

# Simulation from the Tail of the Univariate and Multivariate Normal Distribution

Zdravko Botev · Pierre L'Ecuyer

Received: date / Accepted: date

**Abstract** We study and compare various methods to generate a random variate or vector from the univariate or multivariate normal distribution truncated to some finite or semi-infinite region, with special attention to the situation where the regions are far in the tail. This is required in particular for certain applications in Bayesian statistics, such as to perform exact posterior simulations for parameter inference, but could have many other applications as well. We distinguish the case in which inversion is warranted, and that in which rejection methods are preferred.

**Keywords** truncated · tail · normal · Gaussian · simulation · multivariate

## 1 Introduction

We consider the problem of simulating a standard normal random variable  $X$ , conditional on  $a \leq X \leq b$ , where  $a < b$  are real numbers, and at least one of them is finite. We are particularly interested in the situation where the interval  $(a, b)$  is far in one of the tails, that is,  $a \gg 0$  or  $b \ll 0$ . The standard methods developed for the non-truncated case do not always work well in this case. Moreover, if we insist on using inversion, the standard inversion methods break down when we are far in the tail. Inversion is preferable to a rejection method (in general) in various simulation applications, for example to maintain synchronization and monotonicity when comparing systems with

---

Zdravko Botev  
UNSW Sydney, Australia  
Tel.: +(61) (02) 9385 7475  
E-mail: botev@unsw.edu.au

Pierre L'Ecuyer  
Université de Montréal, Canada and Inria–Rennes, France  
Tel.: +(514) 343-2143  
E-mail: lecuyer@iro.umontreal.ca

common random numbers, for derivative estimation and optimization, when using quasi-Monte Carlo methods, etc. [11–14, 5]. For this reason, a good inversion method is needed, even if rejection is faster. We examine both rejection and inversion methods in this paper.

We also consider simulation of a multivariate normal vector  $\mathbf{X}$  with mean zero and covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$ , conditional on  $\mathbf{X} \geq \mathbf{a}$ , where the inequality is componentwise.

These problems occur in particular for the estimation of certain Bayesian regression models and for exact simulation from these models; see [6, 4] and the references given there. The simulation from the Bayesian posterior requires repeated draws from a standard normal distribution truncated to different intervals, often far in the tail. That is, we need fast and reliable algorithms to generate  $X \sim \mathcal{N}(0, 1)$ , conditional on  $a \leq X \leq b$ , for arbitrary real numbers  $a < b$ .

Various methods have already been proposed to do that; see for example [6, 7, 10, 18, 21, 22]. Some methods work well when the interval  $[a, b]$  contains 0 or is not far from it, but not when  $a \gg 0$  or  $b \ll 0$ . Other methods have been designed for the right tail, that is, when  $a \gg 0$  and  $b = \infty$ , and use rejection. These methods may be adapted in principle to a finite interval  $[a, b]$ , but they may become inefficient when the interval  $[a, b]$  is narrow. We also found no reliable inversion method for an interval far in the tail (say, for  $a > 38$ ; see Section 2). Note that to generate  $X$  from a more general normal distribution with mean  $\mu$  and variance  $\sigma^2$  truncated to an interval  $(a', b')$ , it suffices to apply a simple linear transformation to recover the problem studied here.

Efficient and reliable simulation methods based on importance sampling were developed recently in [3–5] for exact simulation from such multivariate conditional distributions and to estimate the conditional probability  $\mathbb{P}[\mathbf{a} \leq \mathbf{X} \leq \mathbf{b}]$ .

The aim of this paper is twofold. The first is to review and compare the most popular methods for the univariate normal distribution, propose new efficient methods for certain situations, and provide reliable software implementation of these methods. In particular, in the univariate setting we propose a new accurate inversion method for arbitrarily large  $a$  and improvements to commonly used methods. In some of our discussion, we assume that  $a > 0$ . The case where  $b < 0$  is covered by symmetry (just change the sign) and the case where  $a \leq 0 \leq b$  can be handled by standard methods.

The second aim is to propose a simple acceptance-rejection method in the multivariate setting, which, under some conditions, can yield an acceptance probability that approaches unity as we move deeper and deeper into the tail region. While [3] proposes a more complex method for simulation from the truncated multivariate normal distribution than the much simpler one considered here, the theoretical results in [3] do not apply when the target pdf is the most general tail density (see (9) in Section 3). Thus, the multivariate aspect in Section 3 of the present paper both fills an important theoretical gap in [3] and described a simpler method, based on a multivariate exponential proposal.

## 2 Simulation from Tail of Univariate Normal

In this section, we use  $\phi$  to denote the density of the standard normal distribution (with mean 0 and variance 1),  $\Phi$  for its cumulative distribution function

(cdf),  $\bar{\Phi}$  for the complementary cdf, and  $\Phi^{-1}$  for the inverse cdf defined as  $\Phi^{-1}(u) = \min\{x \in \mathbb{R} \mid \Phi(x) \geq u\}$ . Thus, if  $X \sim \mathbf{N}(0, 1)$ ,  $\Phi(x) = \mathbb{P}[X \leq x] = \int_{-\infty}^x \phi(y)dy = 1 - \bar{\Phi}(x)$ . Conditional on  $a \leq X \leq b$ ,  $X$  has density

$$\frac{\phi(x)}{\Phi(b) - \Phi(a)} \quad \text{for } a < x < b \quad (1)$$

We denote this truncated normal distribution by  $\text{TN}_{a,b}(0, 1)$ .

It is well known that if  $U \sim U(0, 1)$ , the uniform distribution over the interval  $(0, 1)$ , then

$$X = \Phi^{-1}(\Phi(a) + (\Phi(b) - \Phi(a))U) \quad (2)$$

has exactly the standard normal distribution conditional on  $a \leq X \leq b$ . But even though very accurate approximations are available for  $\Phi$  and  $\Phi^{-1}$ , (2) is sometimes useless for simulating  $X$ . One reason for this is that whenever computations are made under the IEEE-754 double precision standard (which is typical), any number of the form  $1 - \epsilon$  for  $0 \leq \epsilon < 2 \times 10^{-16}$  (approximately) is identified with 1.0, any positive number smaller than about  $10^{-324}$  cannot be represented at all (it is identified with 0), and numbers smaller than  $10^{-308}$  are represented with less than 52 bits of accuracy.

This implies that  $\bar{\Phi}(x) = \Phi(-x)$  is identified as 0 whenever  $x \geq 39$  and is identified as 1 whenever  $-x \geq 8.3$ . Thus, (2) cannot work when  $a \geq 8.3$ . In the latter case, or whenever  $a > 0$ , it is much better to use the equivalent form:

$$X = -\Phi^{-1}(\bar{\Phi}(a) - (\bar{\Phi}(a) - \bar{\Phi}(b))U), \quad (3)$$

which is accurate for  $a$  up to about 37, assuming that we use accurate approximations of  $\bar{\Phi}(x)$  for  $x > 0$  and of  $\Phi^{-1}(u)$  for  $u < 1/2$ . Such accurate approximations are available for example in [2] for  $\Phi^{-1}(u)$  and via the error function  $\text{erf}$  on most computer systems for  $\bar{\Phi}(x)$ . For larger values of  $a$  (and  $x$ ), a different inversion approach must be developed, as shown next.

