

CORNELL UNIVERSITY BSCB DEPARTMENT
TECHNICAL REPORT BU-1605-M

Computing tail probabilities by numerical Fourier
inversion: the absolutely continuous case

Robert L. Strawderman
Department of Biological Statistics
and Computational Biology
Cornell University

October 29, 2002

Abstract

The numerical computation of $P\{X > x\}$ can be accomplished in a variety of ways. An appealing class of methods may be derived from the contour integral connecting $P\{X > x\}$ and its Fourier representation. Statisticians have largely focused on deriving saddlepoint approximations for this contour integral. The accuracy of such approximations is generally understood in vague terms only, and perhaps more importantly, is rarely under user control. Numerical integration of the contour integral has received considerably less attention, particularly in the statistics literature. The focus of this paper is on the use of the trapezoidal rule applied to said contour integral along an appropriate path. An exponential bound on the approximation (i.e., discretization) error of the trapezoidal rule as a function of the quadrature node spacing is obtained using results of Stenger (1993). This bound is used in developing a reliable non-iterative method of selecting the trapezoidal rule spacing that guarantees control of the approximation error. The epsilon algorithm is used to accelerate the calculation of the tail of the infinite series that results upon applying the trapezoidal rule to the inversion integral. The resulting “automatic” methodology is shown to produce extremely accurate results in a diverse set of problems.

CORRESPONDING AUTHOR:

Robert L. Strawderman
BSCB Department
437 Warren Hall
Cornell University
Ithaca NY 14853-7801
USA

Email: rls54@cornell.edu

Office: 607-255-1642

Fax: 607-255-4698

KEYWORDS: Analytic function; Characteristic function; Epsilon algorithm; Saddlepoint recentering; Sinc quadrature.

RUNNING TITLE: Tail probabilities via transform inversion

1 Introduction

The “exact” calculation of a tail probability $P\{X > x\}$ can be accomplished in a variety of ways, including direct computation through a closed form expression, numerical integration of a known probability density, summation of point probabilities, and so on. In principle, tail probabilities may also be computed via well-known Fourier inversion theorems linking $P\{X > x\}$ to the characteristic function of X ; see, for example, Kawata (1972) or Lukacs (1970). These inversion theorems have potentially significant value in cases where the characteristic function is easy to obtain and simple expressions for the associated tail probability are not. For example, as pointed out by Mehta *et al.* (1998), numerous problems in exact conditional inference may be well-suited to such methods since the characteristic function is either known or can be computed in polynomial time; in contrast, the direct computation of a tail probability can require enormous computational effort.

In statistics, the Fourier representation of a tail probability has primarily been exploited through saddlepoint approximation and related asymptotic methods; see Strawderman (2000) for a selective review. Widely cited sources detailing statistical uses of these approximations include Barndorff-Nielsen and Cox (1989, 1994), Daniels (1954, 1983, 1987), Jensen (1995), Kolassa (1997), Lugannani and Rice (1980), and Field and Ronchetti (1990). Spectacular results have been achieved in the case where X represents a sum of independent and identically distributed random variables. Outside of this comparatively straightforward setting, the conditions required for the validity of these asymptotic approximations become substantially more difficult to write down. Generally, the approximation error is only vaguely understood, typically being characterized in terms of some parameter controlling the peakedness of the integrand (e.g., sample size). The error of saddlepoint-type approximations is therefore not under user control in most statistical applications. These factors place significant obstacles before the trustworthy, routine use of saddlepoint approximations.

Numerical integration of the Fourier inversion integral has received considerably less attention from statisticians. Abate and Whitt (1992) provide an extensive and interesting discussion of various methods for numerical transform inversion, the focus primarily being applications in probability involving distributions supported on the positive real line. Later work providing additional details and/or refinements of certain algorithms discussed in Abate and Whitt (1992)

includes Abate and Whitt (1995) and Abate, Choudhury, and Whitt (1999). The quadrature-based methods considered in these paper are, for the most part, “Fourier series” (i.e., trapezoidal rule) approximations. Waller, Turnbull and Hardin (1995) provide a more limited, statistically oriented review of similar methods and in particular illustrate the use of one method previously described in Böhman (1975). Together, the respective reference lists of these two papers covers most of the work by statisticians on numerical transform inversion. The attraction of numerical quadrature is apparent: in principle, an exact answer can be obtained to any specified level of precision, the resulting error being largely controlled by the number and placement of quadrature nodes. However, the problem is also beset with numerical difficulties, most of which can be traced to the rapidly oscillating Fourier integrand (Davis and Rabinowitz, 1984). A recent survey of existing numerical quadrature schemes for oscillatory integrands can be found in Krommer and Ueberhuber (1998).

The Fourier integrands arising in the representation of tail probabilities have properties that are neither exploited by existing methods for transform inversion nor by more general quadrature schemes for oscillatory integrands. For example, when the moment generating function (MGF) of the random variable of interest exists, the corresponding characteristic function is analytic (Lukacs, 1970, Theorem 7.1.1). The analytic behavior of the integrand has a number of useful implications. Importantly, the path of integration in the contour integral may be altered, allowing substantial flexibility; an important and useful choice of path is that which passes through a saddlepoint (e.g., Rice, 1980; Helstrom, 1983). Upon such path modification, the resulting integral reduces to the Fourier transform of an analytic function, opening up a number of new possibilities from a computational point of view. Using the results of Stenger (1993), it is proved that the trapezoidal rule can be used to derive an extraordinarily accurate approximation to the tail probability inversion integral. This approximation suffers from an unfortunate computational drawback: it requires the summation of an infinite series. It will be shown, however, that this series can be accurately summed using a modest number of terms provided “convergence acceleration” is also employed. This paper exploits these observations in order to derive an effective quadrature rule for tail probability computations involving differentiable CDFs supported on any subset of the real line. The method to be described applies to the computation of almost any univariate tail probability, the main proviso being existence of the MGF (a condition also required for saddlepoint methods).

In essential respects the conclusions of this paper reflect those of Abate and Whitt (1992): the trapezoidal rule, possibly combined with convergence acceleration, is a powerful and effective tool for computing tail probabilities via numerical inversion. However, the implementation of the “Fourier series method” to be described here differs from the various versions described in Abate and Whitt (1992) in some important respects. For example, numerical integration is conducted along a path that depends on the ordinate (i.e., x) under consideration, and the proposed method neither requires explicit information on the support of X nor the tail behavior of its CDF. These differences stem primarily from the assumed existence of the MGF, which has important implications for the mathematical behavior of the characteristic function. One particularly useful consequence of this assumption is that the results of Stenger (1993) may be used to bound the approximation error of the trapezoidal rule, yielding a straightforward non-iterative method for selecting an appropriate spacing for the trapezoidal rule. The error bound depends primarily on the MGF and path of integration selected, and shows that the latter can have important implications for determining both the number of and spacing between quadrature nodes. In contrast, the error bounds for most of the quadrature methods detailed in Abate and Whitt (1992) are derived via the Poisson summation formula. For general probability distributions, these bounds require information about the tail behavior of the CDF (i.e., the quantity being computed) in order to determine an appropriate choice of spacing; however, it is possible to avoid this requirement in the case of probability distributions supported on the positive real line (see Remark D below). Another useful consequence of assuming existence of the MGF is the resulting asymptotic regularity of the zero pattern in the Fourier integrand. Such regularity is valuable, as it substantially increases the chances for successful convergence acceleration. For these purposes, the method of choice in Abate and Whitt (1992) is Euler summation. The epsilon algorithm is used here (cf. Brezinski and Redivo Zaglia, 1991), and is shown to significantly outperform Euler summation in essentially every example considered in Section 5.

The remainder of this paper will proceed as follows. In Section 2, the inversion integral representation of a tail probability is briefly reviewed. It is then shown that the desired integral may be written as a standard Fourier transform in the case of an absolutely continuous random variable. In Section 3.1, some important results for trapezoidal rules applied to general Fourier transforms are reviewed. In Section 3.2, this theory is used to obtain an exponential bound on the approximation error of the trapezoidal rule when applied to the tail probability inversion

integral. Section 4 provides detailed discussion of one possible implementation of this theory. In Section 5, the proposed methodology is applied to some diverse problems previously considered in the literature. The paper closes with a discussion in Section 6.

2 Fourier representation of tail probabilities

2.1 Preliminaries

Let X be a random variable with cumulative distribution function $F(\cdot)$. The random variable X may be a single random variable; alternatively, it can represent a convolution, a maximum, the ordinate of interest in a conditional inference problem, etc . . . For $t \in \mathbb{R}$, the characteristic function of a random variable X is defined as $\xi(t) = E[e^{itX}]$, where $i = \sqrt{-1}$; this complex-valued function always exists. Throughout, it is assumed that the MGF $M(t) = E[e^{tX}]$ exists for $t \in \mathcal{E}$, where \mathcal{E} is an open interval that contains $t = 0$. Then, the characteristic function and moment generating function are related via $\xi(t) = M(it)$. Finally, $K(t)$ shall denote the associated cumulant generating function (CGF).

Much of what follows requires an extension of the definitions of these quantities to complex-valued arguments. Let $z = x + iy \in \mathbb{C}$, where $x, y \in \mathbb{R}$ and \mathbb{C} denotes the complex plane; define $\Re(z) = x$ and $\Im(z) = y$ to respectively be the real and imaginary parts of z . Let $D_d = \{z \in \mathbb{C} : z = x + iy, x \in \mathbb{R}, |y| < d\}$ denote an infinite strip of width d containing the real axis. Define the complex-valued function $\xi(z) = E[e^{izX}] = M(iz)$; notice that the characteristic function of X is recovered from $\xi(z)$ upon setting $\Im(z)$ equal to zero. Importantly, existence of the MGF $M(\cdot)$ implies that $\xi(z) = M(-y + ix)$ is an analytic function for $z \in D_d$ (Lukacs, 1970, Theorem 7.1.1). If the random variable X has bounded support, then D_d becomes the complex plane \mathbb{C} (i.e., $\xi(z)$ is entire) (Lukacs, 1970, Theorem 7.2.3). Otherwise, the strip of regularity is of the form $-\alpha < \Im(z) < \beta$ for $\alpha, \beta > 0$, where $-i\alpha$ and $i\beta$ represent the singularities of $\xi(z)$ closest to the origin (Lukacs, 1970, Theorem 7.1.1).

2.2 The absolutely continuous case

In the case where X is absolutely continuous, there are numerous forms of the so-called “inversion integral” for $P\{X > x\}$; see, for example, Abate and Whitt (1992). When the MGF

exists, a general form providing a useful starting point is the following:

$$P\{X > x\} = \mathcal{H}(-c) + \frac{e^{\nu(c)}}{2\pi i} \int_{c-i\infty}^{c+i\infty} z^{-1} \exp\{\nu(z) - \nu(c)\} dz \quad (1)$$

where $c \in \mathcal{E}$, $\nu(z) = K(z) - xz$ for $z \in \mathbb{C}$ and $\mathcal{H}(w)$ respectively equals 0, $\frac{1}{2}$, or 1 if $w < 0$, $w = 0$, or $w > 0$. Saddlepoint approximations for this contour integral may be derived under suitable conditions. Typically, one selects $c = \hat{u}_x$, where $K'(\hat{u}_x) = x$; doing so,

$$P\{X > x\} = \mathcal{H}(-\hat{u}_x) + \frac{e^{\nu(\hat{u}_x)}}{2\pi i} \int_{\hat{u}_x-i\infty}^{\hat{u}_x+i\infty} z^{-1} \exp\{\nu(z) - \nu(\hat{u}_x)\} dz. \quad (2)$$

As $\hat{u}_x \rightarrow 0$ (i.e., as $x \rightarrow E[X]$), a pole occurs in the integrand on the right-hand side of (2) at $z = 0$. In general, solutions for this problem have been devised with a view towards generating valid asymptotic expansions; see, for example, Daniels (1987) and Kolassa (1997, Ch. 5). However, in view of Cauchy's theorem (cf. Bak and Newman, 1996) the choice of c is basically arbitrary and other selections are certainly possible. Helstrom (1983) proposed setting $c = \bar{u}_x$, where

$$\nu'(\bar{u}_x) - \bar{u}_x^{-1} = 0. \quad (3)$$

The solution to (3) maximizes the entire integrand in (2) (i.e. as opposed to just the exponential term). Taking the path of integration through \bar{u}_x is less desirable from an asymptotic point view because of the resulting difficulties associated with applying Watson's Lemma; for related discussion see Kolassa (1997, §5.6). However, it is an interesting and useful choice since

- $\bar{u}_x \approx \hat{u}_x$ as x moves towards the limits of the support of X , maintaining whatever benefits are afforded by saddlepoint methods for "extreme" x ;
- the solution to (3) remains bounded away from zero as $x \rightarrow E[X]$, resolving the potential numerical instability associated with using \hat{u}_x for x near $E[X]$.

The integral (1) with $c = \bar{u}_x$ will now be expressed in a less familiar form. Specifically, parameterizing the path in \mathbb{C} as $\bar{u}_x + it$ for $t \in \mathbb{R}$, one may write

$$P\{X > x\} = \mathcal{H}(-\bar{u}_x) + \frac{e^{\nu(\bar{u}_x)}}{2\pi} \int_{-\infty}^{\infty} g_x(t) e^{-ixt} dt, \quad (4)$$

where

$$g_x(t) = \frac{\exp\{K(\bar{u}_x + it) - K(\bar{u}_x)\}}{\bar{u}_x + it}. \quad (5)$$

The integral appearing in (4) takes the form of an (inverse) Fourier transform of the complex-valued function $g_x(t)$. Consequently, any quadrature rule appropriate for a Fourier transform can in principle be used to compute $P\{X > x\}$. The focus in this paper will be on the trapezoidal rule, which is explored in detail in the remaining sections. Notably, $g_x(t)e^{-ixt}$ generally has both a nonzero real and imaginary part, the latter being an odd function of t . It follows that (4) may be expressed in terms of $\Re\{g_x(t)e^{-ixt}\}$ only. However, this fact will not be exploited until later.

