_0^\infty e^{-u} \varphi(iux^{1/\alpha}; \alpha, -\beta) du \\ &= \frac{\alpha}{\pi} \sum_{k=1}^{n-1} \frac{\Gamma(\alpha n)}{\Gamma(n)} (-1)^{k+1} (1 + \zeta^2)^{k/2} \sin([\pi\alpha/2 - \arctan \zeta]k) x^k + R_n^\infty. \end{aligned} \quad (2.27)$$

Here, the first n terms of the power series of the characteristic function were used to approximate the integral. Therefore,

$$R_n^\infty := -\frac{1}{\pi} \operatorname{Im} \int_0^\infty \left[\exp(-(1 - i\zeta)xu^\alpha) - \sum_{k=0}^{n-1} \frac{(-[1 - i\zeta]xu^\alpha)^k}{k!} \right] e^{-u} du. \quad (2.28)$$

The bound of $|R_n^\infty|$ is attained in a similar manner as the one for R_n^0 in Lemma 1. Rearranging the last line of (2.27), and substituting $x - \zeta$ for x yields the series (2.25) in the **M**-parameterization. \square

Expression (2.25) converges to the density for $\alpha < 1$, and can be used as an asymptotic expansion for $\alpha > 1$, $\beta \neq -1$. In the case $\alpha > 1$, $\beta = -1$, (2.25) is identically zero, while the density decreases to zero exponentially as $x \rightarrow \infty$ [22]. As with the series above however, we can still guarantee absolute accuracy compared to the true density. Indeed, as a consequence of the lemma, the series (2.25) is accurate to precision ϵ for any x satisfying

$$\begin{aligned} |x - \zeta| &\geq \left[\frac{\alpha}{\pi\epsilon} (1 + \zeta^2)^{\frac{n}{2}} \frac{\Gamma(\alpha n)}{\Gamma(n)} \right]^{1/(\alpha n - 1)} \\ &:= B_{n-1}^\infty(\alpha, \beta). \end{aligned} \quad (2.29)$$

Notably, if for some α, β , we take $n_0(\alpha, \beta)$ terms of (2.18) and $n_\infty(\alpha, \beta)$ terms of (2.25), then it only remains to show that values of x in the range

$$B_{n_0}^0 \leq x - \zeta \leq B_{n_\infty}^\infty \quad (2.30)$$

can be evaluated efficiently. We will elaborate on the details of our scheme in Section 4.

2.4 Derivatives of stable densities

The integral representation (2.9) admits reasonably simple expressions for the Fourier transform of the derivatives of the density with respect to x , α , and β . We will derive these expressions here. For brevity, let $h = h(t, x; \alpha, \beta)$, $\zeta = \zeta(\alpha, \beta)$, and $\partial_x = \partial/\partial x$. Similarly for differentiation with respect to α and β . First, note that for $\alpha \neq 1$,

$$\begin{aligned}\partial_\alpha \zeta &= -\frac{\pi}{2}\beta \left[\left(\tan \frac{\pi\alpha}{2} \right)^2 + 1 \right], \\ \partial_\beta \zeta &= -\tan \frac{\pi\alpha}{2},\end{aligned}\tag{2.31}$$

and

$$\begin{aligned}\partial_x h &= t, \\ \partial_\alpha h &= (t^\alpha - t)\partial_\alpha \zeta + t^\alpha \log(t)\zeta, \\ \partial_\beta h &= (t^\alpha - t)\partial_\beta \zeta.\end{aligned}\tag{2.32}$$

In order to obtain expressions for $\partial_\alpha h$ and $\partial_\beta h$ at $\alpha = 1$, we compute the limit as $\alpha \rightarrow 1$ of the corresponding expressions in (2.32):

$$\begin{aligned}\lim_{\alpha \rightarrow 1} \partial_\alpha h &= \frac{\beta}{\pi} t \log^2 t, \\ \lim_{\alpha \rightarrow 1} \partial_\beta h &= \frac{2}{\pi} t \log t.\end{aligned}\tag{2.33}$$

Since h is continuous in all parameters at $\alpha = 1$ and both one-sided limits exist, the values of $\partial_\alpha f$ and $\partial_\beta f$ are well-defined when $\alpha = 1$. Finally, we have

$$\begin{aligned}\partial_x f(x; \alpha, \beta) &= -\frac{1}{\pi} \int_0^\infty t \sin h e^{-t^\alpha} dt, \\ \partial_\alpha f(x; \alpha, \beta) &= -\frac{1}{\pi} \int_0^\infty (\sin h \partial_\alpha h + t^\alpha \cos h \log t) e^{-t^\alpha} dt, \\ \partial_\beta f(x; \alpha, \beta) &= -\frac{1}{\pi} \int_0^\infty \sin h \partial_\beta h e^{-t^\alpha} dt.\end{aligned}\tag{2.34}$$

The partial derivatives in (2.34) have a relatively compact form. In contrast, the partial derivatives with respect to α and β of the stationary phase integral (2.13) and the series expansions (2.18) and (2.25) become rather unwieldy. Nevertheless, ∂_α of the stationary phase integral and series representation of f was computed in [8] (for the symmetric case). However, this approach becomes cumbersome in the general case, as numerous applications of the product and chain rule make the expressions impractically long.

In order to compute the derivatives of the series representations (2.18) and (2.25) with respect to x , differentiation can be done term-by-term. See Appendix A for this calculation.