## 2.1 Inversion far in the Right Tail

When  $\bar{\Phi}(x)$  is too small to be represented as a floating-point double, we will work instead with the Mills' [25] ratio, defined as  $q(x) \stackrel{\text{def}}{=} \bar{\Phi}(x)/\phi(x)$ , which is the inverse of the hazard rate (or failure rate) evaluated at  $x$ . When  $x$  is large, this ratio can be approximated by the truncated series (see [1]):

$$q(x) \approx \frac{1}{x} + \sum_{n=1}^r \frac{1 \times 3 \times 5 \times \cdots \times (2n-1)}{(-1)^n x^{2n+1}}. \quad (4)$$

In our experiments with  $x \geq 10$ , we compared  $r = 5, 6, 7, 8$ , and we found no significant difference (up to machine precision) in the approximation of  $X$  in 3 by the method we now describe. In view of (3), we want to find  $x$  such that

$\bar{\Phi}(x) = \Phi(-x) = \bar{\Phi}(a) - (\bar{\Phi}(a) - \bar{\Phi}(b))u$ , for  $0 \leq u \leq 1$ , when  $a$  is large. This equation can be rewritten as  $h(x) = 0$ , where

$$h(x) \stackrel{\text{def}}{=} \bar{\Phi}(a) - \bar{\Phi}(x) + (\bar{\Phi}(b) - \bar{\Phi}(a))u \quad (5)$$

To solve  $h(x) = 0$ , we start by finding an approximate solution and then refine this approximation via Newton iterations. We detail how this is achieved. To find an approximate solution, we replace the normal cdf  $\Phi$  in (3) by the standard Rayleigh distribution, whose complementary cdf and density are given by  $\bar{F}(x) = \exp(-x^2/2)$  and  $f(x) = x \exp(-x^2/2)$  for  $x > 0$ . Its inverse cdf can be written explicitly as  $F^{-1}(u) = (-2 \ln(1 - u))^{1/2}$ . This choice of approximation of  $\bar{\Phi}^{-1}$  in the tail has been used before (see for example [2] and Section 4). It is motivated by the facts that  $F^{-1}(u)$  is easy to compute and that  $\bar{\Phi}(x)/\bar{F}(x) \rightarrow 1$  rapidly when  $x \rightarrow \infty$ . By plugging  $\bar{F}$  and  $F^{-1}$  in place of  $\bar{\Phi}$  and  $\bar{\Phi}^{-1}$  in (3), and solving for  $x$ , we find the approximate root

$$x \approx \sqrt{a^2 - 2 \ln(1 - u + u \exp((a^2 - b^2)/2))}, \quad (6)$$

which is simply the  $u$ -th quantile of the standard Rayleigh distribution truncated over  $(a, b)$ , with density

$$f(x) = \frac{x \exp(-(x^2 - a^2)/2)}{1 - \exp(-(b^2 - a^2)/2)} \quad \text{for } a < x < b. \quad (7)$$

The next step is to improve the approximation (6) by applying Newton's method to (5). For this, it is convenient to make the change of variable  $x = \xi(z)$ , where  $\xi(z) \stackrel{\text{def}}{=} \sqrt{a^2 - 2 \ln(z)}$  and  $z = \xi^{-1}(x) = \exp((a^2 - x^2)/2)$ , and apply Newton's method to  $g(z) \stackrel{\text{def}}{=} h(\xi(z))$ . Newton's iteration for solving  $g(z) = 0$  has the form  $z_{\text{new}} = z - g(z)/g'(z)$ , where

$$\begin{aligned} \frac{g(z)}{g'(z)} &= \frac{h(\xi(z))}{h'(\xi(z))} \cdot \frac{1}{\xi'(z)}, \quad (\text{by the chain rule}) \\ &= z\xi(z) \frac{\bar{\Phi}(\xi(z)) - \bar{\Phi}(a) + u(\bar{\Phi}(a) - \bar{\Phi}(b))}{\phi(\xi(z))} \\ &= x \left( zq(x) - q(a)(1 - u) - q(b)u \exp\left(\frac{a^2 - b^2}{2}\right) \right), \end{aligned}$$

and the identity  $x = \xi(z)$  was used for the last equality. A key observation here is that, thanks to the replacement of  $\bar{\Phi}$  by  $q$ , the computation of  $g(z)/g'(z)$  does not involve extremely small quantities that can cause numerical underflow, even for extremely large  $a$ .

The complete procedure is summarized in Algorithm 1, which we have implemented in Java, MATLAB<sup>®</sup>, and **R**. According to our experiments, the larger the  $a$ , the faster the convergence. For example, for  $a = 50$  one requires at most 13 iterations to ensure  $\delta_x \leq \delta^* = 10^{-10}$ , where  $\delta_x$  represents the relative change in  $x$  in the last Newton iteration.

---

**Algorithm 1** : Returns the  $u$ -quantile of  $\text{TN}_{a,b}(0, 1)$ 


---

**Require:** Input  $u \in (0, 1)$ ,  $\delta^*$  $q_a \leftarrow q(a)$  $q_b \leftarrow q(b)$  $c \leftarrow q_a(1 - u) + q_b u \exp(\frac{a^2 - b^2}{2})$  $\delta_x \leftarrow \infty$  $z \leftarrow 1 - u + u \exp(\frac{a^2 - b^2}{2})$  $x \leftarrow \sqrt{a^2 - 2 \ln(z)}$ **repeat** $z \leftarrow z - x(zq(x) - c)$  $x_{\text{new}} \leftarrow \sqrt{a^2 - 2 \ln(z)}$  $\delta_x \leftarrow |x_{\text{new}} - x|/x$  $x \leftarrow x_{\text{new}}$ **until**  $\delta_x \leq \delta^*$ **return** Quantile  $x$ 


---

We note that for an interval  $[a, b] = [a, a + w]$  of fixed length  $w$ , when  $a$  increases the conditional density concentrates closer to  $a$ . In fact, there is practically no difference between generating  $X$  conditional on  $a \leq X \leq a + 1$  and conditional on  $X \geq a$  when  $a \geq 30$ , but there can be a significant difference for small  $a$ .

## 2.2 Rejection Methods

We now examine *rejection* (or *acceptance-rejection*) methods, which can be faster than inversion. A large collection of rejection-based generation methods for the normal distribution have been proposed over the years; see [6, 7, 10, 22] for surveys, discussions, comparisons, and tests. Most of them (the fastest ones) use a change of variable and/or precomputed tables to speedup the computations. In its most elementary form, a rejection method to generate from some density  $f$  uses a hat function  $h \geq f$  and rescales  $h$  vertically to a probability density  $g = h / \int_{-\infty}^{\infty} h(y) dy$ , often called the proposal density. A random variate  $X$  is generated from  $g$ , is accepted with probability  $f(X)/h(X)$ , is rejected otherwise, and the procedure is repeated until  $X$  is accepted as the retained realization. In practice, more elaborate versions are used that incorporate transformations and partitions of the area under  $h$ .