REMARK A: In deriving (4), the path of integration is taken to be a straight line passing through the saddlepoint \bar{u}_x . This is not equivalent to integrating along a path of steepest descent, which in principle could lead to rather substantial numerical gains. Integrating along a path of steepest descent is usually not done in practice because explicitly characterizing such paths is often impossible. Saddlepoint approximations are typically obtained by making a locally quadratic approximation to the path of steepest descent in some neighborhood of the saddlepoint. The interval on which this approximation holds need not be explicitly specified, and as a result contributions of the integrand outside this region are ignored in deriving the approximation. For tail probabilities involving a mean of iid random variables, these contributions vanish at a geometric rate depending on the sample size. This is not necessarily true more generally, and ignoring such contributions can lead to substantial inaccuracies in the final approximation.

3 Trapezoidal rules for tail probabilities

3.1 A review of important results

A trapezoidal rule is typically derived as the exact integral of a piecewise linear approximation to a given integrand. This derivation assumes the integrand has two continuous derivatives, leading one to conclude that the associated approximation error is $O(h^2)$, where h is the spacing between quadrature nodes. Interesting alternative derivations of the trapezoidal rule exist that better highlight the role of integrand smoothness and in particular provide substantial refinements of the $O(h^2)$ error bound. For example, using the Poisson summation formula (e.g., Feller, 1971, pp. 626-631), it is possible to establish the following:

Theorem 1 *Let $f(\cdot)$ be continuous, integrable, and of bounded variation on \mathbb{R} . Let $\hat{f}(\cdot)$ be the*

Fourier transform of f . Then, for $h > 0$,

$$\left| \int_{-\infty}^{\infty} f(t)dt - h \sum_{k=-\infty}^{\infty} f(kh) \right| = \sqrt{2\pi} \left| \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \widehat{f}(2k\pi h^{-1}) \right|.$$

The above shows that the exact approximation error depends directly on the tail behavior of the Fourier transform of the integrand. It is a well-known result in Fourier analysis that the smoother the signal, the faster the tails of its Fourier transform decay to zero (e.g., Kawata, 1972, §2.7). In other words: the smoother the integrand, the better the approximation afforded by the trapezoidal rule. Butzer and Stens (1983) exploit this fact to show that the approximation error in Theorem 1 is $o(h^r)$ if the (1) integrand has $r - 1$ absolutely continuous integrable derivatives; and, (2) r^{th} derivative is integrable and satisfies a Lipschitz-type condition.

The trapezoidal rule may also be derived as the exact integral of an interpolatory expansion of the integrand based on the Whittaker cardinal (or sinc) function; see, for example, Stenger (1993). When the integrand is analytic (i.e., differentiable when considered as a function of $z \in \mathbb{C}$) and other integrability conditions hold, contour integration may be used to show that the approximation error vanishes exponentially fast. The following result establishes the rate of convergence of the trapezoidal rule when applied to the Fourier transform of an analytic function; the error bound refines that given in Theorem 3.3.1 of Stenger (1993):

Theorem 2 For $d > 0$, let $D_d = \{z \in \mathbb{C} : z = x + iy, x \in \mathbb{R}, |y| < d\}$, and define $B(D_d)$ to be the set of functions f satisfying the following conditions:

- (i) For $z \in D_d$, $f(z)$ is analytic;
- (ii) $\int_{-d}^d |f(t + iy)| dy \rightarrow 0$ as $|t| \rightarrow \infty$.
- (iii) $\mathcal{N}_\omega(f, D_d) < \infty$, where $\mathcal{N}_\omega(f, D_d) = \lim_{y \rightarrow d^-} \int_{-\infty}^{\infty} [e^{\omega d} |f(t - iy)| + e^{-\omega d} |f(t + iy)|] dt$.

Suppose $f \in B(D_d)$. Then, for $0 < h < \frac{2\pi d}{\log 2}$ and $|\omega| < \pi h^{-1}$,

$$\left| \int_{-\infty}^{\infty} f(t)e^{i\omega t} dt - h \sum_{k=-\infty}^{\infty} f(kh)e^{i\omega kh} \right| \leq 2\mathcal{N}_\omega(f, D_d) \exp\{-2\pi d/h\}.$$

This result says that the approximation error of the trapezoidal rule vanishes exponentially fast as $h \rightarrow 0$ for functions $f(\cdot)$ that satisfy certain differentiability and integrability conditions when considered as a function of $z \in \mathbb{C}$. The convergence rate is affected by transform ordinate (i.e., $|\omega|$) and the width of the region in which f is analytic (i.e. d). In the next section, this result is used to bound the approximation error of the trapezoidal rule when computing tail probabilities via (4).

3.2 The approximation error for tail probabilities

Provided the required conditions are met, Theorem 2 applies directly to the integral appearing in (4) upon making the identifications $\omega = -x$ and

$$f(t) = g_x(t) = \frac{\exp\{K(\bar{u}_x + it) - K(\bar{u}_x)\}}{\bar{u}_x + it}.$$

The following result, proved in Appendix A, provides sufficient conditions under which Theorem 2 applies and, consequently, an exponential bound on the approximation error of the trapezoidal rule applied to (4).

Theorem 3 *Let X be an absolutely continuous random variable with CDF $F(\cdot)$ and density $F'(\cdot)$. Suppose $M(t) = \exp\{K(t)\} < \infty$ for $t \in \mathcal{E}$, where \mathcal{E} is an open interval containing $t = 0$. Let $\delta = \min\{|s| : s \notin \mathcal{E}\} > 0$, and for a given $x \in \mathbb{R}$, suppose $\min\{F(x), 1 - F(x)\} > 0$. Finally, suppose that for β such that $M(\beta) < \infty$ and some $r > 0$,*

$$\frac{M(\beta + is)}{M(\beta)} = O(s^{-r}) \tag{6}$$

as $|s| \rightarrow \infty$. Then, conditions (i)-(iii) of Theorem 2 are satisfied by $g_x(z)$ for $z \in D_{d^*}$, where:

- for $\mathcal{E} = \mathbb{R}$, $0 < d^* < |\bar{u}_x|$;
- for $\mathcal{E} = (-\infty, \delta)$, $0 < d^* < \min\{|\delta - \bar{u}_x|, |\bar{u}_x|\}$;
- for $\mathcal{E} = (-\delta, \infty)$, $0 < d^* < \min\{|\delta + \bar{u}_x|, |\bar{u}_x|\}$;
- for \mathcal{E} a bounded interval, $0 < d^* < \min\{|\delta + \bar{u}_x|, |\delta - \bar{u}_x|, |\bar{u}_x|\}$.

Let $h = \pi/\Delta$, where $\Delta > \max\{|x|, \frac{\log 2}{2d^*}\}$. Then,

$$\left| P\{X > x\} - \left[\mathcal{H}(-\bar{u}_x) + \frac{e^{\nu(\bar{u}_x)}}{\pi} \mathcal{T}_x(h) \right] \right| \leq \eta(\Delta, x)$$

where

$$\mathcal{T}_x(h) = \frac{h}{2} \sum_{k=-\infty}^{\infty} g_x(kh) e^{-ixkh}, \quad (7)$$

$$\eta(\Delta, x) = \frac{e^{\nu(\bar{u}_x)}}{\pi} \mathcal{N}_x(g_x, D_{d^*}) \exp\{-2\Delta d^*\},$$

and $\mathcal{N}_x(g_x, D_{d^*}) = \lim_{y \rightarrow d^{*-}} \int_{-\infty}^{\infty} [e^{-xd^*} |g_x(t - iy)| + e^{xd^*} |g_x(t + iy)|] dt$.

The approximation error $\eta(\Delta, x)$ decays exponentially as $\Delta \rightarrow \infty$ (i.e., $h \rightarrow 0$). This observation is useful precisely because $h = \pi/\Delta$ is the main parameter under user control. In particular, to bound the approximation error above by $\epsilon_a > 0$, one need only select Δ satisfying $\eta(\Delta, x) < \epsilon_a$. This requires specifying d^* and computing $\mathcal{N}(g_x, D_{d^*})$, problems dealt with in the next section.

REMARK B: The assumptions that X has a probability density and MGF are easily dealt with. The key assumption is (6), which arises naturally. In particular, it is not difficult to show that

$$\frac{M(\beta + is)}{M(\beta)} = \int_{-\infty}^{\infty} e^{isu} \frac{e^{\beta u}}{M(\beta)} F'(u) du,$$

demonstrating that (6) requires the characteristic function of the exponentially tilted density

$$q_\beta(u) = \frac{e^{\beta u} F'(u)}{M(\beta)}$$

to decay to zero at least algebraically. This relatively weak condition is guaranteed to hold with $r = 1$, for example, provided $F''(\cdot)$ exists and is integrable; see also Lemma 6.2.1 of Bleistein and Handelsman (1975). However, (6) also holds under much weaker conditions; see, for example, Theorem 126 of Titchmarsh (1948), Lemma 12.3 of Olver (1974) or Theorem 1 of Wong (1989, §IV.2). Finally, it is noted that (6) can be guaranteed by considering instead the CDF of the convolution $X + Z_\sigma$, where $Z_\sigma \sim N(0, \sigma^2)$ and $X \perp Z_\sigma$. In this case, Theorem 3 provides an accurate approximation to the CDF of $X + Z_\sigma$; the accuracy of the latter as an approximation to the CDF of X depends primarily on σ and, to a lesser extent, $F(\cdot)$.

REMARK C: The results of Theorem 3 do not specifically depend upon the use of \bar{u}_x . In particular, Theorem 3 is easily shown to hold with \bar{u}_x replaced by any other constant c , provided

c is restricted to be interior to \mathcal{E} . This observation has some useful practical implications. Typically, \bar{u}_x approaches the boundary points of \mathcal{E} as the ordinate x approaches the limits of the support of X ; see, for example, Daniels (1954, §6). In this situation and when $\delta < \infty$ (i.e., \mathcal{E} is either semi-infinite or finite), it can happen that $d^* \rightarrow 0$. To see this, suppose $\bar{u}_x > 0$ (i.e., $x > E[X]$) and $\mathcal{E} = (-\infty, \delta)$ where $\delta < \infty$. Then, as $\bar{u}_x \rightarrow \delta$ (i.e., x approaches the upper bound of the support of X), it follows that $\min\{|\delta - \bar{u}_x|, |\bar{u}_x|\} \rightarrow 0$ and hence $d^* \rightarrow 0$. This is problematic because the error bound $\eta(\Delta, x)$ depends on d^* through the term $\mathcal{N}_x(g_x, D_{d^*}) \exp\{-2d^*\Delta\}$. If, for example, $\mathcal{N}_x(g_x, D_{d^*})$ changes slowly as $d^* \rightarrow 0$, Δ will typically need to be increased (i.e., h decreased) in order to maintain $\eta(\Delta, x)$ at a given level. Since decreasing h tends to increase the number of quadrature nodes required, this can adversely affect computation. Intuitively, bounding \bar{u}_x away from the finite boundary(ies) of \mathcal{E} should alleviate such difficulties. However, as the discussion in Section 4.1 will reveal, this is not quite sufficient; one must also be careful in fixing the value of d^* . One simple solution addressing both of these issues is provided in Lemma 1 of Section 4.1.

REMARK D: The bound of Theorem 3 relies only on the behavior of the MGF, which is a known quantity, and holds for general distribution functions. Using the Poisson summation formula, Abate and Whitt (1992, §5) establish a different bound on the approximation error of the trapezoidal rule. In case of a probability distribution supported on \mathbb{R} , this bound requires detailed understanding of the tail behavior of the cumulative distribution function (i.e., the quantity being computed). However, for distributions whose support is bounded on the left, Abate and Whitt (1992, 1995) show that one may bound the approximation error without such knowledge using “exponential damping.” Essentially, Abate and Whitt (1992, 1995) work with the exponentially damped probability density $e^{-as}F'(s)$, and show that an appropriate application of the Poisson summation formula leads to an exponential bound on the approximation error of the trapezoidal rule approximation to $P\{X \leq x\}$ that depends primarily on a and x . The resulting computational algorithm, referred to by these authors as “EULER,” then proceeds by fixing $h = \pi/(2x)$ and attempting to sum the resulting infinite series using Euler summation. It is easily shown that the “damping parameter” a is simply $-c$ in the parameterization (1). Consequently, in the case of probability distributions whose support is bounded on the left, there is an important distinction between the EULER algorithm and the algorithm being proposed here. In particular, instead of bounding the approximation error by selecting h for the fixed choice $c = \bar{u}_x$, Abate and Whitt (1992, 1995) bound the approximation error by

selecting c for the fixed choice $h = \pi/(2x)$.

In addition to approximation error, one must also consider truncation error; that is, the error incurred by truncating the infinite sum (7). This is important in the present setting because the integrand, hence tail of the summation (7), generally decays to zero at a rate governed by the number of integrable derivatives of F on \mathbb{R} (e.g., Feller, 1971, Lemma 4, p. 514). When the number of such derivatives is small, the resulting slow rate of decay may require the computation of a large number of terms in (7) in order to attain a desired level of accuracy. It is possible to place conservative bounds on the truncation error under minimal assumptions; see, for example, Abate and Whitt (1992, §6). These bounds typically yield inefficient quadrature schemes, and a more adaptive approach often proves beneficial. In Section 4.4 a hybrid approach is proposed that involves selecting an initial truncation point based on properties of the integrand. The initial approximation is then improved using convergence acceleration techniques.

