SCALABLE RANDOM NUMBER GENERATION FOR TRUNCATED LOG-CONCAVE DISTRIBUTIONS

Michele Lambardi di San Miniato
Department of Economics and Statistics
University of Udine
Udine, Italy
michele.lambardi@uniud.it

Euloge Clovis Kenne Pagui
Department of Statistical Sciences
University of Padova
Padova, Italy
kenne@stat.unipd.it

ABSTRACT

Inverse transform sampling is an exceptionally general method to generate non-uniform-distributed random numbers, but can be rather unstable when simulating extremely truncated distributions. Many famous probability models share a property called log-concavity, which is not affected by truncation, so they can all be simulated via rejection sampling using Devroye's approach. This sampler is based on rejection and thus more stable than inverse transform, and uses a very simple envelope whose acceptance rate is guaranteed to be at least 20%. The aim of this paper is threefold: firstly, to warn against the risk of wrongly simulating from truncated distributions; secondly, to motivate a more extensive use of rejection sampling to mitigate the issues; lastly, to motivate Devroye's automatic method as a practical standard in the case of log-concave distributions. We illustrate the proposal by means of simulations based on some Tweedie distributions, for their relevance in regression analysis.

Keywords Rejection sampling · Inverse transform sampling · Log-concave distributions · Random number generation · Scalable simulations · Truncated distributions

1 Introduction

We consider the problem of non-uniform random number generation, when the aim is to simulate a univariate, log-concave, and extremely truncated probability distribution. Formally, let X be a random variable with support $\mathcal{X} \subseteq \mathbb{R}$ and probability mass or density function $f(x)$, and let $\mathcal{I} \subset \mathcal{X}$ be such that $\mathbb{P}(X \in \mathcal{I}) > 0$. The distribution is said *log-concave* if $\log f(x)$ is a concave function. Moreover, the distribution of X *truncated to* $\mathcal{I}$ can be defined as the distribution of X conditional to $X \in \mathcal{I}$ and denoted by $f_{\mathcal{I}}(x)$; henceforth, $f(x)$ is referred to as the *base* distribution of X. It is possible to define $f_{\mathcal{I}}(x)$ as

$$f_{\mathcal{I}}(x) = f(x)/P(\mathcal{I}), \quad x \in \mathcal{I},$$

where $P(\mathcal{I}) = \mathbb{P}(X \in \mathcal{I})$. Samples of $f_{\mathcal{I}}(x)$ can be drawn and re-drawn from $f(x)$ until they lie in $\mathcal{I}$, as per hit-or-miss. We consider the case when $\mathcal{I} = \mathcal{X} \cap]a, +\infty[$ for some $a \in \mathbb{R}$, but the software we provide supports truncation over intervals in the form $\mathcal{I} = \mathcal{X} \cap]a, b]$, as typical in the other R software.

The condition of interest is called "extreme truncation" in the following and occurs when $P(\mathcal{I}) \approx 0$, thus as a approaches $\sup \mathcal{X}$. Extreme truncation makes hit-or-miss inefficient at simulating $f_{\mathcal{I}}(x)$, because only a fraction $P(\mathcal{I})$ of samples from $f(x)$ is accepted on average, so a workaround is in order. The inverse transform method is a popular sampling algorithm, but it can be unreliable under extreme truncation [1]. Rejection sampling is more trustworthy [2] but cannot be deployed as straightforwardly in general. Devroye proposed an automatic rejection sampling method whose efficiency is uniformly lower-bounded over all the eligible distributions, which are the log-concave ones [3, 4].

An idea flashes through one's mind when noticing that many classical distributions, including Tweedie ones, are log-concave and therefore so are truncations of these distributions [5]. In this paper, we analyze some risks in simulating truncated distributions via inverse transform. We show via simulations that Devroye's method can mitigate the issues at least with the eligible distributions, as it breaks down only under far more extreme truncation than inverse transform does. As a remark, the problem is not overcome at all, but only mitigated. We restrict the analysis to Tweedie distributions,

which are of immediate relevance for regression analysis tasks, but our conclusions may also apply to other log-concave distributions as well.

We provide an R package that implements Devroye's sampler for truncated log-concave distributions, not just Tweedie ones; the package only requires to specify a base probability mass or density function $f(x)$, the base cumulative distribution function $F(x) = \mathbb{P}(X \leq x)$, and a base mode function that returns the mode of the base distribution $m = \arg\sup_{x \in \mathcal{X}} f(x)$. All such functions can be consistently parameterized, like for instance the Poisson distribution with mean $\lambda > 0$. Moreover, both $f(x)$ and $F(x)$ must be available in logarithmic form for meaningful evaluation in the case of extreme truncation.

The article is organized as follows. In Section 2, we review concepts like truncated and log-concave distributions. In Section 3, we recall some additional definitions involved in the inverse transform and Devroye sampling methods. In Section 5, a simulation study is presented for each implemented distribution. In Section 6 we validate the diagnostic tools provided by our package. Conclusions are drawn in Section 7.

2 Background

Truncation is a common stratagem in statistical modeling of tail events. For instance, in risk prevention, the behavior of a system is analyzed conditional to rare and disruptive events [7, 8]. Related to truncation, though distinct, are censoring [9], binning and quantization [10]. Each probability distribution comes with a built-in tail behavior that implies an assumption about censoring [11]. Truncated distributions are involved in process start-up, inventory control [12] and masking microdata [13]. Truncation allows to expand modeling possibilities, by manipulating known distributions [14]. For instance, stochastic frontier analysis [15, 16] involves positive-valued random variables, defined by truncating either normal, Laplace or other log-concave distributions [17, 18]. The truncated normal is especially popular [19, 20, 21, 22], followed by truncated binomial and Poisson distributions for zero-inflated modeling [23]. Truncated distributions are also needed to assess some robust estimators [24].

Extremely truncated distributions are especially challenging to simulate. In normal applications, stepping into them is considered a mistake, though, which motivates an ostrich approach. For instance, iterative sampling of Bayesian posterior distributions can be troubled by model misspecification or poor initialization [2], so one can just step into parameter configurations that imply extremely truncated distributions, but some trial and error will suffice to get the analysis back on track. Other errors are more substantial than merely computational. The 2008 financial crisis was many sigma away from normality according to some light tail models in use at the time [25]. Nonetheless, there are occasions when simulating extremely truncated distributions is a necessity more than a mistake. The software should be able to accommodate for legitimate prompts and make reliable simulations accordingly.

In applications, interest lies often on multivariate distributions that are analytically intractable. Simulation and Monte Carlo methods are just one option to speed up the analysis and sometimes even the only way out of it [26]. Often, in Bayesian inference, the posterior distribution can be summarized only based on samples [2]. In addition, bootstrap is frequently resorted to in approximating standard errors associated with frequentist estimates [27, 28, 29]. It is thus crucial that random number generators provide samples that reflect the desired distributions, but this aspect should not be taken for granted. In R, many simulation-oriented packages [30, 31] rely on the inverse transform sampling method, which fails at simulating extreme tail events [1]. Breakdowns can be due to inaccuracies [32] and algorithmic limitations [33] that affect the evaluation of probability-related functions. Our aim here is to make it possible to carry out reliable simulations within standard finite-precision arithmetic, where the choice of algorithms is still crucial [34].

Truncated models are higher-dimensional than their base counterparts, since ruled by the additional parameters defining $\mathcal{I}$, which makes diagnostics and prevention even more challenging. This is also why we restrict our discussion to truncation in the specific form used throughout this article, which is still relevant in practice. More complicated truncations can be defined in terms of mixtures. The scope of our discussion is further restricted, as we address the case of log-concave distributions. All such distributions have necessarily light, sub-exponential, tails [35]. This feature makes tail probabilities especially small and hard to evaluate meaningfully. Tweedie distributions, involved in the generalized linear model, are all log-concave; only gamma with shape parameter less than unit is not log-concave [36], but it can be simulated indirectly based on a log-concave exponential power distribution [3]. Log-concavity implies a concentration inequality and an envelope (or "hat") distribution to perform rejection sampling, with execution time bounded in probability [3]. This sampler is very general, can be tuned automatically, has acceptance rate greater or equal than 20%, with eligible distributions, and is also numerically stable since based on rejection sampling [2]. The envelope is a mixture involving uniform and exponential distributions in the continuous case, uniform and geometric distributions in the discrete one, thus it is easy to simulate.