Any of these proposed rejection methods can be applied easily if  $\Phi(b) - \Phi(a)$  is large enough, just by adding a rejection step to reject any value that falls outside  $[a, b]$ . The acceptance probability for this step is  $\Phi(b) - \Phi(a)$ . When this probability is too small, this becomes too inefficient and something else must be done. One way is to define a proposal  $g$  whose support is exactly  $[a, b]$ , but this could be inefficient (too much overhead) when  $a$  and  $b$  change very often. Chopin [6] developed a rejection method specially adapted to this situation. It is based on a hat function defined by juxtaposing a large number of vertical rectangles of different heights but equal surface over some finite interval  $[a_{\min}, a_{\max}]$ , and use an exponential proposal with rate  $a = a_{\max}$  (the

RejectTail variant of Algorithms 2 below) for the tail above  $a_{\max}$  or when  $a > a'_{\max}$ . The fastest implementation uses 4000 rectangles,  $a_{\max} \approx 3.486$ ,  $a'_{\max} \approx 2.605$ . This method is fast, although it requires the storage of very large precomputed tables, which could actually slow down computations on certain type of hardware for which memory is limited, like GPUs.

Simple rejection methods for the standard normal truncated to  $[a, \infty)$ , for  $a \geq 0$ , have been proposed long ago. Marsaglia [19] proposed a method that uses for  $g$  the standard Rayleigh distribution truncated over  $[a, \infty)$ . An efficient implementation is given in [7, page 381]. Devroye [7, page 382] also gives an algorithm that uses for  $g$  an exponential density of rate  $a$  shifted by  $a$ . These two methods have exactly the same acceptance probability,

$$\alpha(a) = a\sqrt{2\pi} \exp(a^2/2)\bar{\Phi}(a), \quad (8)$$

which converges to 1 when  $a \rightarrow \infty$ . Geweke [8] and Robert [21] optimized the acceptance probability to

$$\beta(a) = \lambda\sqrt{2\pi} \exp(a\lambda - \lambda^2/2) \bar{\Phi}(a)$$

by taking the rate  $\lambda = (a + \sqrt{a^2 + 4})/2 > a$  for the shifted exponential proposal. However, the gain with respect to Devroye's method is small and can be wiped out easily by a larger computing time per step. For large  $a$ , both are very close to 1 and there is not much difference between them.

We will compare two ways of adapting these methods to a truncation over a finite interval  $[a, b]$ . The first one is to keep the same proposal  $g$  which is positive over the interval  $[a, \infty)$  and reject any value generated above  $b$ . The second one truncates and rescales the proposal to  $[a, b]$  and applies rejection with the truncated proposal. We label them by *RejectTail* and *TruncTail*, respectively. *TruncTail* has a smaller rejection probability, by the factor  $1 - \bar{\Phi}(a)/\bar{\Phi}(b)$ , but also entails additional overhead to properly truncate the proposal. Typically, it is worthwhile only if this additional overhead is small and/or the interval  $[a, b]$  is very narrow, so it improves the rejection probability significantly. Our experiments will confirm this.

Algorithms 2, 3, 4, state the rejection methods for the *TruncTail* case with the exponential proposal with rate  $a$  [7], with the rate  $\lambda$  proposed in [21], and with the standard Rayleigh distribution, respectively, extended to the case of a finite interval  $[a, b]$ . For the *RejectTail* variant, one would remove the computation of  $q$ , replace  $\ln(1 - qU)$  by  $\ln U$ , and add  $X \leq b$  to the acceptance condition. Algorithm 5 gives this variant for the Rayleigh proposal.

---

**Algorithm 2** :  $X \sim \text{TN}_{a,b}(0, 1)$  with exponential proposal with rate  $a$ , truncated

---

```

 $K_a \leftarrow 2a^2$ 
 $q \leftarrow 1 - \exp(-(b - a)a)$ 
repeat
  Generate  $U, V \sim \text{U}(0, 1)$ , independent
   $X \leftarrow -\ln(1 - qU)$ 
   $E \leftarrow -\ln(V)$ 
until  $X^2 \leq K_a V$ 
return  $a + X/a$ 

```

---

---

**Algorithm 3** :  $X \sim \text{TN}_{a,b}(0, 1)$  with exponential proposal with rate  $\lambda$ , truncated

---

```

 $\lambda \leftarrow (a + \sqrt{a^2 + 4})/2$ 
 $q \leftarrow 1 - \exp(-(b - a)\lambda)$ 
repeat
  Generate  $U, V \sim \text{U}(0, 1)$ , independent
   $X \leftarrow a - \ln(1 - qU)/\lambda$ 
until  $V \leq \exp((X - \lambda)^2/2)$ 
return  $a + X/a$ 

```

---



---

**Algorithm 4** :  $X \sim \text{TN}_{a,b}(0, 1)$  with Rayleigh proposal, truncated

---

```

 $c \leftarrow a^2/2$ 
 $q \leftarrow 1 - \exp(c - b^2/2)$ 
repeat
  Simulate  $U, V \sim \text{U}(0, 1)$ , independently.
   $X \leftarrow c - \ln(1 - qU)$ 
until  $V^2 X \leq a$ 
return  $X \leftarrow \sqrt{2X}$ 

```

---



---

**Algorithm 5** :  $X \sim \text{TN}_{a,\infty}(0, 1)$  with Rayleigh proposal and RejectTail

---

```

 $c \leftarrow a^2/2$ 
repeat
  Simulate  $U, V \sim \text{U}(0, 1)$ , independently.
   $X \leftarrow c - \ln(U)$ 
until  $V^2 X \leq a$  and  $2X \leq b * b$ 
return  $\sqrt{2X}$ 

```

---

When the interval  $[a, b]$  is very narrow, it makes sense to just use the uniform distribution over this interval for the proposal  $g$ . This is suggested in [21] and shown in Algorithm 6. Generating from the proposal is then very fast. On the other hand, the acceptance probability may become very small if the interval is far in the tail and  $b - a$  is not extremely small. Indeed, the acceptance probability in this case is:

$$\frac{\sqrt{2\pi} \exp(a^2/2) (\bar{\Phi}(a) - \bar{\Phi}(b))}{b-a} = \frac{q(a) - q(b) \exp((a^2 - b^2)/2)}{b-a},$$

which decays at a rate of  $1/a$  when  $a \rightarrow \infty$  while  $(b - a)$  remains constant.