4 Implementation

The next 3 sections describe one implementation of the theory of the previous section. It is assumed that an approximation to $P\{X > x\}$ to within $\epsilon_{tot} > 0$ is desired. Previous discussion implies $\epsilon_{tot} = \epsilon_t + \epsilon_a$, where $\epsilon_a > 0$ denotes approximation error and $\epsilon_t > 0$ denotes truncation error; here, it is assumed that $\epsilon_a = \epsilon_t = \frac{\epsilon_{tot}}{2}$. The methodology below is essentially guaranteed to control the level of approximation error. However, guaranteeing control of the truncation error ϵ_t is considerably more difficult, particularly if computational efficiency is of concern. One method for dealing with the problem of truncation error is discussed in Section 4.4. Ultimately, the proposed method is unable to provide an explicit guarantee; however, the results of Section 5 show that the method significantly exceeds the requested accuracy in all cases.

As discussed in Remark C, a modified version of \bar{u}_x may prove useful in handling certain numerical problems. Section 4.1 details a simple proposal that has been found to be empirically effective. Subsequent to making this modification, the essential steps for implementing the theory of Section 3.2 are then as follows:

1. Determine an approximation to $\mathcal{N}_x(g_x, D_{d^*})$.
2. Determine an appropriate choice of spacing $h = \frac{\pi}{\Delta}$.

3. Accurately compute (7).

The first two relate to the control of approximation error; the last primarily reflects the main practical difficulty, which is the control of truncation error. Sections 4.2 and 4.3 below contain the major supporting details for Steps 1 & 2; Section 4.4 contains the key details required for understanding how Step 3 will be handled. The main computational algorithm is then summarized in Section 4.5.

4.1 A modification of \bar{u}_x

The discussion following Theorem 3 suggests that the magnitude of d^* plays a significant role in determining Δ , hence h and possibly the number of quadrature nodes required for the accurate computation of (7). These difficulties may arise when the set \mathcal{E} is bounded above or below (or both). Since \mathcal{E} is dictated by the problem under consideration, alleviating the aforementioned difficulties evidently requires modification of \bar{u}_x , or equivalently, the path of integration; Lemma 1 provides one possible choice.

Lemma 1 *Let $\kappa = I\{\bar{u}_x > 0\}$. Then, Theorem 3 continues to hold with \bar{u}_x replaced by $\bar{\bar{u}}_x$ and $d^* = \frac{|\bar{\bar{u}}_x|}{2}$, where*

- for $\mathcal{E} = \mathbb{R}$, $\bar{\bar{u}}_x = \bar{u}_x$;
- for $\mathcal{E} = (-\infty, \delta)$, $\bar{\bar{u}}_x = \kappa \min\{\bar{u}_x, \frac{\delta}{2}\} + (1 - \kappa)\bar{u}_x$;
- for $\mathcal{E} = (-\delta, \infty)$, $\bar{\bar{u}}_x = (1 - \kappa) \max\{\bar{u}_x, -\frac{\delta}{2}\} + \kappa\bar{u}_x$;
- for \mathcal{E} a bounded interval, $\bar{\bar{u}}_x = (1 - \kappa) \max\{\bar{u}_x, -\frac{\delta}{2}\} + \kappa \min\{\bar{u}_x, \frac{\delta}{2}\}$.

Moreover, $\mathcal{N}_x(g_x, D_{d^*}) = \int_{-\infty}^{\infty} [e^{-xd^*} |g_x(t - id^*)| + e^{xd^*} |g_x(t + id^*)|] dt$.

The proof of Lemma 1 is identical to Theorem 3 and is therefore omitted. Notably, no modification to the path of integration is made when $\mathcal{E} = \mathbb{R}$ since \bar{u}_x is by construction bounded away from zero. In the remaining cases, $\bar{\bar{u}}_x$ is bounded away from any finite boundary of \mathcal{E} , preventing the upper bounds on d^* stated in Theorem 3 from converging to zero as x approaches the bounds of support.

With \bar{u}_x as defined, the various restrictions on d^* detailed in Theorem 3 can be further reduced to the single constraint $0 < d^* < |\bar{u}_x|$. The reason Lemma 1 sets d^* to be one half the distance to $|\bar{u}_x|$ is to avoid further numerical problems when computing $\mathcal{N}_x(g_x, D_{d^*})$. In particular, $\mathcal{N}_x(g_x, D_{d^*})$ can still grow without bound as $d^* \rightarrow |\bar{u}_x|$, resulting in problems similar to those discussed in Remark C. To see this, suppose $\mathcal{E} = (-\infty, \delta)$ for $\delta > 0$, and set $d^* = \bar{u}_x - \epsilon$. In Section 4.2, it is shown that $\mathcal{N}_x(g_x, D_{d^*})$ depends on $M(\bar{u}_x \pm d^*)$. Suppose $\bar{u}_x > \delta/2$; then, $\bar{u}_x = \delta/2$. Since $d^* = \bar{u}_x - \epsilon$, $\mathcal{N}_x(g_x, D_{d^*})$ thus depends on $M(\bar{u}_x + d^*) = M(\delta - \epsilon)$. If $\lim_{\epsilon \downarrow 0} M(\delta - \epsilon) = \infty$, setting d^* to be close to \bar{u}_x (i.e., ϵ close to zero) may result in an extremely large value of $\mathcal{N}_x(g_x, D_{d^*})$. In turn, a large Δ may be required in order to ensure $\eta(\Delta, x) < \epsilon_a$. The remainder of this paper employs \bar{u}_x in place of \bar{u}_x and $d^* = \frac{\bar{u}_x}{2}$.

4.2 Approximating $\mathcal{N}_x(g_x, D_{d^*})$

In order to select Δ to bound the approximation error of the trapezoidal rule, the constant

$$\mathcal{N}_x(g_x, D_{d^*}) = \int_{-\infty}^{\infty} \left[e^{-xd^*} |g_x(t - id^*)| + e^{xd^*} |g_x(t + id^*)| \right] dt$$

must be computed. However, since $\mathcal{N}_x(g_x, D_{d^*})$ appears as part of an upper bound on the exact approximation error, it need not be computed to a high level of precision. Below, it is shown that $\mathcal{N}_x(g_x, D_{d^*})$ may be easily approximated using a low-order quadrature rule.

Define for suitable α the function

$$a(t, \alpha) = \Re [K(\alpha + it) - K(\alpha)]. \quad (8)$$

Since $a(t, \alpha) = a(-t, \alpha)$ for $t \in \mathbb{R}$ and $a(0, \alpha) = 0$, straightforward algebra shows

$$\int_{-\infty}^{\infty} |g_x(t \pm id^*)| dt = 2 \frac{M(\bar{u}_x \mp d^*)}{M(\bar{u}_x)} \int_0^{\infty} \frac{e^{a(t, \bar{u}_x \mp d^*)}}{\sqrt{(\bar{u}_x \mp d^*)^2 + t^2}} dt. \quad (9)$$

These integrals exist provided $e^{a(t, \alpha)} = O(t^{-\psi})$ for some $\psi > 0$ as $t \rightarrow \infty$. This condition is equivalent to requiring that the characteristic function of X decay to zero at least algebraically, and is thus satisfied provided the conditions of Theorem 3 hold.

The integral on the right-hand side of (9) may be difficult to compute numerically if $\exp\{a(t, \bar{u}_x \mp d^*)\}$ decays slowly to zero as $t \rightarrow \infty$. One may alleviate such difficulties substantially by making

the change of variable $t = e^s$:

$$\int_0^\infty \frac{e^{a(t, \bar{u}_x \mp d^*)}}{\sqrt{(\bar{u}_x \mp d^*)^2 + t^2}} dt = \int_{-\infty}^\infty e^{a(e^s, \bar{u}_x \mp d^*)} \frac{e^s}{\sqrt{(\bar{u}_x \mp d^*)^2 + e^{2s}}} ds. \quad (10)$$

To see why, let α equal one of $\bar{u}_x \pm d^*$, and note that $\alpha \neq 0$ by the choice of d^* . Then, since $0 < e^{a(e^s, \alpha)} \leq 1$ for $s \in \mathbb{R}$, the integrand on the right-hand side of (10) is $O(e^{-|s|})$ as $s \rightarrow -\infty$. In addition, since $e^{a(t, \alpha)} = O(t^{-\psi})$ as $t \rightarrow \infty$, the integrand on the right-hand side of (10) is $O(e^{-\psi s})$ as $s \rightarrow \infty$. The right-hand side of (10) is now easier to integrate numerically since the tails of the integrand vanish exponentially.

In computing (10), the lower and upper tails of the integrand are handled differently. Define

$$I_1(\alpha) = \int_{-L_\alpha}^U e^{a(e^s, \alpha)} \frac{e^s}{\sqrt{\alpha^2 + e^{2s}}} ds$$

and

$$I_2(\alpha) = \int_U^\infty e^{a(e^s, \alpha)} \frac{e^s}{\sqrt{\alpha^2 + e^{2s}}} ds \equiv \int_0^{1/U} \frac{e^{a(e^{1/s}, \alpha) - 1/s}}{s^2 \sqrt{\alpha^2 + e^{2/s}}} ds,$$

where $U > 0$ is arbitrary and

$$L_\alpha = \max \left[0, \frac{1}{2} \log \left(\frac{1 - \tanh^2(\gamma)}{\alpha^2 \tanh^2(\gamma)} \right) \right] \quad (11)$$

for $0 < \gamma < 1$. The representation of $I_2(\alpha)$ as an integral on $[0, U^{-1}]$ is valid under the assumptions made thus far. It follows that

$$I(\alpha) = 2 \frac{M(\alpha)}{M(\bar{u}_x)} [I_1(\alpha) + I_2(\alpha)]$$

is guaranteed to approximate $\int_{-\infty}^\infty |g_x(t \pm id^*)| dt$ to within γ . This is easily seen upon noting

$$\int_{-\infty}^{-L} e^{a(e^s, \alpha)} \frac{e^s}{\sqrt{\alpha^2 + e^{2s}}} ds \leq \int_L^\infty \frac{e^{-s}}{\sqrt{\alpha^2 + e^{-2s}}} ds = \tanh^{-1} \left(\frac{1}{\sqrt{\alpha^2 e^{2L} + 1}} \right)$$

for $L > 0$. Hence, only the lower tail of (10) is truncated; this is done because successful truncation of the upper tail requires additional information about the decay rate of $e^{a(t, \alpha)}$ as $t \rightarrow \infty$. In summary:

Proposition 1 For $0 < \gamma < 1$,

$$\widehat{\mathcal{N}}_x = e^{-xd^*} I(\bar{u}_x + d^*) + e^{xd^*} I(\bar{u}_x - d^*) \quad (12)$$

approximates $\mathcal{N}(g_x, D_{d^*})$ to within 2γ .

Since $I_j(\alpha), j = 1, 2$ involve finite limits of integration and integrands that decay exponentially, (12) may be easily approximated using numerical quadrature. In the examples of Section 5, L_α is computed as in (11) with $\gamma = 0.01$. Then, with $U = 10$, the integrals $I_1(\bar{u}_x + d^*)$ and $I_1(\bar{u}_x - d^*)$ are computed using a simple Clenshaw-Curtis-type quadrature rule; see, for example, Krommer and Ueberhuber (1998, §5.2.6) or Press *et al.* (1989, §5.7). The selection $U = 10$ is made both to ensure that $I_1(\bar{u}_x \pm d^*)$ is the dominant contribution to $I(\alpha)$ and to fix the computation of $I_2(\bar{u}_x \pm d^*)$ over a relatively small interval. The integrals $I_2(\bar{u}_x + d^*)$ and $I_2(\bar{u}_x - d^*)$ are each computed with a 3 point Simpson rule, requiring a total of four additional function evaluations since $\lim_{s \rightarrow 0} s^{-2} e^{a(e^{1/s}, \bar{u}_x \pm d^*)} = 0$.

4.3 Bounding the appropriate choice of h

Define $\Delta_{min} = \max\{\frac{\log 2}{2d^*}, |x|\}$. Then, proceeding as in Theorem 3, Δ should be taken as the larger of Δ_{min} and the solution to

$$\frac{e^{\nu(\bar{u}_x)}}{\pi} \mathcal{N}_x(g_x, D_{d^*}) \exp\{-2\Delta d^*\} < \epsilon_a.$$

Substituting (12) in place of $\mathcal{N}_x(g_x, D_{d^*})$ in the previous inequality, one then finds that selecting

$$\Delta > \max\left(\Delta_{min}, \frac{1}{2d^*} \log\left(\frac{e^{\nu(\bar{u}_x)} \hat{\mathcal{N}}_x}{\pi \epsilon_a}\right)\right). \quad (13)$$

ensures that the approximation error is either close to or bounded above by ϵ_a . Since $h = \pi/\Delta$, selecting Δ in this way clearly places an upper bound on the choice of h . A specific choice of Δ , hence h , satisfying the required inequalities is proposed in Section 4.5.1.