3   Sampling truncated distributions

In non-uniform random number generation [3], each sampling method needs some descriptor of the target distribution, like $f(x)$ in rejection sampling. Some descriptors can be troubled at times. For instance, consider the Laplace transform of the distribution [37]: it works as a distribution descriptor, since in one-to-one correspondence with $f(x)$, which is a descriptor itself, but is numerically troubled when the support $\mathcal{X}$ is bounded (this problem is known as Gibbs phenomenon). Sampling methods based on the Laplace transform are expected to be unreliable in such occasions. Other issues may affect other distribution descriptors as well, thus in random number generation it is crucial to know which descriptor fails and *when*. Inaccuracies and algorithmic limitations [33], which we broadly refer to as numerical issues, can affect many statistical tasks apart from random number generation [38, 39], but we address that issue in later sections.

Firstly, we recall few basic distribution descriptors for sampling tasks. In inverse transform sampling, the fundamental distribution descriptor is the quantile function, denoted by $q(p)$ and defined as the pseudo inverse of $F(x)$, that is,

$$q(p) = \inf\left\{x \in \mathbb{R} : F(x) \geq p\right\}, \quad p \in [0, 1].$$

When X is an absolutely continuous random variable, it holds $q(p) = F^{-1}(p)$. The quantile function $q_{\mathcal{I}}(p)$ of X truncated to $\mathcal{I}$ is defined as

$$q_{\mathcal{I}}(p) = q\{F(a) + p\,P(\mathcal{I})\}.$$

3.1   Inverse transform sampling

If $q(p)$ is tractable, simulating a distribution can be based on the inverse transform sampling method (ITS). This technique can be relatively safe in normal applications. When considering extremely truncated distributions, though, this method may fail severely, to an extent that depends also on the specific implementation of the relevant distribution descriptor. This technique is nonetheless a standard one in the absence of more targeted samplers, due to its unrivaled generality. If $q(p)$ can be evaluated efficiently, and if a uniform-distributed random variable U can be simulated, then also X can be simulated, via ITS. One has to exploit the fact that

$$X \sim q(U),$$

where "$\sim$" denotes equivalence in distribution. The ITS will first generate samples of U and then transform them via $q(p)$, in order to simulate samples of X.

For ITS to sample a truncated distribution, little modifications are needed, when $F(x)$ can be evaluated efficiently too. Consider X being truncated to $\mathcal{I}$. Then, U can be transformed via $q_{\mathcal{I}}(p)$ instead of $q(p)$, that is,

$$X \sim q_{\mathcal{I}}(U) \,|\, X \in \mathcal{I}.$$

ITS scales well *in theory*, but the actual scalability depends on software implementations of $q(p)$ and $F(x)$. In R software, as of version 4.1.1, the package `stats` cannot evaluate standard normal quantiles beyond nine sigma. More extreme quantiles will be automatically imputed as $+\infty$. The package `truncnorm` [40] instead imputes $q_{\mathcal{I}}(p) \approx a$, which is more meaningful after concentration inequalities for log-concave distributions, as discussed later. Truncating the standard normal to $]+9, +\infty[$ is thus not possible with ITS. Instead, $F(x)$, or, at least, its logarithm, can still be evaluated reliably far beyond nine sigma, so it is not responsible for the failure of ITS, which is on $q(p)$ alone.

3.2   Rejection sampling

The safety of ITS depends on the precision at which quantiles can be assessed. If more refined implementations of $q(p)$ become available, the safety of ITS can be improved. Other sampling techniques that depend on different distribution descriptors might have a qualitatively distinct breakdown behavior. Rejection sampling (RS) is an alternative simulation method, but it is less general than ITS and more tailored down to specific distributions.

One Markov Chain-Monte Carlo approach related to RS is called hit-or-miss. For simulating $f_{\mathcal{I}}$, samples can be drawn from $f(x)$ until they lie within $\mathcal{I}$. The acceptance rate is $P(\mathcal{I})$. In independent sampling, it takes on average $1/P(\mathcal{I})$ samples from f to generate a single sample from $f_{\mathcal{I}}$, the number of samples is geometric-distributed. This technique is thus inefficient under extreme truncation, because such a condition makes $P(\mathcal{I}) \approx 0$.

In RS, the generic distribution $f(x)$ is the target from which one wishes to draw samples, but a so-called envelope distribution must be found, denoted by $g(x)$, which must be (i) easier to simulate from and (ii) such that

$$f(x) \leq k\,g(x), \quad x \in \mathcal{X},$$

for some known constant $k \geq 1$. Here, $g(x)$ must be defined with respect to the same measure as $f(x)$, that is, the former must be either discrete or continuous in the same fashion as the latter. In general it will hold $k > 1$, while $k = 1$ holds when $g(x) = f(x)$, but $f(x)$ cannot be sampled directly by assumption. In RS terms, hit-or-miss targets $f_{\mathcal{I}}$ via an envelope f, which implies $k = 1/P(\mathcal{I})$. As a matter of efficiency, k should be small, as it takes k samples from $g(x)$ to generate just one sample from $f(x)$, on average. The RS requires to draw a sample, say x, from $g(x)$ and accept it with probability $p(x)$, defined as

$$p(x) = \frac{f(x)}{k \, g(x)} \in [0, 1] \,.$$

If the sample is rejected, a new one is drawn, and so on, until acceptance. The RS relies on the fact that, if $X \sim g(\cdot)$ and $U \sim \text{uniform}[0, 1]$ are independent, then the distribution of X conditional to $U \leq p(X)$ is $f(x)$.

ITS relies on $q(p)$ as a distribution descriptor, RS on $f(x)$. Under truncation, RS relies on $f_{\mathcal{I}}(x)$ and thus also on $F(x)$. Even so, RS must have a different breakdown behavior, depending on how the distribution descriptors are implemented. One can evaluate $q(p)$ by solving $F(x) = p$ for x; even so, some iterative procedure must be involved, which can be either cumbersome or unreliable in tails [37]. As to RS, a major drawback is that finding an envelope is rarely trivial, at least when efficiency is a concern. Typical envelopes are often efficient around the mode m, in the sense that $p(m) \approx 1$, but also heavy tailed and thus potentially inefficient if reused via hit-or-miss under extreme truncation. In such a case, an entirely different envelope must be devised. For instance, the R package `tmvtnorm` [41] features a hit-or-miss default sampler, which repeatedly draws samples from the non-truncated distribution until they fall within the target interval or region, thus may perform poorly in tails. The package also offers a Gibbs sampler as a replacement for such situations, as it is known to be effective even under complicated truncation [42]; the problem with this approach is though that it is not exact method and yields potentially correlated samples that call for inspection.

3.3 Devroye sampling

RS has the capability to outlive ITS under extreme truncation, because it relies on qualitatively different distribution descriptors, $f(x)$ and $F(x)$, which are numerically more stable in standard software. Yet, RS requires a smart enveloping strategy, while ITS works automatically. There exists, though, a wide range of probability models for which a rather general envelope works, with acceptance rate lower-bounded and greater than 20%. Some of the most famous distributions are log-concave [5]. To begin with, all Tweedie distributions, like binomial and normal, are log-concave. The gamma distribution is especially tricky [36] as it is not log-concave when the shape parameter is less than unity, but it can be simulated based on the exponential power distribution [3, Example 2.1]. Devroye provides two lists with some log-concave distributions for the continuous case [3] and the discrete one [4]. Also, Devroye addresses the ideal setting, where "real-value arithmetic can be performed with infinite accuracy" [43] but we show that, even in finite-precision calculus, DS is superior to ITS wherein applicable.

Formally, a probability mass or density function $f(x)$ is said to be log-concave if the function $\ell(x) = \log f(x)$ is concave. In the continuous case, the second derivative of $\ell(x)$ can be checked. Cusps like in the Laplace case are an exception. In the discrete case, it must hold

$$2\ell(x) \geq \ell(x - 1) + \ell(x + 1), \quad \text{for all} \quad x \in \mathcal{X} \,.$$

Now, let m be the mode of $f(x)$. This is a well-defined distribution summary in the case of log-concavity. Rejection sampling à-la Devroye (DS), in the continuous case [44], relies on the following notable inequality:

$$f(x) \leq f(m) \min \left[1, \exp \{ 1 - f(m) |x - m| \} \right] , \tag{1}$$

or, for the unconventionally standardized variable $Y = (X - m) \, f(m)$,

$$Y \sim f(y) \leq \min \{ 1, \exp(1 - |y|) \} \,.$$

This bound can be read in terms of a uniform-exponential mixture that can serve as an envelope for all log-concave densities, with acceptance rate exactly equal to 25%. The simplicity of the algorithm deserves noticing [45]. A concentration inequality is also implied, where (i) the mode turns out to be a meaningful central tendency index and (ii) the tails are necessarily sub-exponential [35]. In the discrete case, analogously, an enveloping uniform-geometric mixture can be found, with acceptance rate $\{4 + f(m)\}^{-1}$ and thus between 20% and 25%, the two bounds being attained only in the limit as $f(m)$ approaches either 1 or 0, respectively.