3 Generalized Gaussian quadrature

In this section we briefly discuss what are known as *generalized Gaussian quadrature rules*. These integration schemes are analogous to the Gaussian quadrature rules for orthogonal polynomials, except that they are applicable to wide classes of functions, not merely polynomials. See [4] for a description of classical Gaussian quadrature with regard to polynomial integration. Generalized Gaussian quadrature schemes were first rigorously introduced in [7, 21]. Recently, a more efficient scheme for their construction was developed in [1]. It is this more recent algorithm that we base our calculations on, and outline the main ideas here. See both of these references for a detailed description of these quadrature rules.

3.1 Gaussian quadrature

A k -point quadrature rule consists of a set of k nodes and weights, which we will denote by $\{x_j, w_j\}$. These nodes and weights are chosen to accurately approximate the integral of a function f with respect to a positive weight function ω :

$$\int_a^b f(x) \omega(x) dx \approx \sum_{j=1}^k w_j f(x_j). \quad (3.1)$$

Many different types of quadrature rules exist which exhibit different behaviors for different classes of functions f . In short, if a k -point quadrature rule exists which integrates k linearly independent functions f_1, \dots, f_k we say that the quadrature rule is a *Chebyshev quadrature*. If the k -point rule is able to integrate $2k$ functions f_1, \dots, f_{2k} then we say that the rule is *Gaussian*.

In the case where the f_ℓ are polynomials, the nodes and weights of the associate Gaussian quadrature can be determined from the class of orthogonal polynomials with corresponding weight function ω . However, in the case where the f_ℓ 's are arbitrary square-integrable functions, these nodes and weights must be determined in a purely numerical manner.

3.2 Nodes and weight by nonlinear optimization

We now provide an overview of the numerical procedure for constructing a Gaussian quadrature rule for the integrand in equation (2.9) using the procedure of [1]. Recall, we are constructing a quadrature rule to compute:

$$f(x; \alpha, \beta) = \frac{1}{\pi} \int_0^\infty \cos(h(t; x, \alpha, \beta)) e^{-t^\alpha} dt, \quad (3.2)$$

i.e., the goal is to compute integrals of the functions we will denote by

$$\phi(t; x, \alpha, \beta) = \cos(h(t; x, \alpha, \beta)) e^{-t^\alpha}. \quad (3.3)$$

For reasons of clarity, we address computing a generalized Gaussian quadrature scheme for a class of functions $\psi = \psi(t; \eta)$, i.e. those that depend on only one parameter, η . The multi-parameter case is analogous. The following discussion is cursory, and we direct the reader to [1] for more details, as

there are several aspects of numerical analysis, optimization, and linear algebra that would merely distract from the current application.

For a selection of $2n$ linearly-independent functions ψ_ℓ , we note that the corresponding n -point generalized Gaussian quadrature $\{t_j, w_j\}$ is the solution to the following system of $2n$ non-linear equations:

$$\begin{aligned} \sum_{j=1}^n w_j \psi_1(t_j) &= \int \psi_1(t) dt, \\ &\vdots = \vdots \\ \sum_{j=1}^n w_j \psi_{2n}(t_j) &= \int \psi_{2n}(t) dt. \end{aligned} \tag{3.4}$$

Obtaining a solution to this system is the goal of the following procedure.

The scheme proceeds by first finding an orthonormal set of functions u_ℓ such that any $\psi(\cdot, \eta)$ can be approximated, to some specified precision ϵ , as a linear combination of the u_ℓ for any η . Next, an oversampled quadrature scheme that integrates *products of these functions* is constructed, by using, for example, adaptive Gaussian quadrature [17]. Adaptive Gaussian quadrature proceeds by dividing the interval of integration into several segments such that on each segment, the integral is computed to a specified accuracy. The accuracy of each quadrature is determined by comparing with the value obtained on a finer subdivision of the interval.

For $2n$ functions, this means that we have a m -point quadrature rule $\{t_j, w_j\}$ such that

$$\left| \int u_k(t) u_\ell(t) dt - \sum_{j=1}^m w_j u_k(t_j) u_\ell(t_j) \right| \leq \epsilon, \tag{3.5}$$

for all $1 \leq k, \ell \leq 2n$, with $m \geq 2n$. Accurately integrating products of the u_ℓ 's allows for stable interpolation to be done for any $\eta \neq \eta_\ell$ [1]. At this point, the vectors $\mathbf{u}_\ell \in \mathbb{R}^m$ serve as finite dimensional embeddings of the square-integrable functions u_ℓ :

$$\mathbf{u}_\ell = \begin{bmatrix} \sqrt{w_1} u_\ell(t_1) \\ \vdots \\ \sqrt{w_m} u_\ell(t_m) \end{bmatrix}. \tag{3.6}$$

Here, the $u_\ell(t_j)$'s are scaled so that $\mathbf{u}_\ell^T \mathbf{u}_\ell \approx \|u_\ell\|_2^2$. Computing a rank-revealing **QR** decomposition of the matrix \mathbf{U} ,

$$\mathbf{U} = \begin{bmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_m^T \end{bmatrix}, \tag{3.7}$$

allows for the immediate construction of a $2n$ -point Chebyshev quadrature rule. This procedure, equivalently, has selected $2n$ values of t_j that can serve as integration (and interpolation) nodes for all of the u_ℓ 's. Refining this $2n$ -point Chebyshev quadrature down to an n -point quadrature proceeds via a Gauss-Newton optimization. On each step, a single node-weight pair (t_i, w_i) is chosen

to be discarded and the remaining nodes and weights are optimized. The procedure proceeds until roughly n nodes remain, or accuracy in the resulting quadrature starts to suffer. While the weights we obtained as a result of this optimization procedure happened to be positive, no explicit effort was made to ensure this. Theoretical considerations for the existence of positive weights can be found in [21].