4.4 Handling truncation error

Successful implementation of the trapezoidal rule requires computing (7) to within $\pm\epsilon_t$ of its actual value. For an integer $N \geq 2$, straightforward calculations yield

$$\frac{h}{2} \sum_{k=-\infty}^{\infty} g_x(kh) e^{-ixkh} = h \left[g_x(0) + \sum_{k=1}^{N-1} R_x(kh) \right] + E_N(h) \quad (14)$$

where $E_N(h) = h \sum_{k=N}^{\infty} R_x(kh)$ and

$$R_x(kh) = \frac{1}{2} \left(g_x(kh) e^{-ixkh} + g_x(-kh) e^{ixkh} \right)$$

is the real part of $g_x(kh)e^{-ixkh}$. Computing (14) to a guaranteed level of accuracy requires an upper bound on $|E_N(h)|$. Some limited insight here can be gained by rewriting $R_x(t)$ in a different form. Specifically let $a_x(t) = a(t, \bar{u}_x)$, the latter being defined as in (8), and set $b_x(t) = \Im(K(\bar{u}_x + it))$ to be its corresponding imaginary part. Then,

$$R_x(t) = \frac{e^{a_x(t)}}{\bar{u}_x^2 + t^2} \left[\bar{u}_x \cos(b_x(t) - xt) + t \sin(b_x(t) - xt) \right]. \quad (15)$$

Since $|\Re(z)| \leq |z|$ for $z \in \mathbb{C}$,

$$|R_x(t)| \leq \frac{e^{a_x(t)}}{\sqrt{\bar{u}_x^2 + t^2}}.$$

The right-hand side is typically monotone decreasing for $t > 0$; if so, the following bound on the truncation error is easily obtained:

$$|E_N(h)| \leq \int_{(N-1)h}^{\infty} \frac{e^{a_x(t)}}{\sqrt{\bar{u}_x^2 + t^2}} dt. \quad (16)$$

However, the utility of (16) is unclear. In part, this is because its computation in general requires numerical integration (i.e., unless asymptotics are available). More importantly, it is because this bound is typically very conservative. The conservatism occurs because the cancellation arising from oscillations in $R_x(t)$ about zero, captured by the bracketed term on the right hand side of (15), is being ignored.

As discussed in Abate and Whitt (1992, 1995), conservative bounds like (16) can be avoided by instead exploiting oscillatory behavior in the terms of the series (14). However, except under special circumstances, the potentially enormous computational gains that can be achieved typically come at the expense of accuracy guarantees. The basic approach to be taken here will involve fixing N to ensure that the terms in the series $E_N(h)$, properly grouped, exhibits approximately alternating behavior. The tail series $E_N(h)$ is then summed using a convergence acceleration method. The effective summation of (14) thus requires a better understanding of the pattern of oscillation in $R_x(t)$. Consider, then, $b_x(t) = \Im(\log M(\bar{u}_x + it))$. With $\xi_R(t)$ and $\xi_I(t)$ respectively denoting the real and imaginary parts of $M(\bar{u}_x + it)$,

$$b_x(t) = \arg(\xi_R(t) + i\xi_I(t)),$$

where $\arg z$ denotes the unique angle in $(-\pi, \pi]$ such that $z = |z|e^{i \arg(z)}$. For $z = x + iy$, this can be explicitly written as

$$\arg(x + iy) = \text{sign}(y) \left\{ \mathcal{H}(-x)\pi + [2\mathcal{H}(x) - 1] \arctan\left(\left|\frac{y}{x}\right|\right) \right\},$$

where (as before) $\mathcal{H}(w)$ respectively equals 0, $\frac{1}{2}$, or 1 if $w < 0$, $w = 0$, or $w > 0$. Consequently,

$$b_x(t) = \text{sign}(\xi_I(t)) \left\{ \mathcal{H}(-\xi_R(t))\pi + [2\mathcal{H}(\xi_R(t)) - 1] \arctan \left(\left| \frac{\xi_I(t)}{\xi_R(t)} \right| \right) \right\}. \quad (17)$$

The following result characterizes the behavior of $b_x(t)$ as $t \rightarrow \infty$ assuming certain smoothness and integrability conditions hold. This result is important for developing insight into the behavior of $R_x(t)$.

Proposition 2 *Let $\varpi_e(s) = \frac{1}{2} (e^{\bar{u}_x s} F'(s) + e^{-\bar{u}_x s} F'(-s))$ and $\varpi_o(s) = \frac{1}{2} (e^{\bar{u}_x s} F'(s) - e^{-\bar{u}_x s} F'(-s))$ respectively denote the even and odd parts of $\varpi_T(s) = e^{\bar{u}_x s} F'(s)$. As $s \downarrow 0$, suppose*

$$\varpi_e(s) \sim \frac{1}{2} \sum_{j=0}^{\infty} \gamma_{e,j} s^{j+\alpha_e-1} \quad \text{and} \quad \varpi_o(s) \sim \frac{1}{2} \sum_{j=0}^{\infty} \gamma_{o,j} s^{j+\alpha_o-1},$$

where $\alpha_k > 0$, $k = e, o$. Then, under further smoothness and integrability conditions on $\varpi_k(\cdot)$ and its derivatives, $b_x(t) \rightarrow \theta$ as $t \rightarrow \infty$, where $\theta \in (-\pi, \pi]$ is a constant.

The appearance of $\varpi_k(\cdot)$, $k = e, o$ in Proposition 2 stems from the direct dependence of $\xi_R(t)$ on $\varpi_e(\cdot)$ and $\xi_I(t)$ on $\varpi_o(\cdot)$. Implied in the expansions for $\varpi_k(s)$ as $s \downarrow 0$ is the assumption of either boundedness or unbounded integrability. For example, $\varpi_e(s)$ is bounded near $s = 0$ if $\alpha_e = 1$; $\varpi_o(s)$ is unbounded but integrable near $s = 0$ if $\alpha_o = 1/2$. Under these and other conditions on $\varpi_e(\cdot)$ and $\varpi_o(\cdot)$, $\xi_R(t)$ and $\xi_I(t)$ admit asymptotic expansions in t as $t \rightarrow \infty$, and these expansions may subsequently be used for studying the asymptotic behavior of $b_x(t)$.

Precise regularity conditions under which the stated expansions hold may be found e.g. in Wong (1989, §4.2). At present, the following loose interpretation of this result is sufficient: if $\xi_I(t)$ and $\xi_R(t)$ are eventually of constant sign and their ratio tends to a unique (possibly infinite) limit as $t \rightarrow \infty$, then $b_x(t)$ tends to a unique limit $\theta \in (-\pi, \pi]$. Subsequent discussion regarding the oscillation patterns in $R_x(t)$ will be restricted to this case, and covers a large and useful class of interesting distributions. For example, if $\varpi_T(s)$ has nonnegative support (i.e., the random variable $X > 0$ with probability 1), then $\varpi_e(s) = \varpi_o(s) = \frac{1}{2} \varpi_T(s)$. Existence of the MGF of X implies exponential decay of $F(x)$ and hence $F'(x)$ as $|x| \rightarrow \infty$. The existence of

$$\varpi_T(s) \sim \sum_{j=0}^{\infty} \gamma_j s^{j+\alpha-1}$$

for some $\alpha > 0$ as $s \downarrow 0$ is then sufficient for obtaining the stated result since the remaining regularity conditions can be shown to hold (cf. Wong, 1989, §IV.2, Thm. 1). However, this does

not exhaust the set of all possibilities, and one example of a distribution for which $b_x(t)$ has no limiting value is studied in Section 5.

Assuming $b_x(t) \rightarrow \theta$ as $t \rightarrow \infty$, (15) shows that the cases $x = 0$ and $x \neq 0$ should be considered separately. Specifically, with $x = 0$, the oscillation in $R_x(t)$ dies out as $t \rightarrow \infty$, and the efficient computation of (14) relies primarily on the decay rate of $t^{-1}e^{a_x(t)}$. Selecting h as in Section 4.3, one may increase N until the desired number of significant digits in (14) has been achieved. Alternatively, recognizing that (14) is simply a trapezoidal rule approximation to $\int_0^\infty R_x(t) dt$, another quadrature rule may instead be employed.

If $x \neq 0$, the oscillation in $R_x(t)$ does not die out as $t \rightarrow \infty$. Computation of (14) then becomes more difficult, and it is helpful to obtain further insight into the zero patterns of $R_x(t)$. For $t > 0$, the zeros of (15) arise as solutions to

$$b_x(t) - xt + k\pi = \arctan\left(-\frac{\bar{u}_x}{t}\right), \quad (18)$$

where k is integer valued. Suppose $b_x(t) - xt \sim \theta - xt$ as $t \rightarrow \infty$, where $\theta \in (-\pi, \pi]$. Since

$$\arctan\left(-\frac{\bar{u}_x}{t}\right) \sim -\frac{\bar{u}_x}{t}$$

as $t \rightarrow \infty$, the solutions of (18) are thus (approximately) given by the solutions to

$$\theta - xt + k\pi + \frac{\bar{u}_x}{t} = 0,$$

or equivalently, by

$$\hat{t}_k = \frac{\theta + k\pi + \sqrt{(\theta + k\pi)^2 + 4x\bar{u}_x}}{2|x|}$$

for k such that \hat{t}_k is real and positive. It is expected that \hat{t}_k will provide an accurate approximation for t such that $b_x(t)$ is roughly constant and $t > 2|\bar{u}_x|$; the latter condition stems from the observation that $|\arctan(u) - u| \leq 0.037$ for $|u| < 1/2$. Moreover, as $k \rightarrow \infty$,

$$\hat{t}_{k+1} - \hat{t}_k = \frac{\pi}{|x|} - \text{sign}(x)\frac{\bar{u}_x}{\pi k^2} + O(k^{-3}), \quad (19)$$

showing that the spacing between successive zeros settles down quickly.

This last observation is useful precisely because it implies $E_N(h)$, with its terms grouped properly, behaves like an alternating series for N sufficiently large. As mentioned earlier, such alternating behavior can potentially be exploited using so-called ‘‘convergence acceleration’’

methods. Two methods known to be effective for alternating series are Euler summation and the epsilon algorithm (e.g., see Brezinski and Redivo Zaglia, 1991). Abate and Whitt (1992) employ Euler summation to accelerate the calculation of a series similar to $E_N(h)$; as indicated earlier, the epsilon algorithm will instead be used here.

4.5 The main algorithm

The previous sections contain key details underlying the computational algorithm for

$$\widehat{P}\{X > x\} = \mathcal{H}(-\bar{u}_x) + \frac{e^{\nu(\bar{u}_x)}}{\pi} \left[hg_x(0) + h \sum_{k=1}^{N-1} R_x(kh) + E_N(h) \right], \quad (20)$$

where $R_x(\cdot)$ and $E_N(\cdot)$ are defined in Section 4.4. The main algorithm can now be presented. Section 4.5.1 describes how N and h are determined; Section 4.5.2 describes how the epsilon algorithm is used in computing the “tail” of the infinite series (20) i.e., $\pi^{-1}e^{\nu(\bar{u}_x)}E_N(h)$.

4.5.1 Selection of N and h

As discussed in Section 4.4, the initial truncation point N is determined such that the terms in the series $E_N(h)$, when properly grouped, begin to exhibit regular alternating behavior. This is evidently connected to the behavior of $R_x(t)$ for $t \geq Nh$; consequently, the selection of N and h are linked as is described below:

1. Determine $z_0 > 0$ such that (i) z_0 is a zero of $R_x(t)$; and, (ii) the zeros of $R_x(t)$ for $t \geq z_0$ are approximately equally spaced. The point z_0 is determined numerically by computing the successive zeros of $R_x(t)$ for $t > 2|\bar{u}_x|$; call this sequence of zeros $\{t_j, j \geq 1\}$. These computations are stopped once $t_{K+1}/t_K - 1 < 0.1$ for some $K \geq 1$, with $z_0 = t_K$.
2. As shown in (19), the predicted asymptotic spacing between successive zeros is $\varphi_0 = \pi/|x|$. If $|t_{K+1} - t_K - \varphi_0| < 0.1\varphi_0$, the zeros of $R_x(t)$ are assumed to follow the asymptotic behavior described in Section 4.4, and in particular the spacing between successive zeros is set to $\varphi = \varphi_0$. However, if $t_{K+1}/t_K - 1 < 0.1$ and $|t_{K+1} - t_K - \varphi_0| \geq 0.1\varphi_0$, this indicates (19) may be incorrect. In this case, the spacing is initially set to $\varphi_1 = t_{K+1} - t_K$. The function $R_x(t)$ is then replaced by a Chebyshev polynomial expansion on the interval $[z_0, z_0 + \varphi_1]$ (see e.g. Press *et al.*, 1989, §5.6). All roots of this polynomial on this interval

are determined, and the spacing φ is taken to be the distance between z_0 and the closest root of the approximating polynomial that exceeds z_0 .

3. Compute $e_0 = z_0 + \frac{\varphi}{2}$, the first (approximate) extrema subsequent to z_0 .

4. Set $N = \lfloor \frac{e_0 \Delta_0}{\pi} \rfloor + 1$, $\Delta_N = N\pi/e_0$, and $h = \pi/\Delta_N$, where

$$\Delta_0 = \max \left(\Delta_{min}, \frac{1}{2d^*} \log \left(\frac{e^{\nu(\bar{u}_x)} \widehat{\mathcal{N}}_x}{\pi \epsilon_a} \right) \right).$$

Steps 1 and 2 can be computationally demanding; however, these can also be avoided if the user provides z_0 and φ , computations requiring some further asymptotic analysis. Step 4 ensures $Nh = e_0$ and that Δ_N satisfies (13). Implicit in these selections is the assertion that zeros of $R_x(t)$ are approximately equally spaced for $t > e_0$; if true, the terms in $E_N(h)$ may then be grouped to produce an approximately alternating series. The rationale for selecting N such that Nh coincides with an extrema (i.e., e_0) instead of a zero (i.e., z_0) is based on recommendations made in Sauter (2000). With these choices, an initial approximation P_\star to (20) is:

$$P_\star = \mathcal{H}(-\bar{u}_x) + \frac{e^{\nu(\bar{u}_x)}}{\Delta_N} \left[\frac{1}{\bar{u}_x} + \sum_{k=1}^{N-1} R_x \left(\frac{\pi k}{\Delta_N} \right) \right]. \quad (21)$$

4.5.2 Computing $\widehat{P}\{X > x\}$

The computation of (20) requires both P_\star and $E_N\left(\frac{\pi}{\Delta_N}\right)$; the former is easily computed via (21), so the latter is now considered. Let $\alpha = \langle \frac{\varphi}{h} \rangle$, the operation $\langle s \rangle$ denoting the nearest integer to s . Then,

$$E_N \left(\frac{\pi}{\Delta_N} \right) = \frac{\pi}{\Delta_N} \sum_{j=0}^{\infty} \beta_j$$

with

$$\beta_j = \sum_{k=N+j\alpha}^{N+(j+1)\alpha-1} R_x \left(\frac{k\pi}{\Delta_N} \right). \quad (22)$$

Because $Nh = e_0$ is a point of extrema and αh approximates the period of $R_x(t)$, each β_j is essentially a trapezoidal rule approximation to the integral of $R_x(t)$ between successive extrema. Provided $|R_x(t)|$ decays monotonically and the zeros of $R_x(t)$ are approximately equi-spaced for $t > Nh$, the β_j should be small in magnitude and alternate in sign.