An appealing aspect of log-concavity is that it is preserved under truncation [5], so one can sample log-concave distributions via DS in both base and truncated versions. As a remark, even in the latter case, the acceptance rate is still at least 20%. The mode of the truncated distribution is obtained by projecting m onto $\mathcal{I}$, namely, $m \leftarrow \max\{a, m\}$ [46].

Algorithms 1 and 2 are for Devroye's sampler in the continuous case and the discrete one, respectively, for the generally useful case $\mathcal{I} = \mathcal{X} \cap]a, b]$. The two algorithms are collected here for completeness and adapted for coherence with our notation. One can refer to Devroye's work, in the specific, Rejection method for log-concave densities. Exponential version [3], and the second unnumbered algorithm from his work on the discrete case [4]. The most relevant shift in truncation is that $F(x)$ is involved in the normalization constant to the truncated distribution and acts as a potential source of numerical nuances. In the limit as $x \to \infty$ diverges, it holds $f(x), 1 - F(x) \to 0$, so in R software the logarithmic versions of these functions are often provided, which are numerically more reliable. The relative (instead of absolute) evaluation error of $f(x)$ and $1 - F(x)$ is bounded, by bounding the absolute error on their logarithms, which ensures a reasonable number of significant digits. As a general recommendation, such logarithmic version of $f(x)$ and $1 - F(x)$ should be implemented meaningfully in this respect.

Data: $m, a, b, f_{\cdot,\cdot}(\cdot)$.
$m \leftarrow \max\{a, \min(b, m)\}$;
$c \leftarrow f_{\mathcal{I}}(m)$;
repeat
    generate $U \sim \text{uniform}[0, 2], V \sim \text{exponential}(1)$;
    if $U \leq 1$ **then**
        $X \leftarrow U$;
        $Z \leftarrow -E$;
    else
        generate $E^* \sim \text{exponential}(1)$;
        $X \leftarrow 1 + E^*$;
        $Z \leftarrow -E - E^*$;
    end
    generate $S \sim \text{uniform}[-1, +1]$;
    $S \leftarrow \text{sign}(S)$;
    $X \leftarrow m + S\,X/c$;
until $\exp(Z) \leq f_{\mathcal{I}}(X)$;
return X;

Algorithm 1: DS for a truncated continuous distribution.

Data: $m, a, b, f_{\cdot,\cdot}(\cdot)$.
$a \leftarrow \lfloor a \rfloor$;
$b \leftarrow \lfloor b \rfloor$;
$m \leftarrow \max\{a + 1, \min(b, m)\}$;
$c \leftarrow f_{\mathcal{I}}(m)$;
$w \leftarrow 1 + c/2$;
repeat
    generate $U, W \sim \text{uniform}[0, 1]$;
    if $U \leq w/(1 + w)$ **then**
        generate $V \sim \text{uniform}[0, 1]$;
        $X \leftarrow w\,V/c$;
    else
        generate $E \sim \text{exponential}(1)$;
        $X \leftarrow (w + E)/c$;
    end
    generate $S \sim \text{uniform}[-1, +1]$;
    $S \leftarrow \text{sign}(S)$;
    $X \leftarrow S\,\text{round}(X)$;
until $W \min\{1, \exp(w - c\,X)\} \leq f_{\mathcal{I}}(m + X)/c$;
$X \leftarrow m + X$;
return X;

Algorithm 2: DS for a truncated discrete distribution.

When a sampler fails, variates can be imputed reasonably in the light of probability calculus. For log-concave distributions, an excess sequence Y_a can be defined as $Y_a \sim X - a | X > a$, whose both mean and variance are non-increasing with a [47, 48]. Imputing a truncated random variable as infinite is not a sound choice, since Y_a is

bounded in probability, but this is what happens with the R package `truncdist`. This pastiche only occurs because the package works as a wrapper and uses probability-related functions provided by other packages in turn, so it does not address log-concave distributions differently from any other. A more meaningful imputation would be the mode m, after the concentration inequality in (1). This latter kind of imputation is adopted, for instance, by the R package `truncnorm`. As an alternative, stochastic, imputation, one may consider exponential approximations at least for normal [49] and perhaps other sub-exponential tails.

3.4 Means of comparison

Typically, the functions $f(x)$, $F(x)$ and $q(p)$ will depend on some parameter vector $\theta \in \mathbb{R}^d$. For instance, in the binomial case, $\theta = (n, p)^\top$, where n is the size and p is the event probability. We should thus write $f(x; \theta)$, $F(x; \theta)$, $q(x; \theta)$ and $m(\theta)$ but, for the sake of readability, we refrain from doing so. Still, all considerations should be considered dependent on the parameter θ.

As per (1), it is useful to think of a base log-concave distribution in simpler terms, based on a central tendency index $\mu = \mu(\theta)$ and a dispersion parameter $\sigma = \sigma(\theta)$, defined as

$$\mu = m, \quad \sigma = 1/f(m).$$

In the normal-distributed case and other location and scale families, it is natural to define σ as the base standard deviation, thus only *proportional* to $1/f(m)$. This choice is meaningful also for other distributions that can be approximated to the normal, like binomial and Poisson ones.

With some distributions, and depending on their implementations, the breakdown of $q(p)$ precedes the one of $F(x)$, which in turn precedes the breakdown of $f(x)$. Call $\bar{a} = q(\bar{p})$ the maximum quantile that can be evaluated, while $F(x)$ is unreliable for $x > \bar{a}'$, and $f(x)$ for $x > \bar{a}''$. One can expect $\bar{a}' < \bar{a}''$ when $F(x)$ is implemented based on summation or numeric integration of $f(x)$. When $q(p)$ is implemented by solving $F(x) = p$ for x, one can expect $\bar{a} < \bar{a}'$.

All the distributions considered here share there being a maximum quantile or critical value, here denoted by $\bar{a}$, that can be computed reliably. More extreme quantiles are thence imputed equal to $\sup \mathcal{X}$ in some R packages. This behavior implies that simulating $f_{\mathcal{I}}$ is not possible when $a > \bar{a}$. With ITS, the safety ratio η is defined here as

$$\eta = \frac{\bar{a} - \mu}{\sigma},$$

which can be read as the maximum sigma of events that can be simulated reliably. If $a < \bar{a}$, a fraction $1 - F_{\mathcal{I}}(\bar{a})$ of samples from $f_{\mathcal{I}}(x)$ will be corrupted. When simulating n variates, the probability of stepping into corrupted samples is $1 - \{1 - F_{\mathcal{I}}(\bar{a})\}^n$. Sample corruption may thus occur even with non-extreme truncation, if n is sufficiently large; by converse, if n is small, issues will arise mostly only when $a \approx \bar{a}$, due to light tails. As a remark, η depends on θ, if the model is parameterized.

When generating n samples of $f_{\mathcal{I}}(x)$ with DS, the failure is either sure or impossible, as ruled by $\bar{a}'$, and essentially independent of n. With RS, we define the safety ratio η' as

$$\eta' = \frac{\bar{a}' - \mu}{\sigma},$$

without risk of overstating the algorithm's safety. We argue that DS allows for $\eta' > \eta$, that is, additional safety with respect to ITS.

There are at least two aspects in the definition of η that are worth discussing. Firstly, it is necessary to test for the distribution of generated samples to assess its trustworthiness, but testing is as tricky as random number generation itself, because both tasks are based on some distribution descriptor that is potentially troubled. Secondly, the standardization of $\bar{a}$ into η is meaningful at least in the normal case, since the evaluation of $f(x)$, $F(x)$ and $q(p)$ involves an internal standardization. The same definition still makes sense at least for some other distributions that can be approximated to the normal and that are treated this way in software, like Poisson and binomial.