Remark 1. Note that in the symmetric case, we must obtain a selection of $x_{\ell_1}, \alpha_{\ell_2}$ that yield a (possibly redundant) basis for all ϕ . This could be done via adaptive discretization in these variables, but in practice, we merely sample x, α at Chebyshev points in Region I in Figure 3a. The parameter α is sampled at roughly 100 Chebyshev points in $[0.5, 2.0]$, and then for each of these values α_{ℓ_2} , x is sampled at roughly 100 Chebyshev point in the interval $[0, B_{40}^\infty(\alpha_{\ell_2})]$. This yields an initial set of 10,000 functions which are then compressed and integrated. In order to make sure this sampling in x and α provides a suitable set of functions to span the space of all ϕ , we rigorously test the quadrature at many thousands of random locations in Region I, comparing against adaptive quadrature. The asymmetric case is analogous, with equispaced sampling in x, α , and β . Experimentally, in order to obtain high accuracy in the resulting quadrature, more nodes are required in the sampling of x than in α or β .

3.3 A rank-reducing transformation

In the brief description of the above algorithm, we assumed that the interval of integration for the quadrature was finite. In our case, the interval in (3.2) is infinite, but can be truncated given that it decays quickly due to the term e^{-t^α} . A common interval of integration for all $\phi(\cdot; x, \alpha)$ can be obtained based on the decay of e^{-t^α} for the smallest α under consideration. In fact, for a particular precision ϵ , we can set the upper limit of integration to be $T_\alpha = (-\log \epsilon)^{1/\alpha}$. Using this limit, we can redefine each integrand ϕ under a linear transformation:

$$\begin{aligned}
 f(x; \alpha, \beta) &\approx \frac{1}{\pi} \int_0^{T_\alpha} \phi(t; \alpha, \beta) dt \\
 &= \frac{T_\alpha}{\pi} \int_0^1 \phi(\tau T_\alpha; \alpha, \beta) d\tau \\
 &= \frac{T_\alpha}{\pi} \int_0^1 \cos(h(\tau T_\alpha; \alpha, \beta)) e^{-(\tau T_\alpha)^\alpha} d\tau \\
 &= \frac{T_\alpha}{\pi} \int_0^1 \tilde{\phi}(\tau; \alpha, \beta) d\tau.
 \end{aligned} \tag{3.8}$$

Computing generalized Gaussian quadratures for the functions $\tilde{\phi}$ turns out to be much more efficient due to the similarity of numerical support (i.e. those t such that $|\tilde{\phi}(t)| > \epsilon$). This change of variables can significantly reduce the rank obtained in the rank-revealing **QR** step of the previous non-linear optimization procedure. For example, the generalized Gaussian quadrature for functions in Region I in Figure 3a consisted of 100 nodes/weights before the change of variables, and only 43 nodes/weights afterward. The resulting quadrature can be applied to the original function ϕ via a straightforward linear transformation of the nodes and scaling of the weights.

To be fair, the stationary phase integral (2.13) too, permits such a rank-reducing transformation.

However, it turns out to be much less efficient because it relies on an *a priori* zero-finding procedure. Notably, (2.13) has changing numerical support for different choices of the input parameters x , α , and β . This is true even after accounting for the parameter-dependent interval of integration by composing the integrand with a linear map from $[0, 1]$ to $[-\theta_0, \pi/2]$. Experimentally, even though the integrand decays to zero at both $-\theta_0$ and $\pi/2$, the differing numerical support is primarily caused by the exponential behavior of the integrand on the side of the interval of integration where $e^{-g(\theta)} \rightarrow 0$. Making use of this observation, for $\alpha \leq 1$ solving for θ_ϵ such that $g(\theta_\epsilon)e^{-g(\theta_\epsilon)} < \epsilon$ allows the integral in (2.13) to be approximated as:

$$f(x; \alpha, \beta) = \frac{\alpha}{\pi|\alpha - 1|} \frac{1}{(x - \zeta)} \int_{\theta_\epsilon}^{\pi/2} g(\theta; x, \alpha, \beta) e^{-g(\theta; x, \alpha, \beta)} d\theta, \quad (3.9)$$

and for $\alpha > 1$:

$$f(x; \alpha, \beta) = \frac{\alpha}{\pi|\alpha - 1|} \frac{1}{(x - \zeta)} \int_{-\theta_0}^{\theta_\epsilon} g(\theta; x, \alpha, \beta) e^{-g(\theta; x, \alpha, \beta)} d\theta. \quad (3.10)$$

A change of variable in these integrals can translate the interval of integration to $[0, 1]$. If the generalized Gaussian quadrature construction procedure is applied to these formulae, we also observe a reduction in the number of nodes and weights required. For example, in the symmetric case, solving for $g(\theta_\epsilon) = 40$ (which yields double-precision decay) the rank of the matrix \mathbf{U} in (3.7) decreased from 290 to 68 for $x \in [10^{-5}, 30]$ and $\alpha \in [.6, .8]$

In practice, the transformation of the stationary phase integral does not prove to be efficient because it relies on an initial zero-finding procedure to construct transformation in θ depending on each parameter x , α , β . In contrast, the change of variables in (3.8) merely requires evaluating \log . Similar changes of variables can be used to simplify the construction of quadratures for evaluating gradients of f .

3.4 Alternatives to generalized Gaussian quadrature

A number of quadrature techniques which are particularly effective for highly oscillatory integrands have been developed relatively recently. See [5, 15, 16] for an informative overview of such methods. Two notable examples which we will discuss here are Filon-type and Levin-type methods. These techniques are applicable to integrals of the form

$$\int f(t) e^{i\omega g(t)} dt, \quad (3.11)$$

where f and g are smooth, non-oscillatory functions, and ω is a scalar. The function g is called the *oscillator*.