Now, for $L \geq 0$, define the sequence

$$P_L = P_\star + \frac{e^{\nu(\bar{u}_x)}}{\Delta_N} \sum_{j=0}^L \beta_j.$$

Notice that (1) P_∞ equals (20) for a specific choice of h and N ; (2) P_L corresponds to a partial sum of the desired infinite series; and, (3) P_L is a linear transformation of $\sum_{j=0}^L \beta_j$. Hence, the sequence of partial sums $\{P_L, L \geq 0\}$ should oscillate about the desired limit P_∞ . If so, the rate at which P_L converges to P_∞ can potentially be accelerated. The epsilon algorithm has been found to be effective when applied to sequences whose terms oscillate about a finite limit; see, for example, Smith and Ford (1982) and more recently Sauter (2000). For readers unfamiliar with this algorithm, a brief introduction is provided in Appendix B; see, for example, Wimp (1981), Brezinski and Redivo Zaglia (1991), or Sauter (2000).

Let ε_L denote the result of applying the epsilon algorithm to a given sequence $\{P_0 \dots P_L\}$, where $L \geq 0$. The exclusion of P_\star as a member of this sequence is deliberate; the intent here is to accelerate the computation of the *tail* series $E_N(\pi/\Delta_N)$, or equivalently, P_∞ . The computations for (20) are thus finished adaptively as follows:

1. Set $L = 0$, compute $\varepsilon_0 = P_0$ and initialize error_{old} to be a large number.
2. Increment L by 2, compute P_{L-1} and P_L .
3. Compute ε_L and $\text{error}_{new} = |\varepsilon_L - \varepsilon_{L-2}|/\varepsilon_L$.
4. If $\frac{1}{3}\text{error}_{old} + \frac{2}{3}\text{error}_{new} < \epsilon_t/1000$, take ε_L as the approximation and stop iteration; otherwise, set $\text{error}_{old} = \text{error}_{new}$ and return to step 3.

The reason for increasing the sequence length by two each time is connected to the difference in behavior of the epsilon algorithm when applied to sequences of even versus odd length; see Appendix B for discussion. The stopping criterion for checking convergence of the computed answer uses relative error and is set significantly smaller than the requested accuracy $\epsilon_t = \frac{1}{2}\epsilon_{tot}$. This is to help ensure that the final approximation matches the desired answer to (at least) the specified level of accuracy. The use of an unequally weighted average of the relative error across two successive iterations is used to guard against being fooled by locally small changes in the computed answer. The most recent iteration receives a larger weight since it is based on more quadrature points.

5 Examples

The six examples below have been chosen to illustrate the performance of the proposed method across a reasonably diverse spectrum of problems. The computations were carried out in MAPLE 6, with floating point precision set at 20 digits (i.e., `Digits:=20`; Waterloo Maple Inc., 2000). A copy of the MAPLE code used for these examples is available upon request.

All computations are carried out using $\epsilon_{tot} = 10^{-8}$. In addition to the CGF $K(t) = \log M(t)$, the only user input required is the ordinate x and the set \mathcal{E} for which $M(t)$, $t \in \mathcal{E}$ is finite. The computation of \bar{u}_x requires $K'(\cdot)$; numerical differentiation is used here. In the tables, the absolute error of the approximations are reported. Also summarized are the values of h , d^* , $\hat{\mathcal{N}}_x$, and the number of points at which the CGF $K(\cdot)$ has been evaluated. This reported number of function evaluations reflects: (1) the computations involved in obtaining $\hat{\mathcal{N}}_x$; and, (2) the quadrature nodes computed in Sections 4.5.1 and 4.5.2. However, it does not reflect those computations needed for determining \bar{u}_x , z_0 and ψ (see Sections 2.2 and 4.5.1). These computations rely on a built-in root finding subroutine, and the author was unable to obtain this information from MAPLE. However, as noted earlier, the numerical computation of z_0 and ψ may also be viewed as optional since it absolves the user from doing the requisite asymptotic analysis. Consequently the reported number of function evaluations may be viewed as a measure of the efficiency of ε_L , given z_0 , ψ and \bar{u}_x .

For comparison, also reported are the absolute errors of the “straight trapezoidal rule” (i.e., unaccelerated) approximation P_L and

$$\text{Euler}(x) = P_\star + \frac{e^{\nu(\bar{u}_x)}}{\Delta_N 2^{L+1}} \sum_{k=0}^{L+1} \binom{L+1}{k} \left[\sum_{j=0}^k \beta_j \right] \equiv P_\star + \frac{e^{\nu(\bar{u}_x)}}{\Delta_N} \sum_{j=0}^{L+1} w_j \beta_j$$

where $w_j = P\{W \geq j\}$ for $W \sim \text{Binomial}(L+1, 1/2)$. The approximation $\text{Euler}(x)$ is obtained by adding to the initial approximation P_\star the (truncated) Euler sum of the series $\sum_{j=0}^{\infty} \beta_j$. Euler summation is expected to be effective since β_j and β_{j+1} should have opposite signs and be decreasing in magnitude. Abate and Whitt (1992) provide a useful discussion on the application of Euler summation in similar problems. In the tables that follow, the approximations ε_L and P_L are respectively referred to as $\text{Epsilon}(x)$ and $\text{Straight}(x)$. Importantly, each of these approximations is ultimately based on the exact same set of quadrature nodes, providing a balanced assessment of both accuracy and the effect of convergence acceleration. Since one

motivation of this paper is to provide a useful alternative to saddlepoint methods, the absolute error of the Lugannani-Rice approximation

$$\text{LR}(x) = 1 - \Phi(\hat{\omega}_x) + \phi(\hat{\omega}_x) [\hat{z}_x^{-1} - \hat{\omega}_x^{-1}]$$

to $P\{X > x\}$ is also provided. Here, $\hat{z}_x = \text{sign}(\hat{u}_x)\sqrt{K''(\hat{u}_x)}$, $\hat{\omega}_x = \text{sign}(\hat{u}_x)\sqrt{2[x\hat{u}_x - K(\hat{u}_x)]}$, and $K'(\hat{u}_x) = x$; see, for example, Kolassa (1997, Chapter 5). Finally, to a large extent, the results in Tables 1-6 speak for themselves; hence, comment is reserved until Section 5.5, where some reasonably general observations can be made.

5.1 Mixtures of independent noncentral χ^2 random variables

Let $\chi_p^2(\omega^2)$ denote a chi-squared random variable with p degrees of freedom and noncentrality parameter ω^2 . The problems considered in this section deal with the mixture $X = \sum_{j=1}^n \lambda_j Y_j$, where $Y_j \sim \chi_{p_j}^2(\omega_j^2)$, $Y_j \perp Y_k$ for all $j \neq k$, $p_j > 0$ are integers, $\omega_j^2 \geq 0$, and $\lambda_j \in \mathbb{R}$ is nonzero. The MGF of X takes the form

$$M_X(t) = \prod_{j=1}^n G(\lambda_j t; p_j, \omega_j^2)$$

where $G(u; \nu, \eta^2) = (1 - 2u)^{-\nu/2} \exp\{\eta^2 u(1 - 2u)^{-1}\}$; see, for example, Johnson, Kotz, and Balakrishnan (1995, Ch. 29). Importantly, the set \mathcal{E} on which $M_X(t) < \infty$ is governed by the sign and magnitude of $\lambda_1 \dots \lambda_n$. As noted in Davies (1980), any quadratic form (or ratio thereof) in independent normal random variables can be reduced to the form $X + \sigma_0 Z$, where $Z \sim N(0, 1)$, $Z \perp X$, and $\sigma_0 \geq 0$. Hence this section covers a significant class of interesting problems. Various methods of numerical inversion of the characteristic function for this problem have been considered previously; see, for example, Imhof (1961), Davies (1973, 1980), Rice (1980), and Helstrom (1983).

5.1.1 An unweighted mixture

Consider $X = Y_1 + Y_2$, where $Y_1 \sim \chi_2^2(0.1)$ and $Y_2 \sim \chi_5^2(0.9)$. In this case, $X \sim \chi_7^2(1)$, and

$$P\{X > x\} = \frac{e^{-1/2}}{2^{7/2}\Gamma(7/2)} \int_x^\infty \left(s^{7/2} - 1\right) e^{-s/2} {}_0F_1\left(\frac{7}{2}, \frac{7s}{4}\right) ds,$$

with ${}_0F_1$ denoting a generalized hypergeometric series. The mean and standard deviation of X are respectively 8 and 4.14. For the purposes of determining the approximation error of the various quadrature rules, the “exact” CDF is computed using MAPLE’s built-in numerical integration routine to a requested accuracy of 20 digits. The results are reported in Table 1.

5.1.2 A weighted mixture supported on $(0, \infty)$

Rice (1980) and Helstrom (1983) consider a complicated random variable arising in radar detection problems. Specifically, these authors require the CDF of $X = \sum_{j=1}^{25} \lambda_j Y_j$, where $Y_j \sim \chi_2^2(0.4)$ and $\lambda_j = 2 \left(1 + \cos \left[\frac{j\pi}{26}\right]\right)$. Notice that $\lambda_j \in (0, 4)$ for $j = 1 \dots 25$, and thus $X > 0$ with probability 1. The mean and standard deviation of X are respectively 120 and 28.79. Rice and Helstrom respectively compute (1) using the trapezoidal rule; however, Rice selects $c \approx \hat{u}_x$ and Helstrom selects $c \approx \bar{u}_x$ (see Section 2.2). Their proposed implementations are otherwise ad-hoc; for example, in both cases the spacing h is selected by repeatedly halving it until convergence in the computed answer is achieved to the desired number of significant digits. Moreover, no attempt is made to select a point of truncation and convergence acceleration is not used to help control truncation error. The results of applying the present method to this problem may be found in Table 2. For the purposes of determining the approximation error of the various quadrature rules, the “exact” CDF is computed numerically. Specifically, with $u_r^* = \bar{u}_x + \frac{1+i\sqrt{3}}{x} \exp(r - e^{-r})$, the integral

$$I(x) = \frac{1 + i\sqrt{3}}{x\pi i} \int_{-\infty}^{\infty} [u_r^*]^{-1} \exp(-u_r^* x + K(u_r^*) + r - e^{-r}) (1 + e^{-r}) dr$$

is computed using MAPLE’s built-in numerical integration routine to a requested accuracy of 20 digits. The desired probability is then given by $P\{X > x\} = \mathcal{H}(-\bar{u}_x) + \Re\{I(x)\}$; see Rice (1980, §5) for further details, including a discussion of the value of this particular representation for numerical calculations.

5.1.3 A weighted mixture supported on \mathbb{R}

Davies (1980, Table 3) considers computing the CDF for

$$X = 7\chi_6^2(6) + 3\chi_2^2(2) - 7\chi_1^2(6) - 3\chi_1^2(2).$$

The support of X is \mathbb{R} ; the mean and standard deviation are respectively 38 and 56.88. The results of applying the present method to this problem may be found in Table 3. In this case

a tractable formula for the CDF of X is not available. The answers obtained from the current procedure and from the FORTRAN algorithm AS155 of Davies (1980) are thus compared, the latter being specifically designed to for this class of problems. Computations were carried out in double precision to a requested accuracy of 10^{-12} .

5.2 Time dependent mean of regulated Brownian motion

Let $B(t)$ denote standard Brownian motion. Then, regulated Brownian motion with drift -1 and diffusion coefficient 1 may for $t \geq 0$ be defined as $B^*(t) = B(t) - t - \min_{s \in [0,t]} \{B(s) - s\}$. It follows that $B^*(t) \geq 0$; moreover, for $t > 0$ (cf. Abate and Whitt, 1987)

$$E[B^*(t)|B^*(0) = 0] = \frac{1}{2} - (t+1)[1 - \Phi(t)] + \sqrt{t} \phi(t), \quad (23)$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ respectively denote the standard normal CDF and density functions. The process $B^*(t)$ is useful for modeling stochastic flow systems arising in queuing theory. Notably, the time dependent mean (23) corresponds to the CDF of a continuous random variable X supported on $(0, \infty)$; call this CDF $F(\cdot)$. The corresponding density function $F'(t) = O(t^{-1/2})$ as $t \downarrow 0$; however, (6) can still be shown to hold with $r = 1/2$. The mean and standard deviation of X are respectively $1/2$ and 0.87 . Abate and Whitt (1992) use this example to demonstrate the performance of various transform inversion methods in a case where the MGF

$$M(t) = \frac{2}{1 + (1 - 2t)^{1/2}}$$

decays slowly. The computational problem is potentially challenging because the terms in (14) decay to zero slowly, and a large number of terms may be required in order to prevent significant truncation error. Essentially, the results of Abate and Whitt for this problem may be summarized as follows: a combination of summation with Euler acceleration is significantly more effective than attempting to sum (14) directly. The results of applying the present method to this problem may be found in Table 4; the exact CDF is computed via (23).