Table 1 summarizes our definition of μ and σ along with the choice of R implementation of all the distributions we checked. In general, the package `Rmpfr` [32] was preferred over `stats`. Both lacked an implementation for the inverse Gaussian, which we retrieved from the R package `actuar` [50].

4 Software

The existing R package of interest is `truncdist`, which allows to simulate truncated distributions based on ITS, in the case when the function $q(p)$ can be evaluated feasibly [30]. In the event of extreme truncation or other situations making

Table 1: Distributions considered in simulations.

distribution	parameters	μ	σ	$f(x)$	$F(x)$	$q(p)$
binomial	n, p	np	$\sqrt{np(1-p)}$	Rmpfr	stats	stats
Poisson	λ	λ	$\sqrt{\lambda}$	Rmpfr	stats	stats
negative binomial	n, p	$\frac{np}{1-p}$	$\frac{\sqrt{np}}{1-p}$	Rmpfr	stats	stats
geometric	p	0	$\frac{\sqrt{1-p}}{p}$	stats	stats	stats
normal	μ, σ	μ	σ	Rmpfr	Rmpfr	stats
gamma	α, λ	$\frac{\alpha}{\lambda}$	$\frac{\sqrt{\alpha}}{\lambda}$	Rmpfr	stats	stats
inverse Gaussian	μ, λ	μ	$\sqrt{\frac{\mu^3}{\lambda}}$	actuar	actuar	actuar

$q(p)$ hard to evaluate, the R package truncnorm [40] replaces ITS with RS in simulating truncated normal distributions. Many other packages will address truncation via specifically tailored RS, but this approach is not automatic enough nor of interest here. The following packages are instead based on DS.

- UnivRNG [51] for Zipf's zeta distribution,
- extraDistr [52] for Dirichlet distribution,
- actuar [50] for inverse Gaussian distribution,
- heavy [53] for multivariate normal and t distributions.

We do not address the multivariate case here, as it likely deserves a more targeted set of strategies, the task being cursed by dimensionality [10, 46]. To our knowledge, the above R packages are the only ones available on the Comprehensive R Archive Network (CRAN) that simulate truncated distributions via DS. Even so, each of these packages addresses only one specific model. It seems thus that DS serves mostly as a placeholder in the absence of a more targeted sampler. Also, these package lack the purpose of addressing a wide variety of truncated distributions by leveraging on the scalability of DS.

We propose a new R package that we call truncLCdist as a hommage to truncdist's aspiration to simply work out of the box. Our package is based entirely on DS and can address all log-concave distributions, provided a definition of $f(x)$, $F(x)$ and m, consistently parameterized. An implementation of $\log f(x)$ must be provided: this can be a foo, which firstly evaluates $f(x)$ and then computes its logarithm or, as recommended, a direct evaluation of $\log f(x)$ that ensures a reasonably high number of significant digits [34]. The two steps evaluation involves firstly computing $f(x)$ and then transforming it via logarithm and may be unreliable at times: the two steps are optimized separately, without foreseeing each other, thus the overall evaluation of $\log f(x)$ is less accurate.

Tweedie distributions are particularly relevant for regression analyses and are included in our package as a default. The user can add custom support for other log-concave distributions just by linking the required functions. Such linkage must be *explicitly* set, via a function that wraps functionalities borrowed from the futile.options package [54]. The linkage also allows to check for issues with specific functions, by handily replacing their implementations. This cannot be done as easily with truncdist, where the linkage is *implicit*, that is, the package retrieves the required functions from the global environment and loaded packages based on names. One should first understand the way in which function masking and conflicts between packages are addressed in R, before being able to feed an alternative distribution descriptor to truncdist.

5 Simulations

Here we enumerate few simulation studies that can be repeated with our software in order to probe into the behavior of ITS and DS in practical random number generation.

5.1 Normal

As a starter, we address a simple case with a normal distribution, with mean μ and variance σ^2. One can consider the standard normal, with mean $\mu = 0$ and variance $\sigma^2 = 1$, without loss of generality.

For the sake of completeness, $F(x)$ translates well under rescaling since it only depends on the standardized input $z = (x - \mu)/\sigma$, while $f(x; \mu, \sigma) = f(z; 0, 1)/\sigma$ is also divided by σ and may thus be troubled in the limit as $\sigma \to 0$ or $\sigma \to \infty$. We assume here that this issue does not occur here and that σ takes non-extreme values, for the sake of simplicity. We defend ourselves with this remark as even our proposal may fail arbitrarily subtle tests, but our approach is still motivated by the usefulness of truncated distributions that are extreme in terms of $\mathcal{I}$.

We probe some integer values of a and look for its largest value $\overline{a}$ for which $q(p)$ is safe to evaluate. Here, it holds $\eta = \overline{a}$, since we probe the standard normal only. Then, ITS has a safety margin $\eta \leq 7$ while DS has $\eta' \leq 38$. These results are conditional to using the implementations of $f(x)$, $F(x)$ and $q(p)$ available in `Rmpfr`, but they do not change when falling back to the original implementations provided in `stats`.

These results can partly explain the safety margins for other distributions, as `R` basic libraries often utilize normal approximations. In particular, base `R` implementations of binomial and Poisson $q(p)$ involve a so-called Edgeworth approximation, which is essentially an enhancement modification to approximate normal quantiles for the case of skewed distributions. We address these approximately normal distributions in the following subsections.

5.2 Binomial

We consider now the binomial case. We ran simulations under some configurations of its parameters, namely, size n and probability p, with mean np and variance $np(1 - p)$. In Figure 1, the values of η and η' were obtained at few selected locations and then interpolated for added readability in maps.

A troubled cone-shaped region is visible in the middle of these plots, with $p = 1/2$ as the most difficult iso-n configuration, and issues spreading to more values of p as n diverges. It seems that, in the limit, `truncdist` has safety margin limited to $\eta \leq 10$, likely due to internal approximations of $q(p)$ to the normal.

The cone visible in plots is mostly a reflection of a condition on n and p that triggers such approximation, so it is partly a side-effect of a specific implementation. In this sense, we trust a sampler as far as it can simulate the approximate, normal, distribution, so there is room for more relevant tests than the ones we present here, which though require a different and more precise binomial implementation.

The proposed alternative `truncLCdist` package not only provides better support for small values of n, but it does achieve a higher safety margin as n diverges, where $\eta' \leq 40$. This result, though close to the one for normal, is instead related to the beta distribution, due to the way in which binomial $F(x)$ is approximated in `stats`. This aspects makes the binomial example not redundant after results related to the normal case.

5.3 Poisson

Poisson distribution is ruled only by a shape parameter $\lambda > 0$ representing both the mean and the variance. In Figure 2, the safety margin η is visibly under ten sigma with `truncdist`. This result too might depend on normal approximations used in `stats`. As to `truncLCdist`, η' is again close to the critical value $\eta' \approx 38$ of the normal case. This result is effectively redundant, since the Poisson distribution ships with an approximation to the normal in `R`.

5.4 Remarks

As a nuance, truncation cannot be arbitrarily extreme in simulations and the safety margin is typically undocumented. This issue does not arise in abstract probability calculus, but in practical finite-precision computing, which is why we had to simulate in order to make our point. Precision in such computations is limited, so $F(x)$ is imputed equal to either zero or one if $X = x$ is a many sigma event. Its inverse $q(p)$ cannot be evaluated meaningfully for all orders, relatedly. This problem must be more acute for the normal and the other light tailed distributions, as their tail probabilities decay faster, but it can be expected to affect even Cauchy and other heavy tailed distributions at some point.

These results serve as a reminder that even basic building blocks in statistics, such as normal quantiles and tail probabilities, can still be a topical issue. When implementing probability-related functions in statistical software, care must be taken for the consequences of low precision in tasks that push these mathematical objects to their computational limits. With truncated distributions, safety should not be taken for granted, the behavior of ITS should be checked, and alternative samplers should be considered.