The integrand of (2.9) is oscillatory, and therefore one could consider applying either Filon- or Levin-type quadrature schemes instead of the generalized Gaussian quadrature rules. Unfortunately, for general parameters ranges, generalized Gaussian quadratures are likely to be the most efficient schemes. We briefly justify this statement with a discussion of Filon and Levin methods.

Filon-type quadratures are interpolatory quadrature rules. That is, they approximate the function

f with a set of functions ψ_k for which an analytical solution of the integral

$$\mu_k = \int \psi_k(t) e^{i\omega g(t)} dt \quad (3.12)$$

exists. This poses a problem for the application of Filon-type quadratures to the integral (2.9). Indeed, note that (2.9) can be written in the form of (3.11) by setting

$$\begin{aligned} f(t) &= e^{-t^\alpha}, \\ g(t) &= (x - \zeta)t + \zeta t^\alpha. \end{aligned} \quad (3.13)$$

Since g depends on α , β , and x , the integrals μ_k have to be recalculated for every evaluation of a stable density with differing parameters. In contrast, the generalized Gaussian quadratures we derived are applicable for a wide range of parameters (see Figure 3).

Turning to Levin-type methods, they can be illustrated with the following observation [16]. Let $u(t)$ be a function which satisfies

$$\frac{d}{dt} \left(u(t) e^{i\omega g(t)} \right) = f(t) e^{i\omega g(t)}. \quad (3.14)$$

We then have

$$\int_a^b f(t) e^{i\omega g(t)} dt = u(b) e^{i\omega g(b)} - u(a) e^{i\omega g(a)}. \quad (3.15)$$

From (3.14), we can also derive the differential equation

$$u' + i\omega g' u = f, \quad (3.16)$$

where u' and g' denote differentiation with respect to t . The problem of computing an oscillatory integral has therefore been converted to that of solving a first-order, linear differential equation on the interval $[a, b]$.

Levin-type methods will require the solution of this ODE *every time* an integral has to be evaluated. Even with high-order convergent ODE solvers, these methods are unlikely to beat generalized Gaussian quadrature methods in terms of floating-point operations (after suitable offline pre-computations). Levin-type methods could be applicable to the parameter range $.9 < \alpha < 1.1$, $\beta \neq 0$, where modestly-sized generalized Gaussian quadratures are not available. For example, if a Chebyshev spectral method is used to solve (3.16), numerical experiments indicate that the condition number of the system can reach $\sim 10^6$ before the solution u can be fully resolved (requiring ~ 1000 Chebyshev terms). Therefore, a significant loss of accuracy with the Chebyshev spectral method is likely. Indeed, sometimes only 9 significant digits are achieved with this approach. A more specialized solver would have to be developed to solve (3.16) efficiently and with higher accuracy.

Levin-type methods could also be used for integral representations of the partial derivatives of stable densities, for which no asymptotic expansion is available. However, it is likely to be cheaper to form a Chebyshev interpolant of the density as a function of the parameter, and then differentiate the series (i.e. perform 2D interpolation and spectral differentiation). This will achieve higher accuracy than a finite difference scheme, with a slightly higher computational cost.

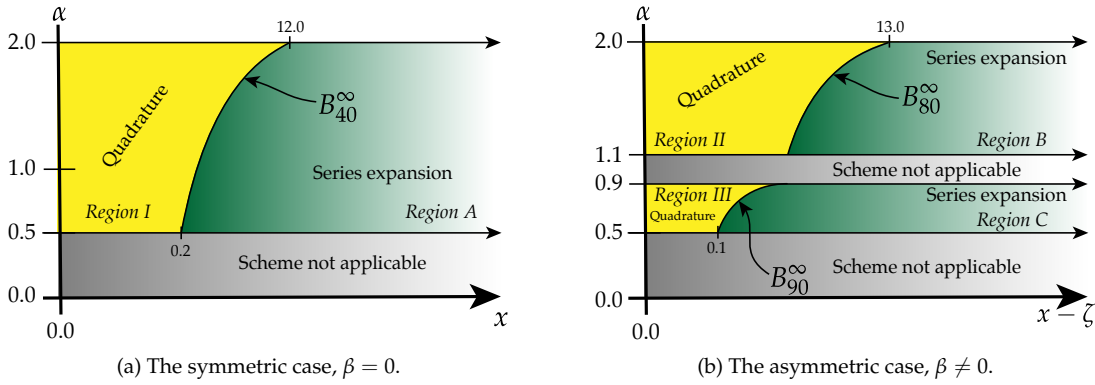


Figure 3: Regions of validity for generalized Gaussian quadrature rules and series approximations for the evaluation of $f(x; \alpha, \beta)$. Asymptotic expansions are used for extreme values of x and generalized Gaussian quadrature routines are able to fill in large regions of the remaining space.

4 Algorithm & Numerical examples

In the following, we will describe the details of our algorithm. In particular, we detail which formula or quadrature should be used depending on values of the parameters x , α , β . We begin with a comparison of the benefits of the two integral representations given by (2.9) and (2.13).

4.1 Choosing an integral representation

It has become clear after several numerical experiments that the stationary phase integral (2.13), while seemingly simpler to evaluate than (2.9), carries several disadvantages. Namely, it cannot be used to reliably evaluate f when $x \sim \zeta$ (the mode of the distributions in the symmetric case), $\alpha \sim 1$, or for very large x . Furthermore, its partial derivatives suffer from the same deficiencies and have rather unwieldy forms. Lastly, the rank-reducing technique of Section 3.3 is not as effective when applied to (2.13). This results in quadratures of much larger sizes when compared to those for (2.9).

In contrast, (2.9) has only one of the aforementioned deficiencies: it cannot be evaluated efficiently when $\alpha \sim 1$ in the asymmetric case. Still, it is important to point out that (2.9) can be easily evaluated at $\alpha \sim 1$ in the symmetric case.