---

**Algorithm 6** :  $X \sim \text{TN}_{a,b}(0, 1)$  with uniform proposal, truncated

---

```

repeat
  Simulate  $U, V \sim \text{U}(0, 1)$ , independently.
   $X \leftarrow a + (b - a)U$ 
until  $2 \ln V \leq a^2 - X^2$ 
return  $X$ 

```

---

Another choice that the user can have with those generators (and for any variate generator that depends on some distribution parameters) is to either

*precompute* various constants that depend on the parameters and store them in some “distribution” object with fixed parameter values, or to *recompute* these parameter-dependent constants each time a new variate is generated. This type of alternative is common in modern variate generation software [15,17]. The first approach is worthwhile if the time to compute the relevant constants is significant and several random variates are to be generated with exactly the same distribution parameters. For the applications in Bayesian statistics mentioned earlier, it is typical that the parameters  $a$  and  $b$  change each time a new variate is generated [6]. But there can be applications in which a large number of variates are generated with the same  $a$  and  $b$ .

For one-sided intervals  $[a, \infty)$ , the algorithms can be simplified. One can use the `RejectTail` framework and since  $b = \infty$ , there is no need to check if  $X \leq b$ . When reporting our test results, we label this the *OneSide* case.

Note that computing an exponential is typically more costly than computing a log (by a factor of 2 or 3 for negative exponents and 10 for large exponents, in our experiments) and the latter is more costly than computing a square root (also by a factor of 10). This means significant speedups could be obtained by avoiding the recomputing of the exponential each time at the beginning of Algorithms 2, 3, and 4. This is possible if the same parameter  $b$  is used several times, or if  $b = \infty$ , or if we use `RejectTail` instead of `TruncTail`.

### 2.3 Speed Comparisons

We report a representative subset of results of speed tests made with the different methods, for some pairs  $(a, b)$ . In each case, we generated  $10^8$  (100 millions) truncated normal variates, added them up, printed the CPU time required to do that, and printed the sum for verification. The experiments were made in Java using the SSJ library [15], under Eclipse and Windows 10, on a Lenovo X1 Carbon Thinkpad with an Intel Core(TM) i7-5600U (single) processor running at 2.60 GHz. All programs were executed in a single thread and the CPU times were measured using the stopwatch facilities in class `Chrono` of SSJ, which relies on the `getThreadCpuTime` method from the Java class `ThreadMXBean` to obtain the CPU time consumed so far by a single thread, and `subtract` to obtain the CPU time consumed between any two instructions.

The measurements were repeated a few times to verify consistency and varied by about 1 to 2 percent at most. The compile times are negligible relative to the reported times. Of course, these timings depend on CPU and memory usage by other processes on the computer, and they are likely to change if we move to a different platform, but on standard processors the relative timings should remain roughly the same. They provide a good idea of what is most efficient to do.

Tables 1 to 2 report the timings, in seconds. The two columns “recompute” and “precompute” are for the cases where the constants that depend on  $a$  and



$b$  are recomputed each time a random variate is generated or are precomputed once and for all, respectively, as discussed earlier.

**Table 1** Time to generate  $n = 10^8$  variates for  $[a, b] = [3.0, 3.1]$  (left pane) and  $[a, b] = [7.0, 8.0]$  (right pane).

Method	CPU time (sec.)		Method	CPU time	
	recom.	precom.		recom.	precom.
Generation in $[a, b]$			Generation in $[a, b]$		
ExponD	6.46	6.22	ExponD	11.70	6.16
ExponDRejectTail	23.04	23.20	ExponDRejectTail	6.04	6.08
ExponR	16.63	9.92	ExponR	15.96	8.98
ExponRRejectTail	32.40	32.40	ExponRRejectTail	9.20	9.09
ExponRRejectTailLog	25.10	25.30	ExponRRejectTailLog	7.03	7.02
Rayleigh	10.29	4.60	Rayleigh	9.86	4.27
RayleighRejectTail	15.23	15.33	RayleighRejectTail	3.91	3.99
Uniform	4.26	4.34	Uniform	25.40	25.68
InverseSSJ	15.14	8.14	InverseSSJ	30.67	8.14
InverseQuickSSJ	18.80	3.31	InverseRightTail	31.12	7.70
InverseRightTail	31.12	7.66	Generation in $[a, \infty)$		
Generation in $[a, \infty)$			ExponDOneSide	5.90	5.96
ExponDOneSide	6.43	6.46	ExponROneSideLog	6.80	6.71
ExponROneSideLog	7.05	6.99	RayleighOneSide	3.74	4.05
RayleighOneSide	4.07	4.41	InverseSSJOneSide	19.00	8.19
InverseSSJOneSide	18.81	8.20	InverseRightTailOneSide	18.76	7.59
InverseRightTailOneSide	18.72	7.64			

**Table 2** Time to generate  $n = 10^8$  variates for  $[a, b] = [100.0, 102.0]$  (left pane) and  $[a, b] = [100.0, 100.0001]$  (right pane).

Method	CPU time (sec.)		Method	CPU time	
	recom.	precom.		recom.	precom.
Generation in $[a, b]$			Generation in $[a, b]$		
ExponD	11.68	6.01	ExponD	12.31	6.83
ExponDRejectTail	5.88	5.91	ExponDRejectTail	543.80	546.58
ExponR	15.79	8.86	ExponR	16.47	10.65
ExponRRejectTail	9.13	9.02	ExponRRejectTail	865.24	865.34
ExponRRejectTailLog	6.93	6.96	ExponRRejectTailLog	651.19	648.99
Rayleigh	9.97	4.16	Rayleigh	10.59	5.07
RayleighRejectTail	3.84	3.90	RayleighRejectTail	323.08	322.41
Uniform	650.62	656.42	Uniform	3.59	3.62
InverseMillsRatio	22.31	15.97	InverseMillsRatio	18.03	12.12
Generation in $[a, \infty)$			Generation in $[a, \infty)$		
ExponDOneSide	5.77	5.82	ExponDOneSide	5.79	5.83
ExponROneSideLog	6.72	6.63	ExponROneSideLog	6.74	6.63
RayleighOneSide	3.67	3.96	RayleighOneSide	3.66	3.99
InverseMillsRatioOneSide	15.62	15.84	InverseMillsRatioOneSide	15.67	15.84

ExponD, ExponR, and Rayleigh refer to the TruncTail versions of Algorithms 2, 3, and 4, respectively. We add “RejectTail” to the name for the RejectTail versions. For ExponRRejectTailLog, we took the log on both sides of the inequality to remove the exponential in the “until” condition. Uniform refers to Algorithm 6.

InversionSSJ refers to the default inversion method implemented in SSJ, which uses [2] and gives at least 15 decimal digits of relative precision, combined with a generic (two-sided) “truncated distribution” class also offered in SSJ. InverseQuickSSJ is a faster but much less accurate version based on a

cruder approximation of  $\bar{\Phi}$  from [20] based on table lookups, which returns about 6 decimal digits of precision. We do not recommend it, due to its low accuracy. Moreover, the implementation we used does not handle well values larger than about 5 in the right tail, so we report results only for small  $a$ . `InverseRightTail` uses the accurate approximation of  $\bar{\Phi}$  together with (3). `InverseMillsRatio` is our new inversion method based on Mills ratio, with  $\delta^* = 10^{-10}$ . This method is designed for the case where  $a$  is large, and our implementation is designed to be accurate for  $a \geq 10$ , so we do not report results for it Table 1. For all the methods, we add “OneSide” for the simplified OneSide versions, for which  $b = \infty$ .