Figure 1: Safety in the binomial (n, p) case. The x-axis is on a logarithmic scale, the y-axis is on a logit scale.
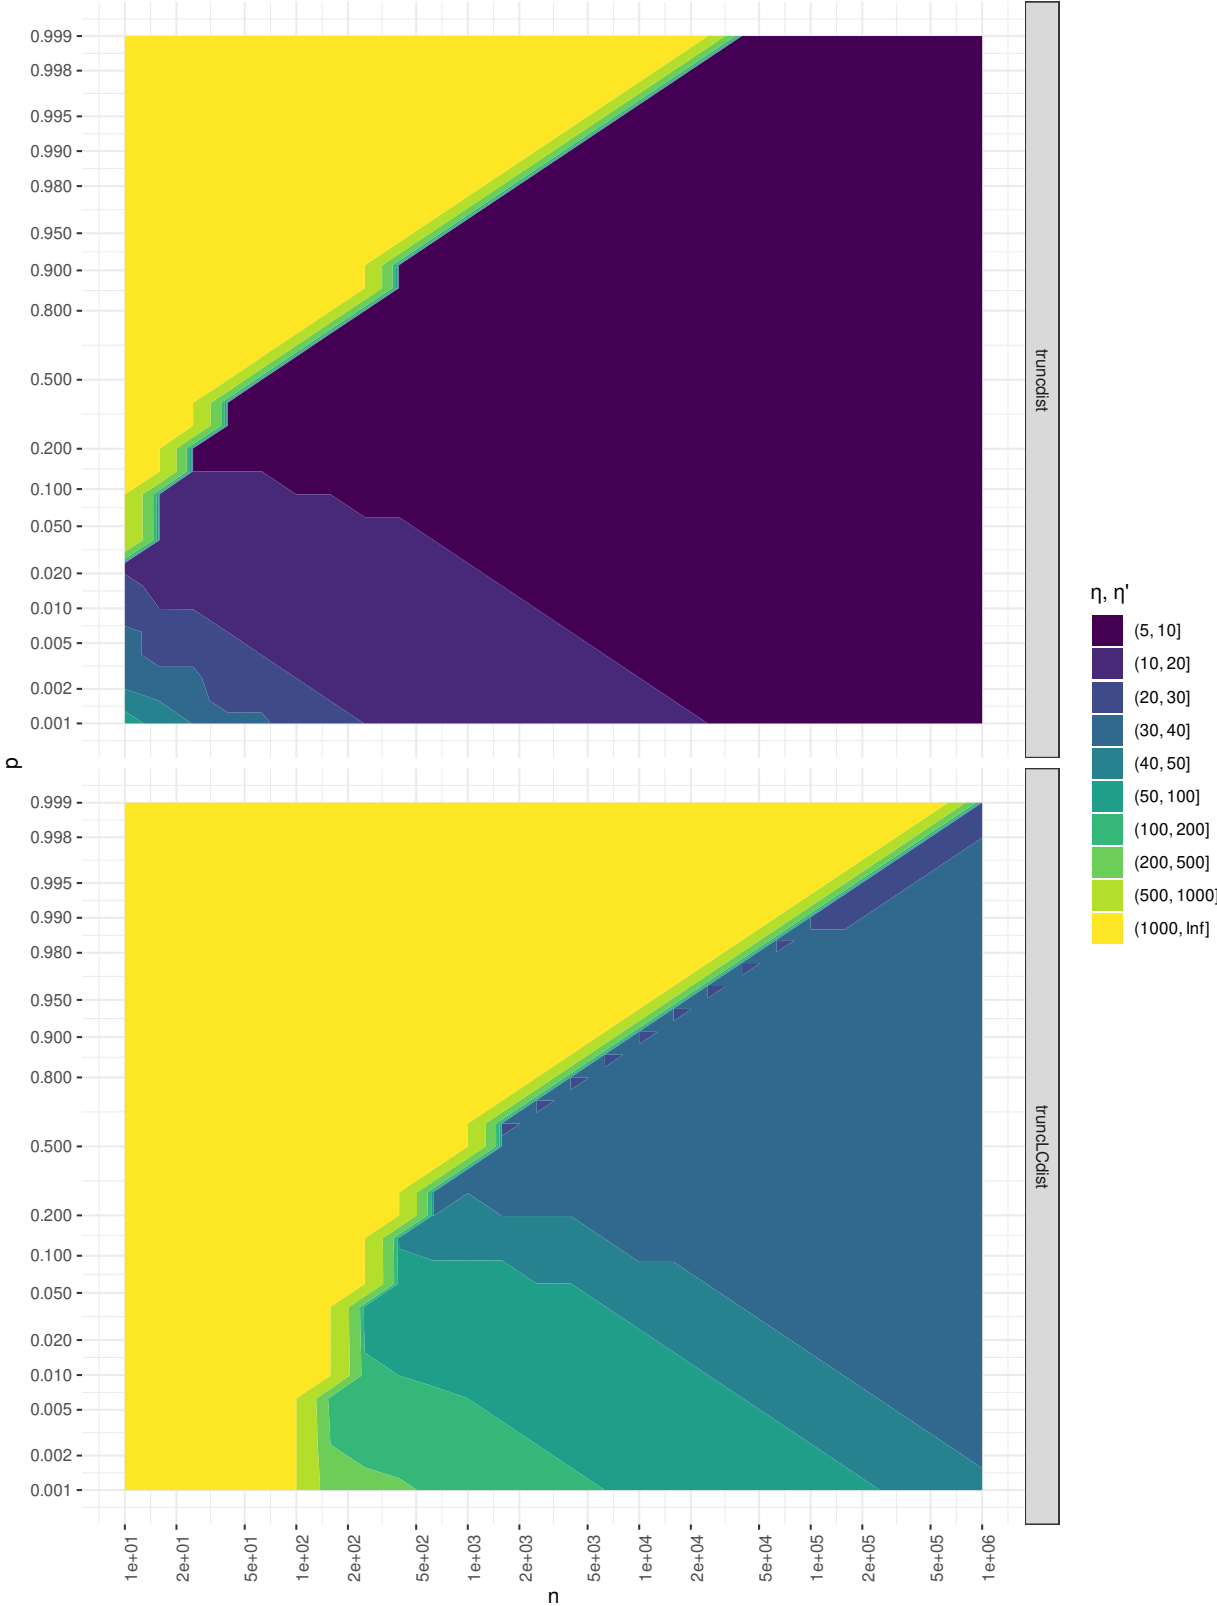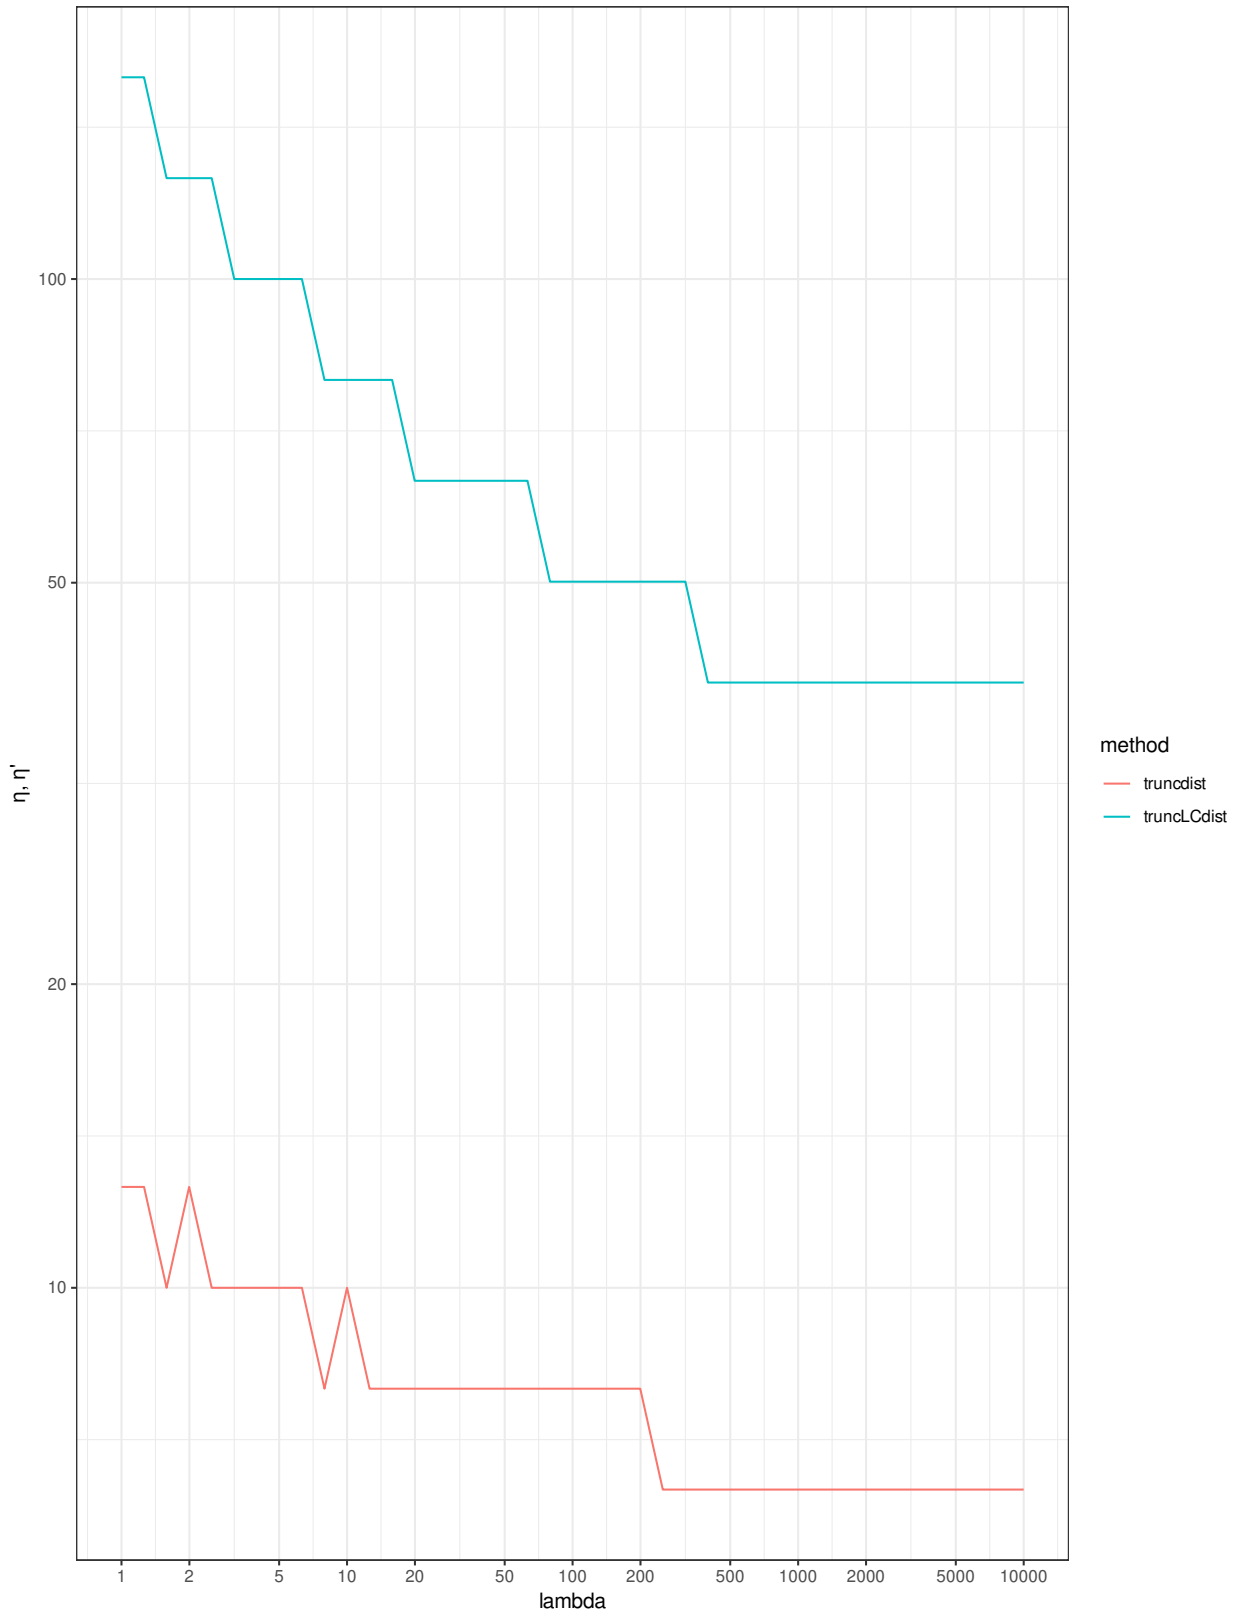
Figure 2: Safety in the Poisson(λ) case. Both x- and y- axis are on a logarithmic scale.