The stationary phase form (2.13) does have two advantages over (2.9). First, it is well behaved for intermediate to large x , whereas (2.9) becomes very oscillatory. Second, it can be evaluated for $\alpha < 0.5$, whereas the relevant interval of integration of (2.9) grows rather fast as $\alpha \rightarrow 0$. However, the series expansion (2.25) is a much more efficient means of evaluation in these regimes. This limits the usefulness of the stationary phase integral for our purpose.

As a result, the only integral representation of the density we use in our algorithm is given by (2.9). One consequence of this choice is that we do not need to use the series expansion around $x = \zeta$ given in (2.18), as the integral is well behaved there. For similar reasons, we use (2.17) to compute F , and the integral representations for the gradient of f given in (2.34).

4.2 The symmetric case $\beta = 0$

We first provide some numerical examples of the accuracy and efficiency of evaluating the symmetric densities

$$f(x; \alpha, 0) = \frac{1}{\pi} \int_0^\infty \cos(xt) e^{-t^\alpha} dt. \quad (4.1)$$

We restrict our attention to values of f for $\alpha \geq 0.5$ for two reasons. First, when α is much smaller than 0.5, and x is close to but not equal to ζ , existing numerical schemes for integral representations and series expansions require prohibitive computational cost to achieve reasonable accuracy. And second, the applications for modeling with stable laws with such small values of the stability parameter α seem to be very rare. Nevertheless, it should be pointed out that when $\alpha < 0.1$, the series (2.25) with $n_\infty = 128$ terms is accurate to double precision for $x - \zeta \geq 10^{-16}$. At $x = \zeta$, the first term of the series (2.18) can be used to obtain an accurate value of the density. Therefore, in this extreme regime, effective numerical evaluations of stable laws is possible using the series expansions alone.

We now move to a description of our evaluation scheme. For a particular value of $\alpha \geq 0.5$, if $x \leq B_{40}^\infty$ (this corresponds to Region I in Figure 3a), we use a 43-point generalized Gaussian quadrature to evaluate the above integral. If $x > B_{40}^\infty$, we use series expansion (2.25). The number of terms in the series expansion was chosen (experimentally, as a precomputation) to roughly equal the number of nodes in the optimized quadrature. A similar method is used for the computation of F and the gradient of f . However, there is a notable difference in the computation of $\partial_\alpha f$ as it does not permit a convenient series expression, as noted in Section 2.4. Instead, we use a finite difference scheme applied to the series (2.25) to compute $\partial_\alpha f$ when $x > B_{43}^\infty$. The accuracy of finite difference schemes depends on the particular scheme used. In practice, a two-point finite difference is accurate to about 10^{-6} , while a fourth order scheme is accurate to about 10^{-10} . The fourth order scheme is listed in the Appendix A.

Accuracy results for f and its gradient are reported in Table 1. The columns are as follows:

- n_{GGQ} : the number of nodes in the generalized Gaussian quadrature scheme,
- n_∞ : the number of terms used for the series (2.25), and
- max error: the maximum absolute L_∞ error relative to adaptive integration.

The accuracy results were obtained by testing our quadrature scheme against an adaptive integration evaluate of (2.9) at 100,000 randomly chosen points in the $x\alpha$ -plane for $x \in [0, B_{n_\infty}^\infty]$ and $\alpha \in [0.5, 2.0]$. We should note that all results are reported in *absolute* precision. When evaluating integrals with arbitrarily sign-changing integrands via quadrature, if the integral is of size δ then it is likely that $\mathcal{O}(|\log \delta|)$ digits will be lost in *relative precision* due to the cancellation effect inherent in

Table 1: Symmetric ($\beta = 0$) stable density evaluation for $\alpha \in [.5, 2.0]$.

	n_{GGQ}	n_∞	max err
f	43	42	5 E-14
$\partial_x f$	44	42	9 E-14
$\partial_\alpha f$	49	42	1 E-13
F	39	42	9 E-14

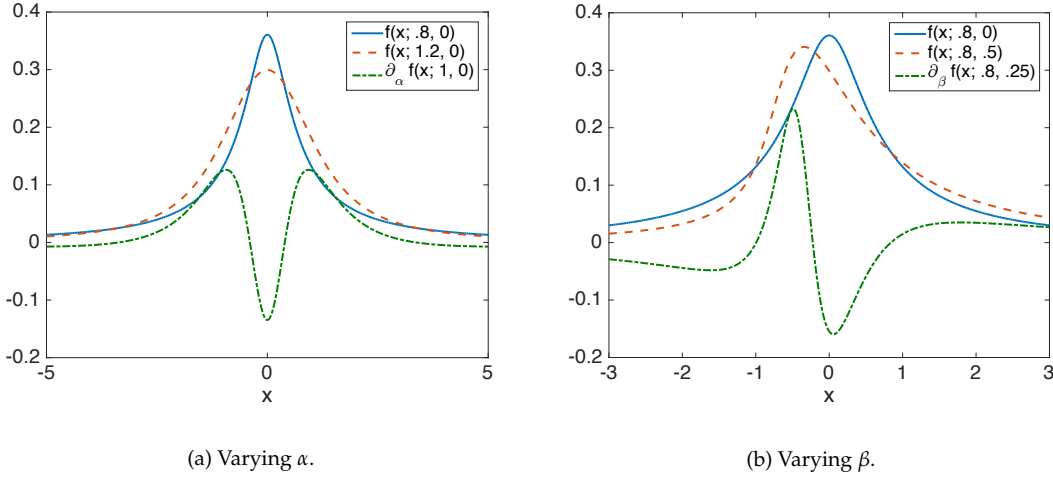


Figure 4: Effect of changing parameters of the density $f(x; \alpha, \beta)$, along with partial derivatives.

floating-point arithmetic. Table 2 contains the 43-point quadrature for evaluating stable densities in Region I of Figure 3a.