For the OneSide case, that is,  $b = \infty$ , the Rayleigh proposal gives the fastest method in all cases, and there is no significant gain in precomputing and storing the constant  $c = a^2/2$ .

For finite intervals  $[a, b]$ , when  $b - a$  is very small so  $\bar{\Phi}(b)/\bar{\Phi}(a)$  is close to 1, the uniform proposal wins and the `RejectTail` variants are very slow. See right pane of Table 2. Precomputing the constants is also not useful for the uniform proposal. For larger intervals in the tail,  $\bar{\Phi}(x)$  decreases quickly at the beginning of the interval and this leads to very low acceptance ratios; see right pane of Table 1 and left pane of Table 2. A Rayleigh proposal with the `RejectTail` option is usually the fastest method in this case. Precomputing and storing the constants is also not very useful for this option. For intervals closer to the center, as in the left pane of Table 1, the uniform proposal performs well for larger (but not too large) intervals, and the `RejectTail` option becomes slower unless  $[a, b]$  is very wide. The reason is that for a fixed  $w > 0$ ,  $\bar{\Phi}(a + w)/\bar{\Phi}(a)$  is larger (closer to 1) when  $a > 0$  is closer to 0.

### 3 Simulation from Tail of Multivariate Normal

Let  $\phi_{\Sigma}(\mathbf{y})$  and

$$\bar{\Phi}_{\Sigma}(\mathbf{a}) = \mathbb{P}[\mathbf{Y} \geq \mathbf{a}], \quad \mathbf{Y} \sim \mathbf{N}(\mathbf{0}, \Sigma),$$

denote the density and tail distribution, respectively, of the multivariate  $\mathbf{N}(\mathbf{0}, \Sigma)$  distribution with (positive-definite) covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$ . In the multivariate extension to (1), we wish to simulate from the pdf ( $\mathbb{I}\{\cdot\}$  is the indicator function):

$$\frac{\phi_{\Sigma}(\mathbf{y}) \mathbb{I}\{\mathbf{y} \geq \mathbf{a}(\gamma)\}}{\bar{\Phi}_{\Sigma}(\mathbf{a}(\gamma))}, \quad (9)$$

where  $\max_i a_i > 0$ , and  $\gamma$  is a *tail parameter* such that at least one component of  $\mathbf{a}(\gamma)$  diverges to  $\infty$  (that is,  $\lim_{\gamma \uparrow \infty} \|\mathbf{a}(\gamma)\| = \infty$ , see [23]). To simulate from this conditional density, we describe an acceptance-rejection algorithm that uses an optimally designed multivariate exponential proposal. Interestingly, unlike the truncated exponential proposal in the one-dimensional setting (see Algorithms 2 and 3), our multivariate exponential proposal is not truncated. Before giving the details of the acceptance-rejection algorithm, we need to introduce some preliminary theory and notation.

### 3.1 Preliminaries and Notation

Define  $\mathbf{P}$  as a permutation matrix, which maps  $(1, \dots, d)^\top$  into the permutation vector  $\mathbf{p} = (p_1, \dots, p_d)^\top$ , that is,  $\mathbf{P}(1, \dots, d)^\top = \mathbf{p}$ . Then,  $\bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) = \mathbb{P}(\mathbf{P}\mathbf{Y} \geq \mathbf{P}\mathbf{a}(\gamma))$  and  $\mathbf{P}\mathbf{Y} \sim \mathbf{N}(\mathbf{0}, \mathbf{P}\Sigma\mathbf{P}^\top)$  for any  $\mathbf{p}$ . We will specify  $\mathbf{p}$  shortly.

First, define the constrained (convex) quadratic optimization:

$$\begin{aligned} \min_{\mathbf{y}} \quad & \frac{1}{2} \mathbf{y}^\top (\mathbf{P}\Sigma\mathbf{P}^\top)^{-1} \mathbf{y} \\ \text{subject to: } & \mathbf{y} \geq \mathbf{P}\mathbf{a}(\gamma) \end{aligned} \quad (10)$$

Suppose  $\boldsymbol{\lambda} \in \mathbb{R}^d$  is the Lagrange multiplier vector, associated with (10). Partition the vector as  $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_1^\top, \boldsymbol{\lambda}_2^\top)^\top$  with  $\dim(\boldsymbol{\lambda}_1) = d_1$  and  $\dim(\boldsymbol{\lambda}_2) = d_2$ , where  $d_1 + d_2 = d$ . In the same way, partition vectors  $\mathbf{y}$ ,  $\mathbf{a}$ , and matrix

$$\Sigma = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \quad (11)$$

We now observe that we can select the permutation vector  $\mathbf{p}$  and the corresponding matrix  $\mathbf{P}$  so that all the  $d_1$  active constraints in (10) correspond to  $\boldsymbol{\lambda}_1 > \mathbf{0}$  and all the  $d_2$  inactive constraints correspond to  $\boldsymbol{\lambda}_2 = \mathbf{0}$ . Without loss of generality, we can thus assume that  $\mathbf{a}$  and  $\Sigma$  are reordered via the permutation matrix  $\mathbf{P}$  as a pre-processing step. After this preprocessing step, the solution  $\mathbf{y}^*$  of (10) with  $\mathbf{P} = \mathbf{I}$  will satisfy  $\mathbf{y}_1^* = \mathbf{a}_1$  (active constraints:  $\boldsymbol{\lambda}_1 > \mathbf{0}$ ) and  $\mathbf{y}_2^* > \mathbf{a}_2$  (inactive constraints:  $\boldsymbol{\lambda}_2 = \mathbf{0}$ ).

We also assume that for large enough  $\gamma$ , the active constraint set of (10) becomes independent of  $\gamma$ , see [23]. An example is given in Corollary 1 below.

### 3.2 Acceptance-Rejection Algorithm

First, we note that simulating  $\mathbf{Y}$  from (9) is equivalent to simulating  $\mathbf{X} \sim \mathbf{N}(-\mathbf{a}(\gamma), \Sigma)$ , conditional on  $\mathbf{X} \geq \mathbf{0}$ , and then delivering  $\mathbf{Y} = \mathbf{X} + \mathbf{a}$ . Thus, our initial goal is to simulate from the target:

$$\pi(\mathbf{x}) = \phi_\Sigma(\mathbf{x} + \mathbf{a}(\gamma)) \mathbb{I}\{\mathbf{x} \geq \mathbf{0}\} / \bar{\Phi}_\Sigma(\mathbf{a}(\gamma))$$