5.3 A non-regular exponential family

Jensen (1995, p. 48) considers the problem of developing saddlepoint approximations for densities of the form

$$f(x) = \frac{x^{-\lambda} e^{-x}}{\Gamma(1 - \lambda, 1)} I\{x > 1\}$$

where $1 < \lambda < 2$ and $\Gamma(a, b) = \int_b^\infty t^{a-1} e^{-t} dt$ denotes the incomplete gamma function. This class of distributions contains the inverse Gaussian distribution ($\lambda = 3/2$) as a special case. Indeed, as shown in Jensen (1995), the choice $\lambda = 3/2$ is quite special in this regard. In particular, the standard saddlepoint approximation to the density function for $\lambda = 3/2$ is exact, but is expected to fare poorly in the extreme tails for any other choice of $\lambda \in (1, 2)$.

Consider the random variable X defined by setting $\lambda = 5/4$, for which

$$M(t) = (1-t)^4 \frac{\Gamma(-\frac{1}{4}, 1-t)}{\Gamma(-\frac{1}{4}, 1)}, \quad t < 1.$$

The mean and standard deviation of X is approximately 1.62 and 0.68 respectively. Perhaps the most interesting aspect of this problem is that, unlike the previous examples, $b_x(t)$ does not tend to a unique limit as $t \rightarrow \infty$. In particular, as $t \rightarrow \infty$, it can be shown that

$$\xi_R(t) \sim -\beta \sin(t) t^{-1} \quad \text{and} \quad \xi_I(t) \sim \beta \cos(t) t^{-1};$$

hence, by (17),

$$b_x(t) \sim \text{sign}(\cos(t)) \{ \mathcal{H}(\sin(t))\pi + [2\mathcal{H}(-\sin(t)) - 1] \arctan(|\cot(t)|) \}.$$

The function $b_x(t)$ is both periodic and discontinuous, exhibiting a sawtooth pattern. Consequently, (19) is no longer valid; however, the zeros of $R_x(t)$ nevertheless exhibit a regular pattern, and hence the proposed algorithm is likely to be effective provided that the associated spacing can be identified. Step 2 of the algorithm described in Section 4.5.1 attempts to do so, and the spacing so determined is used in obtaining the results reported in Table 5. Table 6 compares the spacing predicted by (19) and that determined by the algorithm, and shows that in general the latter exceeds the former. For the purposes of determining approximation error, the “exact” CDF is computed by integrating the density function using MAPLE’s built-in numerical integration routine to a requested accuracy of 20 digits.

5.4 Compound Polya sum

Jensen (1995, §7.4) develops saddlepoint approximations for the distribution function of $X = \sum_{j=1}^N Y_j$, where $N \sim \text{Negative Binomial}(r, p)$ for $r > 0$ and $p \in (0, 1)$, and $\{Y_j, \geq 1\}$ is a sequence of continuous, mutually independent random variables. Random variables like X arise in computing claims distribution in actuarial science. Jensen shows that a saddlepoint

approximation to $P\{X > x\}$ with small relative error should depend on the gamma, not normal, distribution (e.g., see $\text{LR}(x)$ above), suggesting a failure of the usual central limit theory. However, from the point of view of the methods developed in this paper, such nonstandard behavior is not particularly relevant. More relevant is the behavior of the characteristic function of X . In particular, since $P(N = 0) = (1 - p)^r$, the distribution of X is discontinuous at $x = 0$, and thus $M_X(it)$ will not decay to zero as $|t| \rightarrow \infty$. In this case,

$$M_X(t) = \left(1 - \frac{p}{1-p} [M_{Y_1}(t) - 1]\right)^{-r},$$

where $M_{Y_1}(t)$ is the MGF of Y_1 . For Y_1 continuous, $M_{Y_1}(it) \rightarrow 0$ as $|t| \rightarrow \infty$ by the Riemann-Lebesgue Lemma. It follows that $M_X(it) \rightarrow (1 - \beta)^r$ as $|t| \rightarrow \infty$, violating (6) and hence the regularity conditions required for Theorem 3.

It is shown in Jensen (1995) that one may compute the desired tail probability using the fact that, for $x > 0$,

$$P(X > x) = P(X > x | N > 0)P(N > 0) \equiv (1 - p)^r P(\tilde{X} > x)$$

where $\tilde{X} = \sum_{j=1}^{\tilde{N}} Y_j$ and \tilde{N} is a zero-truncated negative binomial random variable. The random variable \tilde{X} is continuous with MGF

$$M(t) = \frac{M_X(t) - (1 - p)^r}{1 - (1 - p)^r},$$

$M(t)$ being defined for $t \in \mathbb{R}$ such that $M(t) > 0$ and $M(t) < \infty$. The methods of this paper may be used to compute the CDF of \tilde{X} and hence of X provided the density of \tilde{X} satisfies (6).

For illustration, suppose $Y_i \sim \text{Exponential}(1)$, $r = 3$, and $p = 1/4$. The exact distribution of X is defined by $P\{X = 0\} = \left(\frac{3}{4}\right)^3 = \frac{27}{64}$ and

$$P\{X > x\} = \frac{27}{64} e^{-x} \sum_{n=1}^{\infty} \left(\sum_{k=0}^{n-1} \frac{x^k}{k!} \right) \binom{3+n-1}{n} \frac{1}{4^n} \quad (24)$$

for $x > 0$. The mean and standard deviation of X are respectively 1 and 1.82. Straightforward calculations show that the MGF of \tilde{X} may be simplified to

$$M(t) = \frac{1}{37} \left(1728 \left[\frac{t-1}{4t-3} \right]^3 - 27 \right),$$

with $\mathcal{E} = (-\infty, 3/4)$. For c interior to \mathcal{E} , it is straightforward to show that $|M(c+it)| \sim \frac{81}{148} |t|^{-1}$ as $|t| \rightarrow \infty$, and hence that (6) is met. The “exact” CDF in this example is computed by truncating the outer summation in (24) at $n = 100$; the results for the above parameter choices are reported in Table 7.

5.5 Summary of results

In general, the tables show that $\text{Epsilon}(x)$ significantly outperforms both $\text{Straight}(x)$ and $\text{Euler}(x)$ in terms of absolute accuracy. It is also evident that the epsilon algorithm substantially improves upon Euler summation as a method for accelerating the computation of the infinite series P_∞ . Perhaps the greatest benefits of using the epsilon algorithm occur in Tables 4 and 5. In these tables, $R_x(t)$ decays comparatively slowly, and $\text{Straight}(x)$ provides a comparatively poor approximation; in contrast, $\text{Euler}(x)$ roughly doubles and $\text{Epsilon}(x)$ roughly triples the number of correct significant digits. Tables 1 and 7 also display impressive gains. The tables further show the following general trends:

- h , d^* and \widehat{N}_x can vary significantly with the ordinate x ;
- the number of quadrature nodes required is typically greatest for ordinates x close to $E[X]$, and decays in the tails;
- $\text{Epsilon}(x)$ meets or significantly exceeds the requested accuracy of $\epsilon_{tot} = 10^{-8}$ in all cases;
- $\text{LR}(x)$ provides an erratic and comparatively poor approximation to $P\{X > x\}$, though tends to improve for more extreme values of x .

In contrast to the others, Table 3 demonstrates a remarkable level of consistency across all of the quadrature-based methods for all ordinates. These results could be interpreted as saying that each method is providing the same number of significant digits and does so to the requested level of accuracy. However, there is also ample evidence in the remaining tables to suggest that $\text{Epsilon}(x)$ achieve significantly greater accuracy than either $\text{Straight}(x)$ or $\text{Euler}(x)$. Thus, a reasonable alternative explanation is that the exact CDF, which lacks a computable form and is approximated using the double precision FORTRAN algorithm AS155 of Davies (1980), is comparatively inaccurate. Further experimentation with various levels of requested accuracy equal to or smaller than 10^{-8} indicate that the answers produced by Algorithm AS155 are probably only trustworthy to 10^{-8} or so, regardless of requested accuracy. This is consistent with the fact that Algorithm AS155 relies on double precision floating point computations. Since MAPLE is able to carry out its computations with effectively arbitrary precision, a more likely explanation is that the answers obtained by $\text{Epsilon}(x)$ are significantly more accurate than Table 3 suggests.

6 Discussion

The trapezoidal rule is well known and has been proposed as way to integrate the tail probability inversion integral in a number of previous papers. However, to the authors knowledge, this is the first paper to use the results of Stenger (1993) for this particular purpose. Theorem 3 shows that the conditions required on the distribution of the random variable X are actually quite weak, the most restrictive being perhaps existence of the MGF. From the perspective of the available literature on sinc quadrature rules (see e.g. Lund & Bowers, 1992; Stenger, 1993), a novel feature of the implementation here is the attention paid to the computation of the order constant (i.e., $\mathcal{N}_x(g_x, D_{d^*})$). The proposals made in Section 4.2 for approximating $\mathcal{N}_x(g_x, D_{d^*})$ are new and facilitate choosing h with confidence. As stated, Proposition 1 is restricted to the specific problem at hand. However, the essential assumptions on the integrand(s) are boundedness at the origin and algebraic decay, and thus the basic ideas apply more generally.

The examples demonstrate that proposed methodology works very well. The results support the conclusions of Abate and Whitt (1992) that numerical transform inversion can be done accurately and (relatively) easily. The current algorithm represents an improvement over existing technology for univariate transform inversion in the sense that the same basic algorithm can be applied without regard to the support of the distribution of X . However, many other choices exist, especially for the inversion of one-sided Laplace transforms; see Abate and Whitt (1992, §15) for a review. Most of these algorithms are also derived from the trapezoidal rule. The previously mentioned algorithm EULER of Abate and Whitt (1992, 1995; see also Remark D) is one elegant example. EULER is surprisingly easy to code and could have been used for all of the examples except that in Section 5.1.3. However, the simplicity of EULER is due at least in part to its ignorance of certain considerations that are addressed in detail here. In particular, EULER requires the specification of two parameters n and m whose roles are analogous to those of N and L ; the former determines at what point in the series acceleration is begun, the latter determines how many terms are used in actually approximating the “tail”. Abate and Whitt (1995) remark that “we typically use $m = 11$ and $n = 15$, increasing n as necessary.” The selection $m = 11$ makes the tacit assumption that the tail series being summed is actually alternating. Since EULER fixes $h = \pi/(2x)$, knowing whether or not this is the case is equivalent to the problem of determining n such that the sign pattern of their integrand has indeed become “regular”; see Section 4.5.1. The basic EULER algorithm expends no effort to

determine whether this is indeed the case.

Practically speaking, tail probabilities are rarely needed to the level of accuracy observed in the tables. Indeed, $\text{LR}(x)$ provides an answer to what may be regarded as an adequate level of precision in most cases. However, Table 4 demonstrates that this is not always the case. Because trustworthy guidelines for successful application of these methods are lacking, one must therefore regard such evidence as anecdotal. In fact, little in the way of informative error bounds exist for saddlepoint approximations; this is particularly true in the case of the Lugannani-Rice approximation. In contrast, the theory presented earlier provides a computable bound on the approximation error associated with the series (7), and in principle the tail probability can be obtained to as many significant digits as desired by summing enough terms. The tables confirm this and show in particular that it is possible to do significantly better than $\text{LR}(x)$ with modest additional computational effort. For the examples considered, $\text{Epsilon}(x)$ significantly increased accuracy compared to either $\text{Straight}(x)$ or $\text{Euler}(x)$. These results are in line with the observations of Sauter (2000), and further investigation into why this is the case would be worthwhile. One difficulty here is that the analysis of the epsilon algorithm, which is a nonlinear sequence transformation, is significantly more difficult than it is for Euler summation (i.e., a linear transformation). Results on the acceleration properties of the epsilon algorithm in some important special cases can be found in Brezinski and Redivo Zaglia (1991, §2.3).

The proposed methodology only requires specifying the cumulant (or moment) generating function of a univariate random variable. This covers a large number of interesting situations, including the computation of conditional distribution functions. For example, the algorithm described here may in principle be applied in problems for which Skovgaard's conditional tail probability approximation can be used (Skovgaard, 1987). However, in order to do so, the corresponding conditional cumulative generating function is needed, and this presents a significant drawback (cf. Kolassa, 1997, §7.2). Put another way, the computations required would parallel those needed for a "single", rather than "double", saddlepoint approximation to the desired conditional probability. Successfully circumventing this problem may require extending Theorem 2 to multivariate transforms; for related work in this direction, see Choudhury, Lucantoni, and Whitt (1994).

The focus of this paper has been on the case in which X is absolutely continuous. Computations for lattice-valued random variables are significantly easier. For example, the inversion integral

(1) can be expressed over a finite range (cf. Kolassa, 1997, §2.7). Consequently the primary challenge is in obtaining a practically useful bound on the approximation error; the difficulties dealt with Section 4.4 do not arise. One approach for handling lattice-valued random variables is described in Abate and Whitt (1992); another can be derived using Stenger's results. However, computations in more general problems, such as when

- X has a density but not a MGF (e.g., Cauchy random variables);
- X is discrete but not lattice-valued;
- the CDF of X is continuous but not everywhere differentiable;

present substantially greater challenges for “Fourier series” methods. The main reasons for this rests in the rate of decay of the Fourier transform and (at points of discontinuity) the so-called Gibbs Phenomenon. Smoothing at some level will usually be required in order to successfully cope with such problems.