We should note that while it is possible to construct more efficient quadratures for smaller regions of the $x\alpha$ -plane, namely for $\alpha > 1$ (distributions with finite expectation), it is useful to obtain a single global quadrature value in a lone region. As shown in Figure 4a, small changes in α induce equivalently small changes in the density (a rather low-rank update in Fourier-space). In particular applications with restricted stability parameters, it may be prudent to construct even more efficient quadratures. There are several parameter combinations or ranges that might benefit from specialized quadrature. For example, the Holtmark distribution ($\alpha = 1.5$, $\beta = 0$) occurs in statistical investigations of gravity [2, 3]. The methods of this paper can be applied to compute this distribution, and others, very efficiently.

4.3 The asymmetric case $\beta \neq 0$

In the asymmetric case, $\beta \neq 0$, we first change variables and evaluate the densities at locations relative to: $x - \zeta$. This ensures that the densities are continuous in all parameters. As in the symmetric case, we restrict our attention to densities with $\alpha \geq 0.5$. Furthermore, due to difficulties in the integral and series formulations near $\alpha = 1$, we partition the α space into two regions: $[0.5, 0.9]$ and $[1.1, 2.0]$. For values of $\beta \neq 0$, $|\zeta| \rightarrow \infty$ as $\alpha \rightarrow 1$. This is the main mode of failure for both integral representations (2.9) and (2.13) near $\alpha = 1$. As a consequence, the integrand in (2.9) becomes highly oscillatory for even small values of $x - \zeta$, and (2.13) becomes spiked, as seen in Figure 2. Quadrature techniques developed to deal with highly oscillatory integrands may be applicable in this regime, and will be investigated in future work.

For calculating asymmetric densities, the parameter space is partitioned in the following manner: for all $0 \leq x - \zeta \leq B_{n_\infty}^\infty$, the densities are calculated via a generalized Gaussian quadrature scheme for the integral (2.9). For $x - \zeta > B_{n_\infty}^\infty$, the series expansion (2.25) is used. As mentioned previously, we have not obtained a convenient series representation of $\partial_\alpha f$ and $\partial_\beta f$. Similar to the computation

Table 2: Nodes and weights for computing the integral in (4.1) in Region I of Figure 3a. Note that the change of variables discussed in Section 3.3 must be used before applying the quadrature.

j	t_j	w_j
1	3.8153503841778930 E-08	1.9462166165433782 E-07
2	1.8621751229398742 E-06	5.6557228645853394 E-06
3	2.3548989111566051 E-05	5.0123980914007912 E-05
4	1.4796873542253231 E-04	2.3484191896467563 E-04
5	5.9719633529811916 E-04	7.3189687338231666 E-04
6	1.7776065804175705 E-03	1.7238717892356147 E-03
7	4.2473152693930051 E-03	3.3181618633886167 E-03
8	8.6062904061371317 E-03	5.4843557934027244 E-03
9	1.5348863951004616 E-02	8.0460517169448388 E-03
10	2.4742939762206897 E-02	1.0741992568943348 E-02
11	3.6794136418563730 E-02	1.3324899124821651 E-02
12	5.1299788260226145 E-02	1.5632319416985788 E-02
13	6.7944092105184303 E-02	1.7598001767457079 E-02
14	8.6382423526857308 E-02	1.9224720756886148 E-02
15	1.0629323929619865 E-01	2.0550906564542663 E-02
16	1.2740084223127754 E-01	2.1626845166204386 E-02
17	1.4948000254495675 E-01	2.2501767869303416 E-02
18	1.7235168105825832 E-01	2.3218324218440851 E-02
19	1.9587547846015377 E-01	2.3811106669646236 E-02
20	2.1994170091684220 E-01	2.4307093498802106 E-02
21	2.4446430088367060 E-01	2.4726814975746716 E-02
22	2.6937507294734536 E-01	2.5085627984821550 E-02
23	2.9461905048621601 E-01	2.5394814769833289 E-02
24	3.2015086713296453 E-01	2.5662404915729992 E-02
25	3.4593177859400515 E-01	2.5893662976614856 E-02
26	3.7192698736533930 E-01	2.6091208338375568 E-02
27	3.9810287487972756 E-01	2.6254650270947675 E-02
28	4.2442340488910107 E-01	2.6379218475411595 E-02
29	4.5084450818929106 E-01	2.6452938570694140 E-02
30	4.7730398807466573 E-01	2.6449899946836126 E-02
31	5.0370157776242630 E-01	2.6317157280117857 E-02
32	5.2986292621392794 E-01	2.5956236923454400 E-02
33	5.5549151318191370 E-01	2.5233135355892482 E-02
34	5.8023336057818919 E-01	2.4253858975533026 E-02
35	6.0420106522936246 E-01	2.3883208046443095 E-02
36	6.2845118361063135 E-01	2.4788564380040439 E-02
37	6.5391666500166423 E-01	2.6135655855085593 E-02
38	6.8067763680759019 E-01	2.7386023987669407 E-02
39	7.0883363435562430 E-01	2.9079584045104245 E-02
40	7.3935214962210505 E-01	3.2403729259281477 E-02
41	7.7501382927296592 E-01	3.9683359210637488 E-02
42	8.1983271443438077 E-01	5.0313579393503942 E-02
43	8.7653187131388799 E-01	6.3807406535572972 E-02

Table 3: Asymmetric ($\beta \neq 0$) stable density evaluation for $\alpha \in [0.5, 0.9]$ and $\alpha \in [1.1, 2.0]$.