5.5 Geometric

The geometric distribution is ruled by the event probability p, the R default implementation has support over all non-negative integers, and it is endowed with a property known as *lack of memory*, namely,

$$X \sim \text{geom}(p) \implies P(X > x + h | X > x) = P(X > h), \quad x, h = 0, 1, \ldots.$$

Truncation under a geometric model just implies an affine translation of the random variable. Even this simple fact from probability calculus may not be reflected in computation due to software-related issues. The package `truncdist` is agnostic as to the probability distributions, giving no special attention to anyone in the specific, so does our `truncLCdist` package.

We probe some configurations of p and check some values of a taken from a geometric progression, until failure is detected. In Figure 3, the value of η is reported as a function of p. One may first notice the asymmetry in the plot, as among opposite log-odds the positive ones are easier to deal with. It is disturbing how little the absence of memory is reflected in ITS-based simulations. While such a property should hold independently of the extremeness of truncation, it is at least reassuring that DS can reflect it to a broader extent.

5.6 Exponential

The exponential distribution is ruled by the rate parameter $\lambda > 0$, whose reciprocal serves as the scale parameter σ, and the latter can be kept fixed in simulations with little loss of generality, in the same fashion as when keeping σ fixed in the normal case. The exponential distribution has no memory, so it is the continuous counterpart to the geometric distribution. In formulas,

$$X \sim \text{exponential}(\lambda) \implies P(X > x + h | X > x) = P(X > h), \quad x, h \geq 0.$$

We simulated the exponential distribution via the gamma with shape parameter equal to one. Checking all integer values of a from 1 to 1000, ITS seems safe up to $\eta \leq 31$, while DS is safe to a broader extent, as $\eta' \leq 745$.

5.7 Exponential power distribution

We shortly introduce the exponential power distribution (EPD), which is supported by our package only for ancillary purposes. Let the EPD be defined as

$$f(x; \beta) = \frac{1}{2\Gamma(1/\beta + 1)} \exp\left(-|x|^\beta\right), \quad x \in \mathbb{R}.$$

When $\beta \geq 1$, the distribution is log-concave, and this is the only case relevant to us, wherein it allows to generate gamma-distributed variates, indirectly.

5.8 Gamma

The gamma distribution is ruled by a shape parameter $\alpha > 0$ and the rate $\lambda > 0$, hence it has expected value α/λ and variance α/λ^2. In simulations, we probe a grid of values for α, while we keep $\lambda = 1$ fixed without loss of generality, as with the scale parameter in the normal case.

The gamma distribution is a special one among those we address in this paper and in the package, since it is a log-concave distribution only for some parameter configurations, namely, when $\alpha \geq 1$. The case $\alpha = 1$ implies an exponential density. As to the case $\alpha < 1$ [3], if X is EPD with shape parameter $\beta > 1$ then $|X|^\beta$ is gamma-distributed with $\alpha = 1/\beta < 1$. In the package we set up an exception handler that allows to override an existing definition depending on the parameter configuration.

Figure 4 reports on η and η' for different configurations of α. When $\alpha < 1$, DS seems not to improve much on traditional ITS. Differences in safety can be appreciated, instead, when $\alpha \geq 1$. This otherwise curious shift in regime around $\alpha = 1$ is only due to the switch in implementation from the indirect EPD-based simulation when $\alpha < 1$ to direct DS in the case when $\alpha \geq 1$.

As a matter of style, this casewise solution may not be elegant, but the gamma case looks actually tricky in the literature, where many other samplers have been proposed to address exclusively one of the two cases. As it concerns the generality of the exception handler, this workaround may look mostly like a patch meant only for the gamma case, but we foresee it may help in implementing other distributions that are log-concave only on subspaces of $\mathcal{X}$ (such as the t-distribution) or even transformations of log-concave-distributed variables (which may not be log-concave-distributed in turn).