Appendix A: Proofs

Proof of Theorem 3

For $h = \pi/\Delta$, the conditions on Δ becomes equivalent to the stated conditions on h in Theorem 2. Also, from (4), observe that

$$\left| P\{X > x\} - \left[\mathcal{H}(-\bar{u}_x) + \frac{e^{\nu(\bar{u}_x)}}{\pi} \mathcal{T}_x(h) \right] \right| = \frac{e^{\nu(\bar{u}_x)}}{2\pi} \left| \int_{-\infty}^{\infty} g_x(t) e^{-ixt} dt - \sum_{k=-\infty}^{\infty} g_x(kh) e^{-ixkh} \right|.$$

Hence, the result follows directly from Theorem 2 provided that conditions (i)-(iii) are satisfied under the stated conditions.

Let $z = s + iy \in \mathbb{C}$, and let $D_a \subset \mathbb{C}$ denote the infinite strip of width $a > 0$ containing the real axis. We must first establish that conditions (i)-(iii) of Theorem 2 hold for $f = g_x$. To show (i), recall first that

$$g_x(z) = \frac{\exp\{K(\bar{u}_x - y + is) - K(\bar{u}_x)\}}{\bar{u}_x - y + is}.$$

Since the composition of two analytic functions is analytic in the intersection of their respective regions of analyticity, it suffices to separately consider the behavior of numerator and denominator.

Since z^{-1} is analytic except at the origin, $(\bar{u}_x - y + is)^{-1}$ is analytic in $\mathbb{C} - \{s = 0, y = \bar{u}_x\}$. Thus, for $y \in (-|\bar{u}_x|, |\bar{u}_x|)$, it follows that $(\bar{u}_x - y + is)^{-1}$ is analytic for $z \in D_{d_1}$, where $0 < d_1 < |\bar{u}_x|$. By results in Lukacs (1970, Chapter 7) and for $\delta > 0$ defined as in the statement of the theorem, $\exp\{K(iz)\}$ is analytic in D_δ provided that $\bar{u}_x - y$ lies interior to \mathcal{E} . By considering the set of all possible cases (here, governed by the four basic forms of \mathcal{E} and whether or not \bar{u}_x is > 0), it is then easy to show $g_x(z)$ is analytic in D_{d^*} provided d^* is strictly less than the upper bounds specified in the statement of the theorem.

Let d^* satisfy the stated conditions. Then, in order to prove (ii), it must be shown that

$$\int_{-d^*}^{d^*} |g_x(s + iy)| dy = \int_{-d^*}^{d^*} \frac{|\exp\{K(\bar{u}_x - y + is) - K(\bar{u}_x)\}|}{\sqrt{(\bar{u}_x - y)^2 + s^2}} dy \quad (25)$$

decays to zero as $|s| \rightarrow \infty$. Since $|y| \leq d^*$, it follows that $|K(\bar{u}_x - y)| < \infty$; moreover,

$$|\exp\{K(\bar{u}_x - y + is) - K(\bar{u}_x - y)\}| \leq 1$$

for all $s \in \mathbb{R}$. Consequently,

$$|\exp\{K(\bar{u}_x - y + is) - K(\bar{u}_x)\}| \leq \exp\{K(\bar{u}_x - y) - K(\bar{u}_x)\} < \infty,$$

and the right-hand side of (25) goes to zero as $|s| \rightarrow \infty$.

To prove (iii), let $y \rightarrow (d^*)^-$. Then, it suffices to establish the boundedness of

$$\int_{-\infty}^{\infty} \left[e^{-xd^*} |g_x(s + id^*)| + e^{xd^*} |g_x(s - id^*)| \right] ds.$$

With $M(u) = \exp\{K(u)\}$ and $z = s \pm id^*$,

$$\begin{aligned} |g_x(z)| &= \frac{|\exp\{K(\bar{u}_x \mp d^* + is) - K(\bar{u}_x)\}|}{\sqrt{(\bar{u}_x \mp d^*)^2 + s^2}} \\ &= \frac{M(\bar{u}_x \mp d^*)}{M(\bar{u}_x)} \times \frac{|\exp\{K(\bar{u}_x \mp d^* + is) - K(\bar{u}_x \mp d^*)\}|}{\sqrt{(\bar{u}_x \mp d^*)^2 + s^2}}. \end{aligned} \quad (26)$$

The restrictions on d^* ensure that $M(\bar{u}_x \mp d^*) < \infty$ and $(\bar{u}_x \mp d^*)^2 > 0$. This, combined with the ridge property of characteristic functions (e.g., Lukacs, 1970 or Daniels, 1954), implies that the integrand is bounded and exists for $s = 0$. Provided

$$|\exp\{K(\bar{u}_x \mp d^* + is) - K(\bar{u}_x \mp d^*)\}| = O(|s|^{-r})$$

as $|s| \rightarrow 0$ for some $r > 0$, (26) is $o(|s|^{-1})$ as $|s| \rightarrow \infty$ and therefore integrable.

Let $F'(\cdot)$ denote the probability density function associated with the CDF $F(\cdot)$ of X ; note that $F'(\cdot)$ exists since X is assumed to be absolutely continuous. Then,

$$\exp\{K(\bar{u}_x \mp d^* + is) - K(\bar{u}_x \mp d^*)\} = \int_{-\infty}^{\infty} e^{isu} \left[\frac{e^{(\bar{u}_x \mp d^*)u} F'(u)}{M(\bar{u}_x \mp d^*)} \right] du. \quad (27)$$

The right-hand side is exactly the same as the left-hand side of (6) with $\beta = \bar{u}_x \mp d^*$. Since $M(\bar{u}_x \mp d^*) < \infty$ by construction, the proof is complete. \square

Proof of Proposition 2

Because $M(\bar{u}_x + it)$ is the Fourier transform of $\varpi_T(s) = e^{\bar{u}_x s} F'(s)$,

$$\begin{aligned} M(\bar{u}_x + it) &= \int_{-\infty}^{\infty} e^{its} \varpi_T(s) ds \\ &= \int_{-\infty}^{\infty} \cos(ts) \varpi_T(s) ds + i \int_{-\infty}^{\infty} \sin(ts) \varpi_T(s) ds \\ &= 2 \int_0^{\infty} \cos(ts) \varpi_e(s) ds + 2i \int_0^{\infty} \sin(ts) \varpi_o(s) ds, \end{aligned}$$

the last line implying that $\xi_R(t) = 2 \int_0^{\infty} \cos(ts) \varpi_e(s) ds$ and $\xi_I(t) = 2 \int_0^{\infty} \sin(ts) \varpi_o(s) ds$.

Consider first $\xi_R(t) = 2 \int_0^{\infty} \cos(ts) \varpi_e(s) ds$, which may also be written $\xi_R(t) = 2\Re\{\int_0^{\infty} e^{its} \varpi_e(s) ds\}$.

Using the results of Wong (1989, §IV.2), the asymptotic expansion for $\varpi_e(s)$ as $s \downarrow 0$ implies

$$\xi_R(t) \sim \frac{\gamma_{e,0} \cos\left(\frac{\pi\alpha_e}{2}\right) \Gamma(\alpha_e)}{t^{\alpha_e}} + \frac{\gamma_{e,1} \cos\left(\frac{\pi(\alpha_e+1)}{2}\right) \Gamma(\alpha_e + 1)}{t^{\alpha_e+1}} + o(t^{-(\alpha_e+1)}), \quad t \rightarrow \infty.$$

Similarly, as $t \rightarrow \infty$,

$$\xi_I(t) \sim \frac{\gamma_{o,0} \sin\left(\frac{\pi\alpha_o}{2}\right) \Gamma(\alpha_o)}{t^{\alpha_o}} + \frac{\gamma_{o,1} \sin\left(\frac{\pi(\alpha_o+1)}{2}\right) \Gamma(\alpha_o + 1)}{t^{\alpha_o+1}} + o(t^{-(\alpha_o+1)}).$$

Notice that the validity of these expansions implies that both $\xi_I(t)$ and $\xi_R(t)$ become of constant sign as $t \rightarrow \infty$. Moreover,

$$\frac{\xi_I(t)}{\xi_R(t)} \sim \frac{\gamma_{o,0} \sin\left(\frac{\pi\alpha_o}{2}\right) \Gamma(\alpha_o) t^{\alpha_e - \alpha_o} + \gamma_{o,1} \sin\left(\frac{\pi(\alpha_o+1)}{2}\right) \Gamma(\alpha_o + 1) t^{\alpha_e - \alpha_o - 1} + o(t^{\alpha_e - \alpha_o - 1})}{\gamma_{e,0} \cos\left(\frac{\pi\alpha_e}{2}\right) \Gamma(\alpha_e) + \gamma_{e,1} \cos\left(\frac{\pi(\alpha_e+1)}{2}\right) \Gamma(\alpha_e + 1) t^{-1} + o(t^{-1})},$$

from which various possibilities now present themselves. In particular, observe

$$\frac{\xi_I(t)}{\xi_R(t)} \sim \begin{cases} \pm\infty & \alpha_e > \alpha_o > 0 \\ \frac{\gamma_{o,0} \sin\left(\frac{\pi\alpha}{2}\right) \Gamma(\alpha) + \gamma_{o,1} \sin\left(\frac{\pi(\alpha+1)}{2}\right) \Gamma(\alpha+1) t^{-1} + o(t^{-1})}{\gamma_{e,0} \cos\left(\frac{\pi\alpha}{2}\right) \Gamma(\alpha) + \gamma_{e,1} \cos\left(\frac{\pi(\alpha+1)}{2}\right) \Gamma(\alpha+1) t^{-1} + o(t^{-1})} & \alpha_e = \alpha_o = \alpha > 0 \\ 0 & 0 < \alpha_e < \alpha_o. \end{cases}$$

The limit in the case where $\alpha_e = \alpha_o$ depends on the structure of the expansions for $\varpi_e(s)$ and $\varpi_o(s)$, and in particular on the coefficient sequences $\gamma_{k,j}$, $j \geq 0$, $k = e, o$. In the present situation, the only important implication is that a unique limit exists, whether or not it is

finite. These observations, combined with the fact that $\xi_I(t)$ and $\xi_R(t)$ become of constant sign as $t \rightarrow \infty$, now imply that

$$\arctan\left(\left|\frac{\xi_I(t)}{\xi_R(t)}\right|\right) \sim \begin{cases} \frac{\pi}{2} & \alpha_e > \alpha_o > 0 \\ \theta_0 & \alpha_e = \alpha_o = \alpha > 0 \\ 0 & 0 < \alpha_e < \alpha_o, \end{cases}$$

where $\theta_0 \in [-\pi/2, \pi/2]$. The stated result now follows from (17). \square

Appendix B: The Epsilon algorithm

Let $\{S_0, S_1, \dots, S_J\}$ be a finite subsequence of $\{S_j, j \geq 0\}$ and let S_∞ denote its limit. Applying the epsilon algorithm to $\{S_0, S_1, \dots, S_J\}$ means the following:

1. For $j = 0 \dots J$, set $\varepsilon_0^{(j)} = S_j$.
2. For $j = 0 \dots J - 1$, set $\varepsilon_1^{(j)} = \left(\varepsilon_0^{(j+1)} - \varepsilon_0^{(j)}\right)^{-1}$.
3. For $k = 2 \dots J$ and $j = 0 \dots J - k$ compute

$$\varepsilon_k^{(j)} = \varepsilon_{k-2}^{(j+1)} + \left(\varepsilon_{k-1}^{(j+1)} - \varepsilon_{k-1}^{(j)}\right)^{-1}.$$

4. Output $\varepsilon_m^{(n)}$ for some $m, n \leq J$.

The behavior of the epsilon algorithm depends strongly on whether J is even or odd; see, for example, Wimp (1981, p. 141). In fact, with $n = 0$, only sequences of the form $\varepsilon_{2M}^{(0)}$ converge to S_∞ as $M \rightarrow \infty$; the sequence members $\varepsilon_{2M+1}^{(0)}$ constitute auxiliary quantities that generally diverge as $M \rightarrow \infty$. It can be shown that if the sequence $\{S_j, j \geq 0\}$ is *totally oscillating* (i.e., alternating, and with sequence terms obeying additional regularity conditions), $\varepsilon_{2M}^{(0)}$ converges to S_∞ at an accelerated rate (Brezinski and Redivo Zaglia, 1991, Theorems 2.19 and 2.25).

Following previously established conventions (e.g., Sauter, 2000), the result of applying the epsilon algorithm to $\{S_0, S_1, \dots, S_J\}$ is therefore taken to be $\varepsilon_{2K_J}^{(0)}$, where $K_J = 2\lfloor J/2 \rfloor$. This implies in particular that $\varepsilon_{2K_J}^{(0)} = \varepsilon_J^{(0)}$ if $J = 2M$ and $\varepsilon_{2K_J}^{(0)} = \varepsilon_{J-1}^{(0)}$ if $J = 2M + 1$. It further follows that applying the epsilon algorithm to $\{S_0, S_1, \dots, S_J\}$ and to $\{S_0, S_1, \dots, S_{J+1}\}$ for $J = 2M$ yields exactly the same result. This explains why two members of the sequence are computed at each step of the algorithm of Section 4.5.