(a) $\alpha \in [.5, .9]$				(b) $\alpha \in [1.1, 2.0]$			
	n_{GGQ}	n_∞	max err		n_{GGQ}	n_∞	max err
f	94	90	5 E−14	f	86	80	2 E−14
$\partial_x f$	110	90	1 E−13	$\partial_x f$	96	80	2 E−14
$\partial_\alpha f$	113	90	9 E−14	$\partial_\alpha f$	98	80	9 E−14
$\partial_\beta f$	109	90	5 E−14	$\partial_\beta f$	93	80	4 E−14
F	181	90	1 E−8	F	88	80	1 E−14

Table 4: Timings for density evaluations

t_{GQ}	$t_{GQ}(\beta = 0)$	t_{AQ1}	t_{AQ2}
0.003 sec	0.001 sec	0.3 sec	2 sec

of $\partial_\alpha f$ in the symmetric case, we use finite differences to approximate values of $\partial_\alpha f$ and $\partial_\beta f$ whenever $x > B_{n_\infty}^\infty$. Depending on the finite difference scheme used, this may lead to reduced accuracy compared to the quadrature method for the computation of f and $\partial_x f$. Similar accuracy reports to those for the symmetric densities are contained in Tables 3a and 3b. Notably, the quadrature rule for F in the regime $\alpha \in [.5, .9]$ is less accurate and has more nodes/weights than for the other functions. This is due to the fact that the integrand in (2.17) for F is singular at $t = 0$ in the asymmetric case. As a consequence, designing highly-accurate quadrature rules for (2.17) without using quadruple precision calculations is not possible. This issue will be investigated in future work. The corresponding quadrature rules are available for download at gitlab.com/s_ament/qastable.

4.4 Efficiency of the method

To test the efficiency of our method, we compare our implementation of the density function evaluation to two different implementations based on adaptive quadrature. All codes are written in MATLAB. The first implementation simply applies MATLAB’s `integral` function to the oscillatory integral (2.9). Note that this function can be called in a vectorized manner by adjusting the `ArrayValued` argument. Without this adjustment, the computations below are about an order of magnitude slower. (I.e. t_{AQ1} and t_{AQ2} are roughly 10 times as large.) The second implementation mimics the approach that was previously taken to compute the stationary phase integral (2.13). Namely, it first locates the peak of the integrand using MATLAB’s intrinsic `fzero` function, and subsequently applies `integral` on the two subintervals created by splitting the original interval of integration at the peak of the integrand.

The validation test proceeds as follows. First, α and β are chosen randomly in the permissible parameter ranges. Then, 10,000 uniformly random x are generated such that $0 \leq x - \zeta \leq 20$. Thereafter, we record the wall-clock time each method takes to calculate the stable density at all 10,000 points. For our tests, we require the absolute accuracy of the adaptive schemes to be 10^{-10} . The results are reported in Table 4. The columns of the table are:

t_{GQ} : the time taken by our scheme to compute the density at all points,

t_{AQ1} : time taken by the first adaptive scheme outlined above, and
 t_{AQ2} : the time taken by the second adaptive scheme outlined above.

We also report a timing for the symmetric case ($\beta = 0$) for our scheme, since it uses a quadrature separate from the one in the asymmetric case. The test was performed on a MacBook Pro with a 2.4 GHz Intel Core i7 and 8 GB 1333 MHz DDR3 RAM. As one can see, our scheme outperforms the adaptive ones by at least two orders of magnitude.

5 Conclusions

In this work, we have presented efficient quadrature schemes and series expansions for numerically evaluating the densities, and derivatives thereof, associated with what are known as stable distributions. The quadratures are of generalized Gaussian type, and were constructed using a non-linear optimization procedure. The series expansions were obtained straightforwardly from integral representations, but seem to have not been previously presented in the computational statistics literature. The methods of this paper are quite efficient, and easily vectorizable. This is in contrast to existing schemes for evaluating these integrals, which were predominately based on adaptive integration – which cannot take full advantage of vectorization schemes due to varying depths of recursion.

Furthermore, while the quadratures that we constructed are (nearly) optimal with respect to the number of nodes and weights required, they do not obtain full double precision accuracy ($\sim 10^{-16}$). We often only achieve absolute accuracies of 12 or 13 digits. While some of the precision loss is due to merely roundoff error in summing the terms in the quadrature, some of the loss of accuracy is due to solving the ill-conditioned linearization of the quadrature problem. The accuracy lost due to this aspect of the procedure could be recovered if the quadrature generation codes were re-written using quadruple precision arithmetic instead of double precision. In most cases the accuracy we obtained is sufficient for general use, but we are investigating a higher precision procedure for constructing the quadrature rules.

The schemes presented in this paper still fail to thoroughly address the evaluation of the density function (and gradient and CDF) for values of $\alpha \approx 1$ in the asymmetric case. One could, however, perform a large-scale precomputation in extended precision in order to tabulate these densities for various values of x and β , store the results, and later interpolate to other values. This approach was beyond the scope of this work. This approach was used for maximum likelihood estimation in [12]. Unless chosen very carefully, a rather large number of interpolation nodes are necessary to achieve high accuracy, and each function ($f, \nabla f, F$) has to be tabulated separately. We are actively investigating approaches to fill in this gap in the numerical evaluation of the density (and gradient and CDF).

A software package written in MATLAB for computing stable densities, their gradients, and distribution functions using the algorithms of this paper is available at gitlab.com/s_ament/qastable, and will be continually updated as we improve the efficiency and accuracy of existing evaluations, and include additional capabilities.

Acknowledgments

The authors would like to thank Sinan Güntürk and Margaret Wright for several useful conversations.