Second, the partitioning into active and inactive constraints of (10) suggests the following proposal density:  $g(\mathbf{x}; \boldsymbol{\eta}) = g_1(\mathbf{x}_1; \boldsymbol{\eta}) g_2(\mathbf{x}_2 | \mathbf{x}_1)$ ,  $\boldsymbol{\eta} > \mathbf{0}$ , where

$$g_1(\mathbf{x}_1; \boldsymbol{\eta}) = \exp(-\boldsymbol{\eta}^\top \mathbf{x}_1) \prod_{k=1}^{d_1} \eta_k, \quad \mathbf{x}_1 \geq \mathbf{0}$$

is a multivariate exponential proposal, and

$$g_2(\mathbf{x}_2 | \mathbf{x}_1) = \phi_\Sigma(\mathbf{x} + \mathbf{a}) / \phi_{\Sigma_{11}}(\mathbf{x}_1 + \mathbf{a}_1)$$

is the multivariate normal pdf of  $\mathbf{x}_2$ , conditional on  $\mathbf{x}_1$  (see [11, Page 146]):

$$\mathbf{X}_2 | (\mathbf{X}_1 = \mathbf{x}_1) \sim \mathbf{N}(-\mathbf{a}_2 + \Sigma_{12}^\top \Sigma_{11}^{-1}(\mathbf{x}_1 + \mathbf{a}_1), \Sigma_{22} - \Sigma_{12}^\top \Sigma_{11}^{-1} \Sigma_{12})$$

With this proposal, the likelihood ratio for acceptance-rejection is:

$$\frac{\pi(\mathbf{x}) \bar{\Phi}_\Sigma(\mathbf{a}(\gamma))}{g(\mathbf{x}; \boldsymbol{\eta})} = \mathbb{I}\{\mathbf{x} \geq \mathbf{0}\} \frac{\phi_{\Sigma_{11}}(\mathbf{x}_1 + \mathbf{a}_1)}{g_1(\mathbf{x}_1; \boldsymbol{\eta})} = \mathbb{I}\{\mathbf{x} \geq \mathbf{0}\} \exp(\psi(\mathbf{x}_1; \boldsymbol{\eta})),$$

where  $\psi$  is defined as:

$$\psi(\mathbf{x}_1; \boldsymbol{\eta}) := -\frac{(\mathbf{x}_1 + \mathbf{a}_1)^\top \Sigma_{11}^{-1}(\mathbf{x}_1 + \mathbf{a}_1)}{2} + \boldsymbol{\eta}^\top \mathbf{x}_1 - \sum_{k=1}^{d_1} \ln(\eta_k) - \frac{\ln |\Sigma_{11}|}{2} - \frac{d_1 \ln(2\pi)}{2}$$

Next, our goal is to select the value for  $\boldsymbol{\eta}$  that will maximize the acceptance rate of the resulting acceptance-rejection algorithm (see Algorithm 7 below).

It is straightforward to show that, with the given proposal density, the acceptance rate for a fixed  $\boldsymbol{\eta} > \mathbf{0}$  is given by

$$\bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) \exp(-\max_{\mathbf{x}_1 \geq \mathbf{0}} \psi(\mathbf{x}_1; \boldsymbol{\eta}))$$

Hence, to maximize the acceptance rate, we minimize  $\max_{\mathbf{x}_1 \geq \mathbf{0}} \psi(\mathbf{x}_1; \boldsymbol{\eta})$  with respect to  $\boldsymbol{\eta}$ . In order to compute the minimizing  $\boldsymbol{\eta}$ , we exploit a few of the properties of  $\psi$ .

The most important property is that  $\psi$  is concave in  $\mathbf{x}_1$  for every  $\boldsymbol{\eta}$ , and that  $\psi$  is convex in  $\boldsymbol{\eta}$  for every  $\mathbf{x}_1$ . Moreover,  $\psi$  is continuously differentiable in  $\boldsymbol{\eta}$ , and we have the saddle-point property (see [3]):

$$\min_{\boldsymbol{\eta} > \mathbf{0}} \max_{\mathbf{x}_1 \geq \mathbf{0}} \psi(\mathbf{x}_1; \boldsymbol{\eta}) = \max_{\mathbf{x}_1 \geq \mathbf{0}} \min_{\boldsymbol{\eta} > \mathbf{0}} \psi(\mathbf{x}_1; \boldsymbol{\eta}) \quad (12)$$

Let  $\psi^* = \psi(\mathbf{x}_1^*; \boldsymbol{\eta}^*)$  denote the optimum of the minimax optimization (12) at the solution  $\mathbf{x}_1^*$  and  $\boldsymbol{\eta}^*$ . The right-hand-side of (12) suggests a method for computing  $\boldsymbol{\eta}^*$ , namely, we can first minimize with respect to  $\boldsymbol{\eta}$  (this gives  $\boldsymbol{\eta} = \mathbf{1}/\mathbf{x}_1$ , where the vector division is componentwise), and then maximize over  $\mathbf{x}_1 \geq \mathbf{0}$ . This yields the concave (unconstrained) optimization program for  $\mathbf{x}_1^*$ :

$$\mathbf{x}_1^* = \operatorname{argmax} \left\{ -\frac{(\mathbf{x}_1 + \mathbf{a}_1)^\top \Sigma_{11}^{-1}(\mathbf{x}_1 + \mathbf{a}_1)}{2} + \sum_{k=1}^{d_1} \ln x_k \right\}, \quad (13)$$

It then follows that  $\boldsymbol{\eta}^* = \mathbf{1}/\mathbf{x}_1^*$ . In summary, we have the following algorithm for simulation from (9).

---

**Algorithm 7** :  $\mathbf{X} \sim \mathbf{N}(\mathbf{0}, \Sigma)$ , conditional on  $\mathbf{X} \geq \mathbf{a}(\gamma)$  for large  $\gamma$ .

---

Solve (10) with  $\mathbf{P} = \mathbf{I}$  and compute the associated Lagrange multiplier  $\boldsymbol{\lambda}$ . Using  $\boldsymbol{\lambda}$ , construct the reordering (permutation) matrix  $\mathbf{P}$ , if needed.  
 $\mathbf{a} \leftarrow \mathbf{P}\mathbf{a}$   
3:  $\Sigma \leftarrow \mathbf{P}\Sigma\mathbf{P}^\top$   
Let  $\mathbf{L}$  be the lower triangular Cholesky factor of  $\Sigma_{22} - \Sigma_{12}^\top \Sigma_{11}^{-1} \Sigma_{12}$ , see (11)  
Solve the concave optimization problem (13) to obtain  $\mathbf{x}_1^*$ .  
6:  $\boldsymbol{\eta}_1^* \leftarrow \mathbf{1}/\mathbf{x}_1^*$   
 $\psi^* \leftarrow \psi(\mathbf{x}_1^*; \boldsymbol{\eta}^*)$   
**repeat**  
9: **repeat**  
Simulate  $U_0, U_1, \dots, U_{d_1} \sim \mathbf{U}(0, 1)$ , independently.  
 $E_k \leftarrow -\ln(U_k)/\eta_k^*$  for  $k = 1, \dots, d_1$   
12:  $\mathbf{X}_1 \leftarrow (E_1, \dots, E_{d_1})^\top$  {simulate  $\mathbf{X}_1 \sim g_1(\mathbf{x}_1; \boldsymbol{\eta}^*)$ }  
 $E \leftarrow -\ln(U_0)$   
**until**  $E > \psi^* - \psi(\mathbf{X}_1; \boldsymbol{\eta}^*)$   
15:  $\mathbf{Z}_2 \leftarrow (Z_1, \dots, Z_{d_2})^\top$ , where  $Z_1, \dots, Z_{d_2} \sim \mathbf{N}(0, 1)$ , independently.  
 $\mathbf{X}_2 \leftarrow \mathbf{L}\mathbf{Z}_2 - \mathbf{a}_2 + \Sigma_{12}^\top \Sigma_{11}^{-1}(\mathbf{X}_1 + \mathbf{a}_1)$  {simulate  $\mathbf{X}_2 \sim g_2(\mathbf{x}_2 | \mathbf{X}_1)$ }  
**until**  $\mathbf{X}_2 \geq \mathbf{0}$   
18:  $\mathbf{X} \leftarrow \mathbf{X} + \mathbf{a}$  {shift to obtain draw from pdf (9)}  
 $\mathbf{X} \leftarrow \mathbf{P}^\top \mathbf{X}$  {reverse reordering, if any}  
**return**  $\mathbf{X}$