Table 1: Unweighted mixture of 2 noncentral chi-square random variables

x	h	d^*	$\hat{\mathcal{N}}_x$	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
0.1	16.785	22.334	4.92	56	1.9×10^{-11}	1.6×10^{-11}	9.8×10^{-12}	1.7×10^{-8}
1.0	0.839	2.087	5.02	101	1.5×10^{-8}	1.7×10^{-14}	2.5×10^{-9}	1.3×10^{-5}
3.0	0.201	0.595	5.04	157	1.2×10^{-7}	1.6×10^{-16}	8.9×10^{-10}	5.6×10^{-5}
5.0	0.097	0.305	5.38	161	1.1×10^{-6}	5.3×10^{-14}	1.3×10^{-7}	5.8×10^{-5}
7.0	0.058	0.187	5.96	200	5.1×10^{-7}	1.9×10^{-15}	5.0×10^{-8}	2.1×10^{-4}
8.0	0.029	0.091	4.52	277	3.6×10^{-8}	2.4×10^{-15}	1.4×10^{-8}	—
9.0	0.030	0.097	4.33	229	9.0×10^{-8}	2.2×10^{-15}	1.2×10^{-8}	2.9×10^{-4}
11.0	0.035	0.109	4.16	168	1.4×10^{-6}	4.1×10^{-15}	8.6×10^{-8}	2.8×10^{-4}
13.0	0.039	0.120	4.31	161	3.7×10^{-8}	1.5×10^{-14}	1.0×10^{-9}	2.2×10^{-4}
15.0	0.041	0.125	4.64	130	6.8×10^{-8}	2.4×10^{-13}	4.3×10^{-9}	1.6×10^{-4}

Table 2: Weighted mixture of 25 noncentral chi-square random variables

x	h	d^*	$\hat{\mathcal{N}}_x$	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
52.682	0.047	0.117	4.59	70	4.4×10^{-14}	3.9×10^{-14}	5.4×10^{-14}	9.6×10^{-7}
90.0	0.013	0.041	4.30	98	2.6×10^{-12}	2.7×10^{-15}	2.1×10^{-10}	4.3×10^{-5}
120.0	0.005	0.015	4.54	234	4.0×10^{-18}	3.9×10^{-18}	5.5×10^{-18}	—
150.0	0.007	0.021	4.10	108	5.1×10^{-11}	5.1×10^{-11}	7.1×10^{-11}	2.7×10^{-5}
295.678	0.010	0.031	58.94	70	3.5×10^{-16}	1.3×10^{-16}	4.1×10^{-15}	1.1×10^{-9}

Table 3: Weighted mixture of 4 noncentral chi-square random variables

x	h	d^*	\hat{N}_x	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
-80.0	0.006	0.016	5.87	156	1.0×10^{-8}	1.0×10^{-8}	1.0×10^{-8}	2.8×10^{-4}
-40.0	0.005	0.014	4.25	275	3.8×10^{-9}	3.8×10^{-9}	3.8×10^{-9}	1.1×10^{-3}
-10.0	0.004	0.012	4.02	706	2.0×10^{-9}	2.0×10^{-9}	2.0×10^{-9}	1.8×10^{-3}
10.0	0.003	0.010	4.17	875	1.7×10^{-8}	1.7×10^{-8}	1.7×10^{-8}	9.9×10^{-4}
40.0	0.003	0.008	4.58	432	1.5×10^{-8}	1.5×10^{-8}	1.5×10^{-8}	1.8×10^{-3}
80.0	0.004	0.011	4.04	211	2.2×10^{-9}	2.2×10^{-9}	3.1×10^{-9}	2.1×10^{-3}
120.0	0.004	0.013	4.29	142	2.1×10^{-8}	2.1×10^{-8}	2.2×10^{-8}	7.2×10^{-4}

Table 4: Time dependent mean of regulated Brownian motion

x	h	d^*	\hat{N}_x	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
0.01	23.094	73.549	12.39	291	3.6×10^{-4}	1.7×10^{-15}	1.2×10^{-7}	1.6×10^{-2}
0.1	2.095	7.037	13.57	345	3.3×10^{-4}	1.5×10^{-17}	6.8×10^{-8}	3.4×10^{-2}
0.5	0.038	0.125	16.05	3313	1.3×10^{-4}	1.1×10^{-16}	4.3×10^{-8}	—
1.0	0.038	0.125	16.23	1591	1.1×10^{-4}	6.1×10^{-16}	2.0×10^{-7}	2.7×10^{-2}
2.0	0.038	0.125	16.80	888	1.4×10^{-4}	3.8×10^{-17}	2.6×10^{-8}	1.3×10^{-2}
3.0	0.039	0.125	17.63	598	1.5×10^{-4}	3.9×10^{-17}	2.6×10^{-8}	6.8×10^{-3}
4.0	0.039	0.125	18.73	454	4.6×10^{-4}	4.0×10^{-17}	2.2×10^{-8}	3.5×10^{-3}
5.0	0.039	0.125	20.13	370	3.7×10^{-5}	3.9×10^{-17}	2.3×10^{-8}	1.8×10^{-3}
6.0	0.039	0.125	21.85	309	3.9×10^{-4}	3.5×10^{-17}	1.7×10^{-8}	9.7×10^{-4}
8.0	0.040	0.125	26.33	275	1.9×10^{-4}	8.6×10^{-18}	2.8×10^{-9}	2.8×10^{-4}
10.0	0.041	0.125	32.47	227	1.0×10^{-4}	2.2×10^{-17}	2.3×10^{-9}	8.5×10^{-5}

Table 5: A nonregular exponential family

x	h	d^*	\hat{N}_x	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
1.1	2.844	9.492	8.19	288	3.1×10^{-5}	1.9×10^{-17}	2.1×10^{-8}	4.9×10^{-3}
2.0	0.078	0.250	9.70	1285	1.3×10^{-5}	1.0×10^{-17}	5.5×10^{-9}	9.4×10^{-3}
3.0	0.080	0.250	10.79	674	1.7×10^{-5}	6.4×10^{-18}	4.2×10^{-9}	3.4×10^{-3}
4.0	0.080	0.250	12.55	363	1.3×10^{-6}	1.4×10^{-17}	5.1×10^{-9}	1.0×10^{-3}
6.0	0.083	0.250	18.61	250	7.5×10^{-6}	4.2×10^{-17}	1.5×10^{-9}	9.4×10^{-5}
8.0	0.086	0.250	29.41	176	4.9×10^{-6}	1.2×10^{-16}	5.4×10^{-10}	8.8×10^{-6}

Table 6: Predicted vs. computed zero spacings in Table 5

	x	Predicted interval	Computed interval
	1.1	2.856	31.545
	2.0	1.571	3.142
Table	3.0	1.047	1.571
5	4.0	0.785	1.051
	6.0	0.524	0.631
	8.0	0.393	0.451

Table 7: Compound Polya sum of Exponential(1) random variables

x	h	d^*	$\widehat{\mathcal{N}}_x$	# of points	Absolute Error			
					Straight(x)	Epsilon(x)	Euler(x)	LR(x)
0.05	7.025	19.875	7.99	201	1.3×10^{-6}	1.7×10^{-15}	9.4×10^{-8}	5.9×10^{-4}
0.5	0.581	1.879	8.25	242	2.7×10^{-5}	1.9×10^{-15}	6.4×10^{-8}	7.0×10^{-4}
1.0	0.271	0.885	8.59	264	1.5×10^{-4}	2.9×10^{-14}	1.2×10^{-7}	5.3×10^{-4}
2.0	0.059	0.188	5.65	552	2.0×10^{-5}	1.7×10^{-16}	3.0×10^{-8}	1.5×10^{-3}
4.0	0.061	0.188	6.60	286	6.1×10^{-6}	3.8×10^{-16}	1.5×10^{-8}	9.0×10^{-4}
8.0	0.064	0.188	11.59	172	8.8×10^{-6}	6.9×10^{-17}	1.1×10^{-9}	9.2×10^{-5}
12.0	0.066	0.188	23.40	131	3.8×10^{-7}	1.8×10^{-16}	3.6×10^{-9}	6.6×10^{-6}
16.0	0.069	0.188	49.01	110	3.1×10^{-7}	9.8×10^{-16}	1.1×10^{-10}	4.2×10^{-7}

References

- Abate, J. and Whitt, W. (1987), “Transient behavior of regulated Brownian motion I: starting at the origin,” *Advances in Applied Probability*, **19**, 560-598.
- Abate, J. and Whitt, W. (1992), “The Fourier series method for inverting transforms of probability distributions,” *Queueing Systems*, **10**, 5-87.
- Abate, J. and Whitt, W. (1995), “Numerical inversion of Laplace transforms of probability distributions,” *ORSA Journal on Computing*, **7**, 36-43.
- Abate, J., Choudhury, G. and Whitt, W. (1999), “An Introduction to Numerical Transform Inversion and its Application to Probability Models,” In *Computational Probability*, W. Grassman (ed.), Boston: Kluwer, pp. 257-323
- Bak, J. and Newman, D.J. (1996), *Complex Analysis, 2nd ed.*, New York: Springer-Verlag.
- Barndorff-Nielsen, O.E. and Cox, D.R. (1989). *Asymptotic Techniques for Use in Statistics*, London: Chapman-Hall.
- Barndorff-Nielsen, O.E. and Cox, D.R. (1994). *Inference and Asymptotics*, London: Chapman Hall.
- Bleistein, N.L. and Handelsman, R.A. (1975), *Asymptotic Expansions of Integrals*. New York: Dover.
- Böhman, H. (1975), “Numerical inversion of characteristic functions,” *Scand. Actuarial Journal*, **2**, 121-124.
- Brezinski, C. and Redivo Zaglia, M., (1991), *Extrapolation methods: Theory and practice*, Amsterdam: North-Holland.
- Butzer, P.L. and Stens, R.L. (1983), “The Poisson summation formula, Whittaker’s cardinal series and approximate integration,” In *Second Edmonton conference on approximation theory*, Canadian Mathematical Society Conference Proceedings (Edmonton, Alberta, 1982), Providence: American Mathematical Society.
- Choudhury, G.L., Lucantoni, D.M., and Whitt, W. (1994), “Multidimensional transform inversion with applications to the transient $M/G/1$ queue,” *Annals of Applied Probability*, **4**, 719-740.

- Daniels, H. (1954), "Saddlepoint approximations in statistics," *Annals of Mathematical Statistics*, **25**, 631-650.
- Daniels, H. (1983), "Saddlepoint approximations for estimating equations," *Biometrika*, **70**, 89-96.
- Daniels, H. (1987), "Tail probability approximations," *International Statistical Review*, **55**, 37-48.
- Davies, R.B. (1973), "Numerical inversion of a characteristic function," *Biometrika*, **60**, 415-417.
- Davies, R.B. (1980), "[Algorithm AS 155] The distribution of a linear combination of χ^2 random variables," *Applied Statistics*, **29**, 323-333.
- Davis, P. and Rabinowitz, P. (1984), *Methods of Numerical Integration*, 2nd ed., New York: Academic Press.
- Feller, W. (1971), *An Introduction to Probability Theory and its Applications*, Vol. 2, 2nd ed., New York: Wiley.
- Field, C.A. and Ronchetti, E. (1990), *Small Sample Asymptotics*, IMS Lecture Notes Vol. 13, Hayward: Institute of Mathematical Statistics.
- Helstrom, C.W. (1983), Comment on "Distribution of quadratic forms in normal random variables—evaluation by numerical integration," *SIAM J. Sci. Statist. Comput.*, **4**, 353-356.
- Imhof, J.P. (1961), "Computing the distribution of quadratic forms in normal variables," *Biometrika*, **48**, 419-426.
- Jensen, J.L. (1995), *Saddlepoint Approximations*, Oxford: Clarendon Press.
- Johnson, N.L., Kotz, S., and Balakrishnan, N. (1995), *Continuous univariate distributions: Vol. 2, 2nd ed.*, New York: Wiley.
- Kawata, T. (1972), *Fourier Analysis in Probability Theory*, New York: Academic Press.
- Kolassa, J. (1997), *Series Approximation Methods in Statistics*, 2nd ed., New York; Springer-Verlag.

- Krommer, A.R. and Ueberhuber, C.W. (1998), *Computational Integration*, Philadelphia : Society for Industrial and Applied Mathematics.
- Lugannani, R. and Rice, S.O. (1980), "Saddlepoint approximation for the distribution of the sum of independent random variables," *Advances in Applied Probability*, **12**, 475-490.
- Lukacs, E. (1970), *Characteristic functions, 2nd ed.*, London: Griffin.
- Mehta, C.R., Senchaudhuri, P., Patel, N. (1998), Discussion of "Approximately exact inference for the common odds ratio in several 2×2 tables" by Strawderman and Wells, *Journal of the American Statistical Association*, **93**, 1313-1316.
- Olver, F.W.J. (1974), *Asymptotics and Special Functions*, New York: Academic Press.
- Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T., (1989), *Numerical recipes: The art of scientific computing (FORTRAN version)*, Cambridge: Cambridge University Press.
- Rice, S.O. (1980), "Distribution of quadratic forms in normal random variables—evaluation by numerical integration," *SIAM J. Sci. Statist. Comput.*, **1**, 438–448.
- Sauter, T. (2000), "Computation of irregularly oscillating integrals," *Applied Numerical Mathematics*, **35**, 245-264.
- Skovgaard, I.M. (1987), "Saddlepoint expansions for conditional distributions," *Journal of Applied Probability*, **24**, 875-887.
- Smith, D.A. and Ford, W.F. (1982), "Numerical comparisons of nonlinear convergence accelerators," *Mathematics of Computation*, **38**, 481-499.
- Stenger, F. (1973), *Numerical Methods Based on Sinc and Analytic Functions*, New York: Springer-Verlag.
- Strawderman, R.L. (2000), "Higher-order asymptotic approximation: Laplace, saddlepoint, and related methods," *Journal of the American Statistical Association*, **95**, 1358-1364.
- Titchmarsh, E.C. (1948), *Introduction to the Theory of Fourier Integrals, 2nd ed.*, Oxford: Clarendon Press.

Waller, L.A., Turnbull, B.W., and Hardin, J.M. (1995), "Obtaining distribution functions by numerical inversion of characteristic functions with applications," *American Statistician*, **49**, 346-350.

Waterloo Maple Inc. (2000), *Maple 6 for Linux*, Waterloo: Waterloo Maple Software.

Wimp, J. (1981), *Sequence transformations and their applications*, New York: Academic Press.

Wong, R. (1989), *Asymptotic approximations of integrals*, Boston: Academic Press.