A Gradients of series expansions

Here we provide formulae for the derivatives of the series expansions presented in Section 2.4. From (2.18),

$$\partial_x f(x; \alpha, \beta) = \frac{1}{\alpha\pi} \sum_{k=0}^{\infty} \frac{\Gamma(\frac{k+2}{\alpha})}{\Gamma(k)} (1 + \zeta^2)^{-\frac{k+2}{2\alpha}} \sin((\pi/2 + (\arctan \zeta)/\alpha)(k+2)) (x - \zeta)^k. \quad (\text{A.1})$$

Using an error bound analogous to the ones in Section 2.3, we have that

$$|x - \zeta| \leq C_n^0(\alpha) := \left[\epsilon \alpha \pi (1 + \zeta^2)^{\frac{n+2}{2\alpha}} \frac{\Gamma(n)}{\Gamma(\frac{n+2}{\alpha})} \right]^{1/n}. \quad (\text{A.2})$$

By differentiating (2.25), we attain

$$\partial_x f(x, \alpha, \beta) = \frac{\alpha}{\pi} \sum_{k=1}^{\infty} (-1)^k \frac{(\alpha k + 1) \Gamma(\alpha k)}{\Gamma(k)} (1 + \zeta^2)^{k/2} \sin((\pi\alpha - \arctan \zeta)k) (x - \zeta)^{-\alpha k - 2}, \quad (\text{A.3})$$

whose radius of convergence to precision ϵ we estimate by

$$|x - \zeta| \geq C_{n-1}^{\infty}(\alpha) := \left[\frac{\alpha}{\pi\epsilon} (1 + \zeta^2)^{\frac{n}{2}} \frac{(\alpha n + 1) \Gamma(\alpha n)}{\Gamma(n)} \right]^{1/(\alpha n - 2)}. \quad (\text{A.4})$$

For the parameter ranges where there is no convenient formulation of the derivatives, we can use a finite difference approximation of the form

$$\partial_x f(x) = \frac{-f(x+2h) + 8f(x+h) - 8f(x-h) + f(x-2h)}{12h} + O(h^4). \quad (\text{A.5})$$

References

- [1] J. Bremer, Z. Gimbutas, and V. Rokhlin. A nonlinear optimization procedure for generalized Gaussian quadratures. *SIAM J. Sci. Comput.*, 32:1761–1788, 2010.
- [2] S. Chandrasekhar. Stochastic problems in physics and astronomy. *Rev. Mod. Phys.*, 15(1):1–89, Jan 1943.
- [3] P. Chavanis. Statistics of the gravitational force in various dimensions of space: from gaussian to lévy laws. *The European Physical Journal B*, 70(3):413–433, 2009.
- [4] G. Dahlquist and A. Björck. *Numerical Methods*. Dover, New York, NY, 2003.

- [5] A. Iserles, S. Nørsett, and S. Olver. Highly oscillatory quadrature: The story so far. In *Numerical mathematics and advanced applications*, pages 97–118. Springer, 2006.
- [6] K. S. Kölbig and B. Schorr. A program package for the Landau distribution. *Comput. Phys. Commun.*, 31:97–111, 1984.
- [7] J. Ma, V. Rokhlin, and S. Wandzura. Generalized Gaussian Quadrature Rules for Systems of Arbitrary Functions. *SIAM J. Numer. Anal.*, 33:971–996, 1996.
- [8] M. Matsui and A. Takemura. Some improvements in numerical evaluation of symmetric stable density and its derivatives. *Multivar. Anal.*, 35:149–172, 2006.
- [9] S. Mittnik and S. T. Rachev. Modeling asset returns with alternative stable distributions. *Econometric Reviews*, 12(3):261–330, 1993.
- [10] C. L. Nikias and M. Shao. *Signal Processing with Alpha-Stable Distributions and Applications*. Wiley, New York, NY, 1995.
- [11] J. P. Nolan. Numerical calculation of stable densities and distribution functions. *Commun. Statist.-Stochastic Models*, 13(4):759–774, 1997.
- [12] J. P. Nolan. *Maximum likelihood estimation and diagnostics for stable distributions.*, chapter Lévy Processes: Theory and Applications. Birkhäuser Boston, 2001.
- [13] J. P. Nolan. Modeling financial data with stable distributions. In S. T. Rachev, editor, *Handbook of Heavy Tailed Distributions in Finance*, pages 105–130. Elsevier/North Holland, New York, NY, 2003.
- [14] J. P. Nolan. *Stable Distributions - Models for Heavy Tailed Data*. Birkhauser, Boston, 2015. In progress, Chapter 1 online at academic2.american.edu/~jpnolan.
- [15] S. Olver. *Numerical Approximation of Highly Oscillatory Integrals*. PhD thesis, University of Cambridge, 2008.
- [16] S. Olver. Fast, numerically stable computation of oscillatory integrals with stationary points. *Bit. Numer. Math.*, 50:149–171, 2010.
- [17] W. H. Press, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes*. Cambridge University Press, New York, NY, 3rd edition, 2007.
- [18] C. E. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- [19] N. G. Shephard. From characteristic function to distribution function: A simple framework for the theory. *Econometric Theory*, 7:519–529, 1991.
- [20] M. Teimouri and H. Amindavar. A novel approach to calculate stable densities. In *Proceedings of the World Congress on Engineering*, volume 1, 2008.

- [21] N. Yarvin and V. Rokhlin. Generalized Gaussian quadratures and singular value decompositions of integral operators. *SIAM J. Sci. Comput.*, 20:699–718, 1998.
- [22] V. M. Zolotarev. *One-dimensional stable distributions*. American Mathematical Society, Providence, RI, 1986.