Figure 3: Safety in the geometric(p) case. The x-axis is on a logit scale, the y-axis on a logarithmic scale.
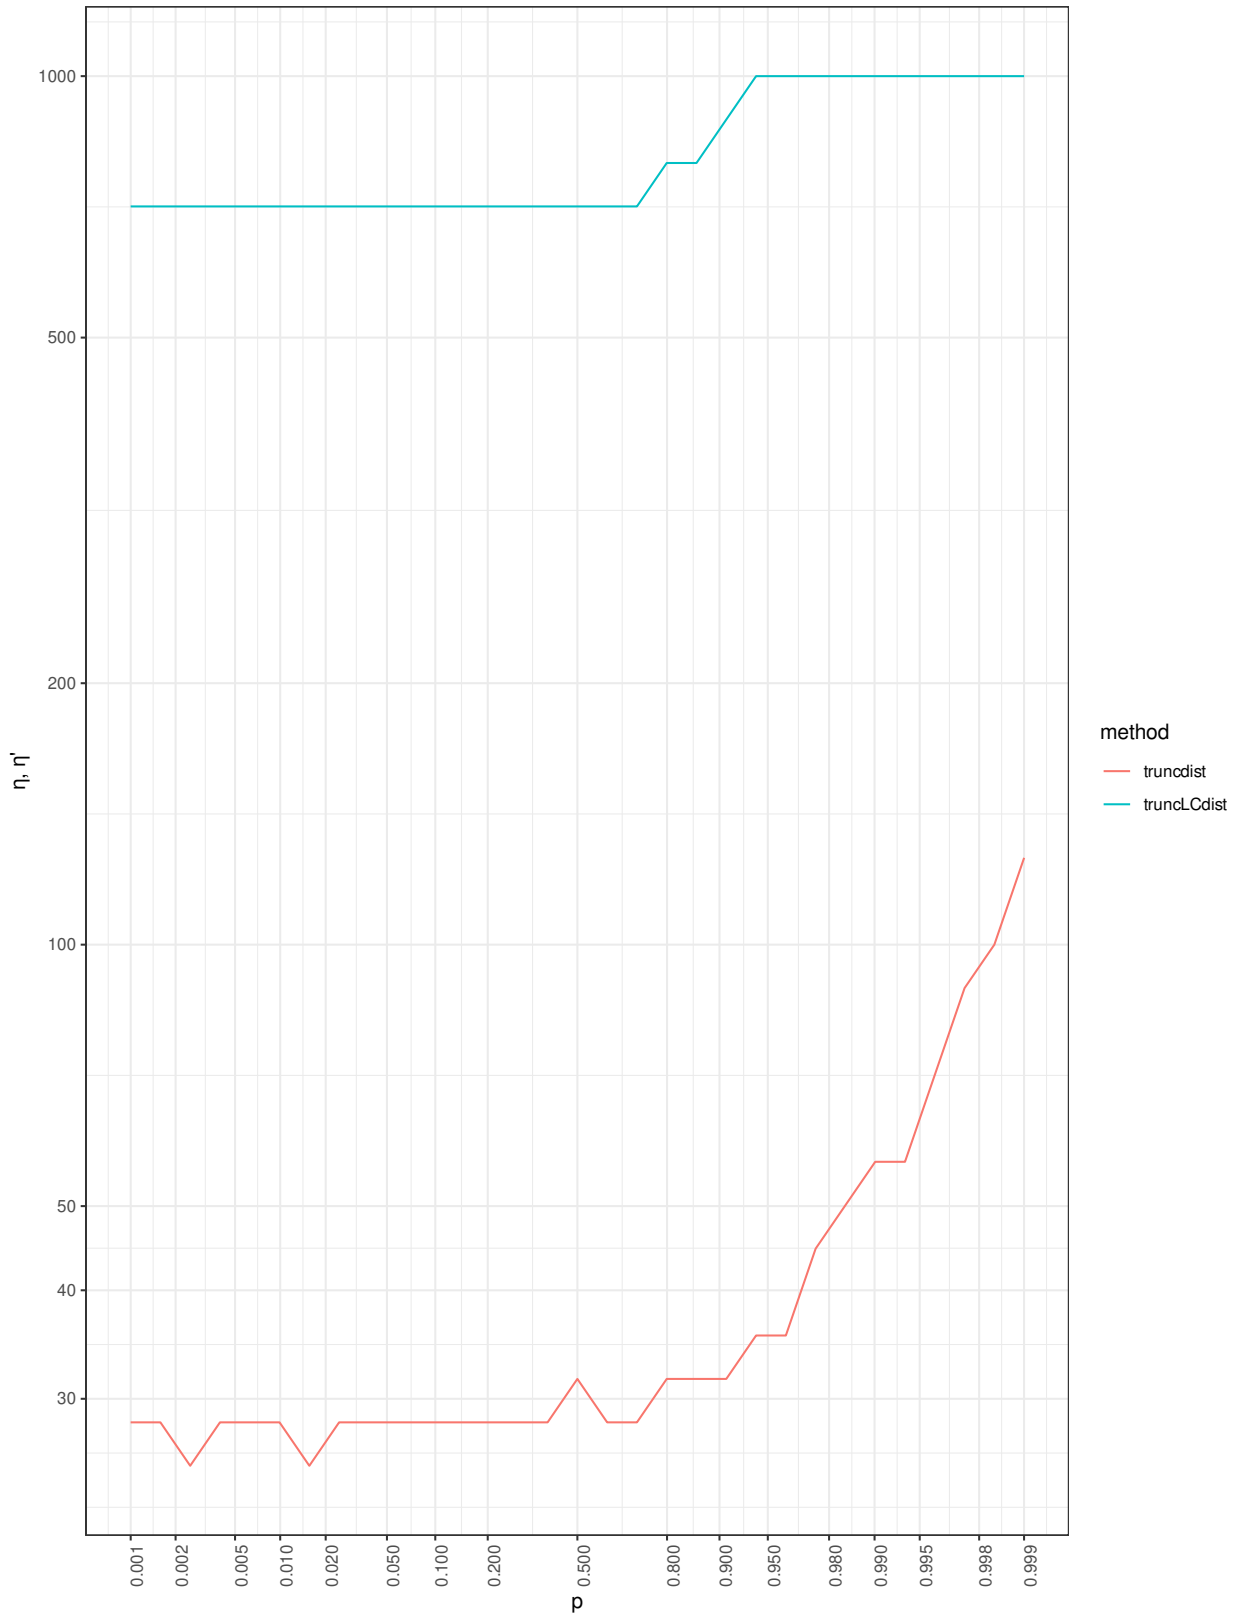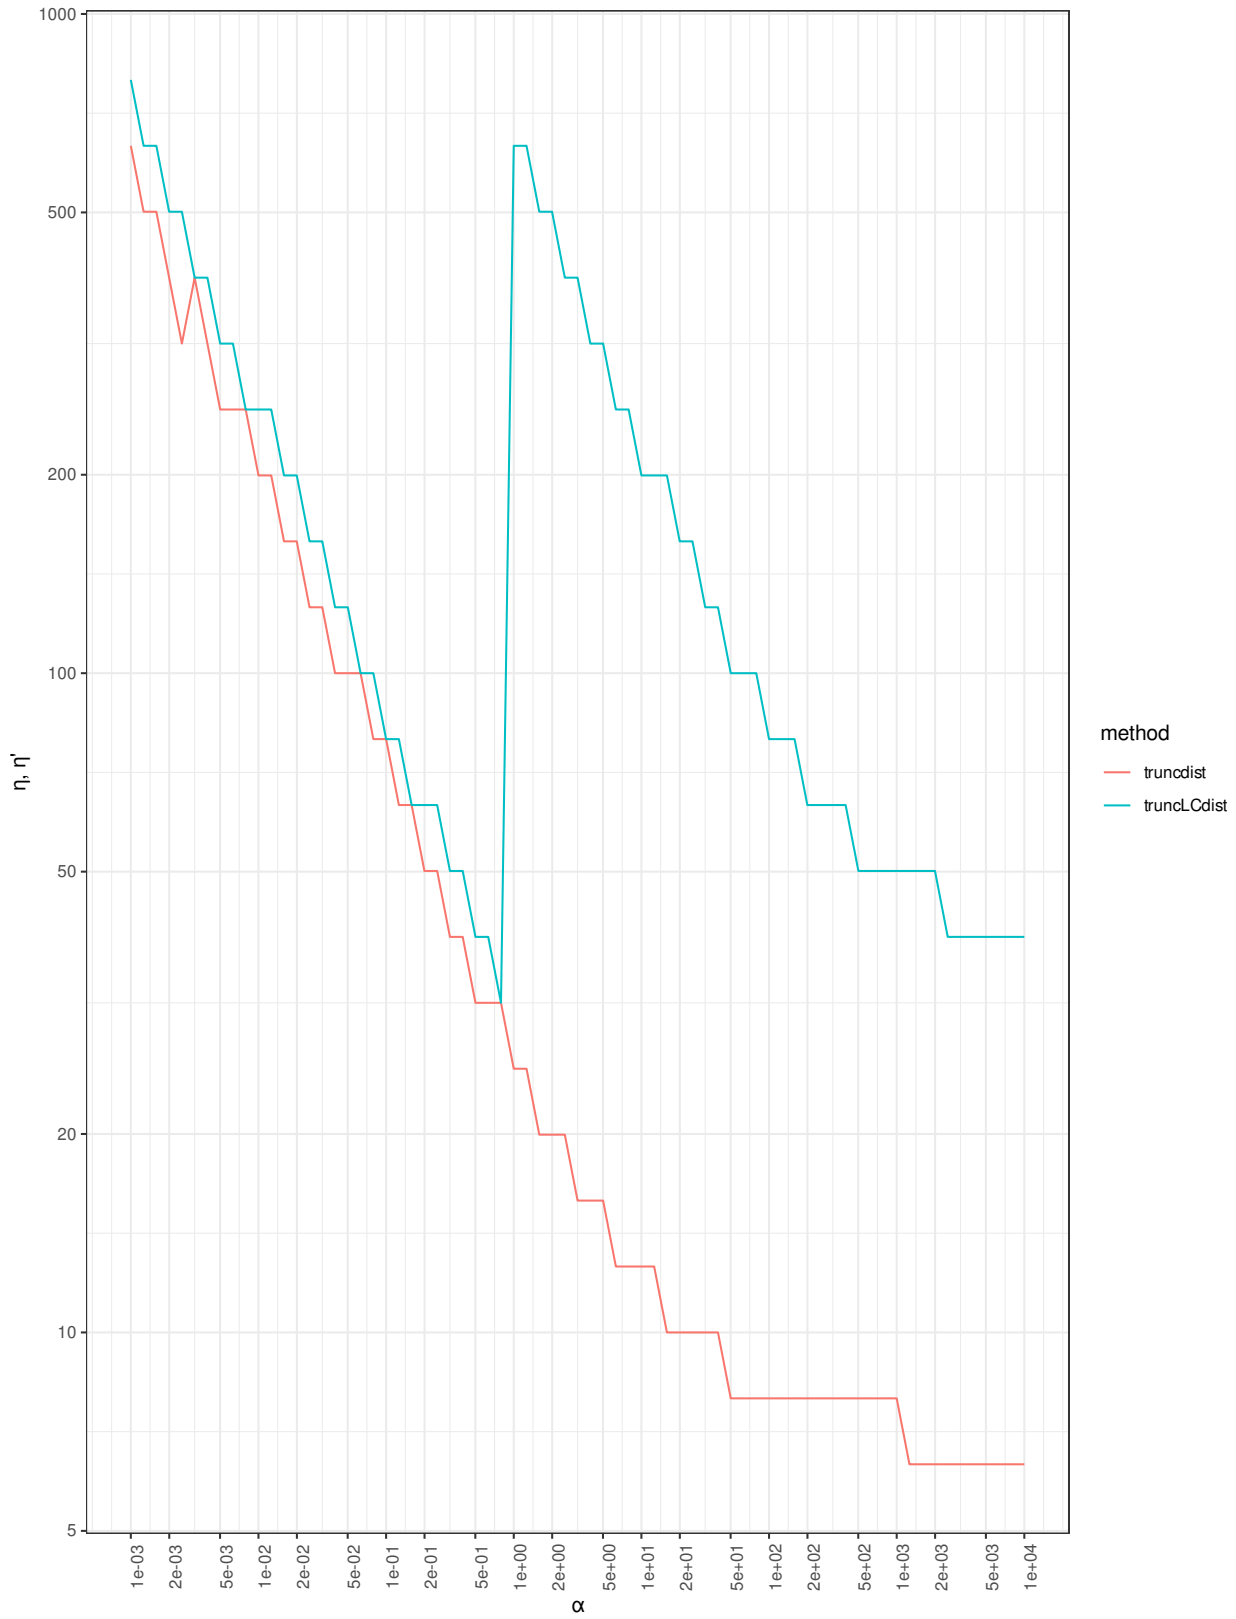
Figure 4: Safety in the gamma$(\alpha, \lambda = 1)$ case. Both x- and y- axis are on a logarithmic scale.