---

### 3.3 Asymptotic Efficiency

The acceptance rate of Algorithm 7 above is

$$\mathbb{P}_g[E > \psi^* - \psi(\mathbf{X}_1; \boldsymbol{\eta}); \mathbf{X}_2 \geq \mathbf{0}] = \bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) \exp(-\psi^*),$$

where  $\mathbb{P}_g$  indicates that  $\mathbf{X}$  was drawn from the proposal  $g(\mathbf{x}; \boldsymbol{\eta}^*)$ . As in the one-dimensional case, see (8), it is of interest to find out how this rate depends on the tail parameter  $\gamma$ . In particular, if the acceptance rate decays to zero rapidly as  $\gamma \uparrow \infty$ , then Algorithm 7 will not be a viable algorithm for simulation from the tail of the multivariate Gaussian. Fortunately, the following result asserts that the acceptance rate does not decay to zero as we move further and further into the tail of the Gaussian.

**Theorem 1 (Asymptotically Bounded Acceptance Rate)** *Let  $\mathbf{y}^*$  be the solution to (10) after any necessary reordering via permutation matrix  $\mathbf{P}$ . Define  $\mathbf{a}_\infty := \lim_{\gamma \uparrow \infty} (\mathbf{a}_2(\gamma) - \mathbf{y}_2^*(\gamma))$  with  $\mathbf{a}_\infty \leq \mathbf{0}$ . Then, the acceptance rate of the accept-reject Algorithm 7 is ultimately bounded from below:*

$$\liminf_{\gamma \uparrow \infty} \bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) \exp(-\psi^*(\gamma)) \geq \mathbb{P}[\mathbf{Y}_2 \geq \mathbf{a}_\infty \mid \mathbf{Y}_1 = \mathbf{0}],$$

where the probability  $\mathbb{P}[\mathbf{Y}_2 \geq \mathbf{a}_\infty \mid \mathbf{Y}_1 = \mathbf{0}]$  is calculated under the original measure (that is,  $\mathbf{Y} \sim \mathbf{N}(\mathbf{0}, \Sigma)$ ) and, importantly, does not depend on  $\gamma$ .

*Proof* First, note that with the assumptions and notation of Section 3.1, Hashorva and Hüsler [23] have shown the following:

$$\bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) = \frac{\mathbb{P}[\mathbf{Y}_2 \geq \mathbf{a}_\infty \mid \mathbf{Y}_1 = \mathbf{0}]}{(2\pi)^{d_1/2} |\Sigma_{11}|^{1/2} \prod_{k=1}^{d_1} \mathbf{e}_k^\top \Sigma_{11}^{-1} \mathbf{a}_1} \exp\left(-\frac{\mathbf{a}_1^\top \Sigma_{11}^{-1} \mathbf{a}_1}{2}\right) (1 + o(1)), \quad \gamma \uparrow \infty,$$

where  $\mathbf{e}_k$  is the unit vector with a 1 in the  $k$ -th position, and  $f(x) = o(g(x))$  stands for  $\lim_{x \rightarrow a} f(x)/g(x) = 0$ .

Second, the saddle-point property (12) implies the following sequence of inequalities for any arbitrary  $\boldsymbol{\eta}$ :  $\psi^* \leq \psi(\mathbf{x}_1^*; \boldsymbol{\eta}) \leq \max_{\mathbf{x}_1} \psi(\mathbf{x}_1; \boldsymbol{\eta})$ . In particular, when  $\boldsymbol{\eta} = \Sigma_{11}^{-1} \mathbf{a}_1$ , then  $\max_{\mathbf{x}_1} \psi(\mathbf{x}_1; \Sigma_{11}^{-1} \mathbf{a}_1) = \psi(\mathbf{0}; \Sigma_{11}^{-1} \mathbf{a}_1)$ , and we obtain:

$$\exp(-\psi^*) \geq \exp(-\psi(\mathbf{0}; \Sigma_{11}^{-1} \mathbf{a}_1)) = \frac{\prod_{k=1}^{d_1} \mathbf{e}_k^\top \Sigma_{11}^{-1} \mathbf{a}_1}{\phi_{\Sigma_{11}}(\mathbf{a}_1)}$$

Therefore,  $\bar{\Phi}_\Sigma(\mathbf{a}(\gamma)) \exp(-\psi^*) \geq \mathbb{P}[\mathbf{Y}_2 \geq \mathbf{a}_\infty \mid \mathbf{Y}_1 = \mathbf{0}] (1 + o(1))$  as  $\gamma \uparrow \infty$ , and the result of the theorem follows.  $\blacksquare$

As a special case, we consider the asymptotic result of Savage [24]:

$$\frac{\bar{\Phi}_\Sigma(\gamma \Sigma \mathbf{c})}{\phi_\Sigma(\gamma \Sigma \mathbf{c})} = \frac{1}{\gamma^d \prod_{k=1}^d c_k} (1 + o(1)), \quad \mathbf{c} > \mathbf{0}, \quad \gamma \uparrow \infty, \quad (14)$$

which is the multivariate extension of the one-dimensional Mills' ratio [25]:  $\frac{\bar{\Phi}(\gamma)}{\phi(\gamma)} = \frac{1}{\gamma} (1 + o(1))$ . Interestingly, the following corollary shows that when the tail is of the Savage-Mills type, the acceptance probability not only remains bounded away from zero, but approaches unity.

**Corollary 1 (Acceptance with Probability One.)** *The acceptance rate of Algorithm 7 for simulation from (9) with  $\mathbf{a} = \gamma \Sigma \mathbf{c}$  for some  $\mathbf{c} > \mathbf{0}$  satisfies:*

$$\lim_{\gamma \uparrow \infty} \bar{\Phi}_{\Sigma}(\gamma \Sigma \mathbf{c}) \exp(-\psi^*(\gamma)) = 1$$

*Proof* Straightforward computations show that the Lagrange multiplier of (10) (with  $\mathbf{P} = \mathbf{I}$ , the identity matrix) is  $\boldsymbol{\lambda} = \Sigma^{-1} \mathbf{a} = \gamma \mathbf{c} > \mathbf{0}$ , so that the set of inactive constraints is empty. Then, repeating the argument in Theorem 1:

$$\exp(-\psi^*) \geq \exp(-\psi(\mathbf{0}, \gamma \mathbf{c})) = \frac{\gamma^d \prod_{k=1}^d c_k}{\phi_{\Sigma}(\gamma \Sigma \mathbf{c})} \stackrel{(14)}{=} \frac{1+o(1)}{\bar{\Phi}_{\Sigma}(\gamma \Sigma \mathbf{c})}, \quad \text{as desired.} \quad \blacksquare$$

As a numerical example, we used Algorithm 7 to simulate  $10^3$  random vectors from (9) for  $d = 10$ ,  $\mathbf{a} = \gamma \mathbf{1}$ , and  $\Sigma = \frac{9}{10} \mathbf{1}\mathbf{1}^{\top} + \frac{1}{10} \mathbf{I}$  (strong positive correlation) for a range of large values of  $\gamma$ .

Table 3 below reports the acceptance rate, estimated by observing the proportion of rejected proposals in line 17 of Algorithm 7, for a range of different  $\gamma$ .

**Table 3** Estimates of the acceptance probability,  $\bar{\Phi}_{\Sigma}(\mathbf{a}(\gamma)) \exp(-\psi^*)$ , as a function of  $\gamma$ .

$\gamma$	10	15	20	25	30	50	100	$10^3$
accept. rate	0.009	0.04	0.0815	0.15	0.19	0.34	0.44	0.50

The table confirms that as  $\gamma$  gets larger, the acceptance rate improves.

## 4 Conclusion

We have proposed and tested both inversion and acceptance-rejection methods to generate a standard normal, truncated to an interval  $[a, b]$ , when  $a \gg 0$ . We have also proposed an acceptance-rejection method for the tail of the multivariate normal distribution.

In the univariate setting, inversion is slower than the fastest rejection method, as expected. However, inversion is still desirable in many situations. Our new inversion method excels in those situations when  $a$  is large (say,  $a \geq 10$ ). For  $a$  not too large (say,  $a \leq 30$ ), the accurate approximation of [2] implemented in InversionSSJ works well.

When inversion is not needed, the rejection method with the Rayleigh proposal is usually the fastest when  $a$  is large enough.

It is interesting to see that, in the univariate setting, using the Rayleigh proposal is faster than using the truncated exponential proposal as in [6, 8, 21]. However, in the multivariate setting, we show that the *truncated* exponential method of [6, 8, 21] can be extended to help simulate from the multivariate normal tail, provided that we use an *untruncated* multivariate exponential proposal (that is,  $\mathbf{X} \geq \mathbf{0}$ ) combined with a Gaussian mean-shift (that is,  $\mathbf{Y} = \mathbf{X} + \mathbf{a}$ ).

## References

1. M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. Dover, New York, 1970.
2. J. M. Blair, C. A. Edwards, and J. H. Johnson. Rational Chebyshev approximations for the inverse of the error function. *Mathematics of Computation*, 30:827–830, 1976.
3. Z. I. Botev. The normal law under linear restrictions: simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2016. doi: 10.1111/rssb.12162.
4. Z. I. Botev and P. L’Ecuyer. Efficient estimation and simulation of the truncated multivariate Student-t distribution. In *Proceedings of the 2015 Winter Simulation Conference*, pages 380–391. IEEE Press, 2015.
5. Z. I. Botev, M. Mandjes, and A. Ridder. Tail distribution of the maximum of correlated Gaussian random variables. In *Proceedings of the 2015 Winter Simulation Conference*, pages 633–642. IEEE Press, 2015.
6. N. Chopin. Fast simulation of truncated Gaussian distributions. *Statistics and Computing*, 21(2):275–288, 2011.
7. L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, New York, NY, 1986.
8. J. Geweke. Efficient simulation of the multivariate normal and Student-t distributions subject to linear constraints and the evaluation of constraint probabilities. In *Computing science and statistics: Proceedings of the 23rd symposium on the interface*, pages 571–578, Fairfax, Virginia, 1991.
9. C. Hans. Model uncertainty and variable selection in Bayesian lasso regression. *Statistics and Computing*, 20(2):221–229, 2010.
10. W. Hörmann, J. Leydold, and G. Derflinger. *Automatic Nonuniform Random Variate Generation*. Springer-Verlag, Berlin, 2004.
11. D. P. Kroese, T. Taimre, and Z. I. Botev. *Handbook of Monte Carlo Methods*. John Wiley and Sons, New York, 2011.
12. P. L’Ecuyer. Variance reduction’s greatest hits. In *Proceedings of the 2007 European Simulation and Modeling Conference*, pages 5–12, Ghent, Belgium, 2007. EUROSIS.
13. P. L’Ecuyer. Quasi-Monte Carlo methods with applications in finance. *Finance and Stochastics*, 13(3):307–349, 2009.
14. P. L’Ecuyer. Random number generation with multiple streams for sequential and parallel computers. In *Proceedings of the 2015 Winter Simulation Conference*, pages 31–44. IEEE Press, 2015.
15. P. L’Ecuyer. SSJ: Stochastic simulation in Java, software library, 2016. <http://simul.iro.umontreal.ca/ssj/>.
16. P. L’Ecuyer and G. Perron. On the convergence rates of IPA and FDC derivative estimators. *Operations Research*, 42(4):643–656, 1994.
17. J. Leydold. *UNU.RAN—Universal Non-Uniform RANdom number generators*, 2009. Available at <http://statmath.wu.ac.at/unuran/>.
18. G. Marsaglia. Generating a variable from the tail of the normal distribution. *Technometrics*, 6(1):101–102, 1964.
19. G. Marsaglia and T. A. Bray. A convenient method for generating normal variables. *SIAM Review*, 6:260–264, 1964.
20. G. Marsaglia, A. Zaman, and J. C. W. Marsaglia. Rapid evaluation of the inverse normal distribution function. *Statistics and Probability Letters*, 19:259–266, 1994.
21. C. P. Robert. Simulation of truncated normal variables. *Statistics and computing*, 5(2):121–125, 1995.
22. D. B. Thomas, W. Luk, P. H. Leong, and J. D. Villasenor. Gaussian random number generators. *ACM Computing Surveys*, 39(4):Article 11, Nov. 2007.
23. E. Hashorva and J. Hüsler. On multivariate Gaussian tails. *Annals of the Institute of Statistical Mathematics*, 55(3):507–522, 2003.
24. R. I. Savage. Mills’ ratio for multivariate normal distributions. *J. Res. Nat. Bur. Standards Sect. B*, 66:93–96, 1962.
25. J. P. Mills. Table of the ratio: area to bounding ordinate, for any portion of normal curve. *Biometrika*, pages, 395–400, 1926