5.9 Negative binomial

The negative binomial distribution, also known as compound Poisson-gamma, can be used in modeling the length of a sequence of Bernoulli random variables, with event probability p, ending upon the n-th event. Like in the binomial case, n and p are shape parameters, so we probe a two-dimensional grid of configurations.

The values of η are reported in Figure 5. DS looks uniformly safer than ITS, though the safety of both samplers visibly erodes as n diverges. Here again $q(p)$ fails sooner than $F(x)$, hence the relation between the contour plots.

5.10 Inverse Gaussian

The inverse Gaussian distribution is ruled by two orthogonal shape parameters, $\mu > 0$ and $\phi > 0$, respectively the mean and dispersion parameters. The safety of DS looks uniformly superior to the one of ITS.

6 Beyond diagnostics

In the previous section, we have investigated the safety margins of two samplers based on some proxy information, namely, diagnostics. Our endpoint with `truncdist` was the presence or absence of execution errors or infinite values. The endpoint for `truncLCdist` was the imputation rate. However, even when diagnostics raise no warning, it is still possible that samples do not reflect the sampled distribution. We now validate the diagnostics as a proxy for actual reliability of our proposed sampler. Since an exhaustive validation is not possible, we only show few examples, namely, normal, Poisson and geometric.

6.1 Normal

The truncated normal distribution is quite popular in statistical modeling, as outlined in the introduction. The normal distribution is also one of the few whose truncation behavior is well known. The Mills ratio recurs in the expression of cumulants [55], and the excess distribution resembles the exponential under extreme truncation [49]. The behavior of other distributions is generally less known, but this knowledge can be surrogated with Monte Carlo simulations. As a remark, this is why the reliability of samplers is so crucial.

We consider the standard normal case, without loss of generality, because the mean and standard deviation serve as location and scale parameters, respectively. We let a range in the lattice $0, 0.05, 0.1, \ldots, 50$. Both the samplers we consider have safety $\overline{a}$ lying between 0 and 50, actually. For each value of a, we simulate $n = 10^5$ independent samples with both samplers. Ideally, it would be best to test the distribution itself in a Kolmogorov-Smirnov fashion, but the problem is that the descriptors in use for random number generation affect both simulation and the test statistic. We use the sample mean

$$\overline{X} = \frac{1}{n} \sum_{i=1}^{n} X_i \,,$$

and standard deviation

$$S = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (X_i - \overline{X})^2} \,,$$

based on n independent variates $X_1, \ldots, X_n \sim f_{\mathcal{I}}$, to compute a Z-statistic defined as

$$Z = \frac{\overline{X} - \mu_a}{S/\sqrt{n}} \overset{.}{\sim} \mathrm{normal}(0, 1) \,.$$

Here, $\mu_a = \mathbb{E} X_i > a$ is the expected value of the truncated normal, dependent on a. If the sampler is reliable, then Z has approximately a standard normal distribution. The aim is to test the mean of the samples, while it would not be meaningful to test the mode for instance, because this is typically the imputed value with our package.

Our test with the Z-statistic is meaningful only as far as μ_a can be computed reliably. The R package `truncnorm` [40] provides a function `etruncnorm` that computes μ_a reliably even under extreme truncation. However, even this function has a rupture point, as for instance it can yield $\mu_a < a$ when $a > 12000$, though this does not look like an impairing defect. The evaluation of truncated variance via the function `vtruncnorm` is instead much less reliable, as it fails for $a > 9$. This is why we estimate the truncated standard deviation empirically, via S.

In Figure 7, the Z-statistic is reported along the vertical axis, and a along the horizontal one. Our proposed sampler works for $a \leq \overline{a} = 38.45$ while `truncdist` is still less safe as predicted. However, the point in this test is to show that our package can actually be trusted when diagnostics report no error.

Figure 5: Safety in the negative binomial(n, p) case. The x-axis is on a logarithmic scale, and the y-axis is on a logit scale, like in the binomial case.
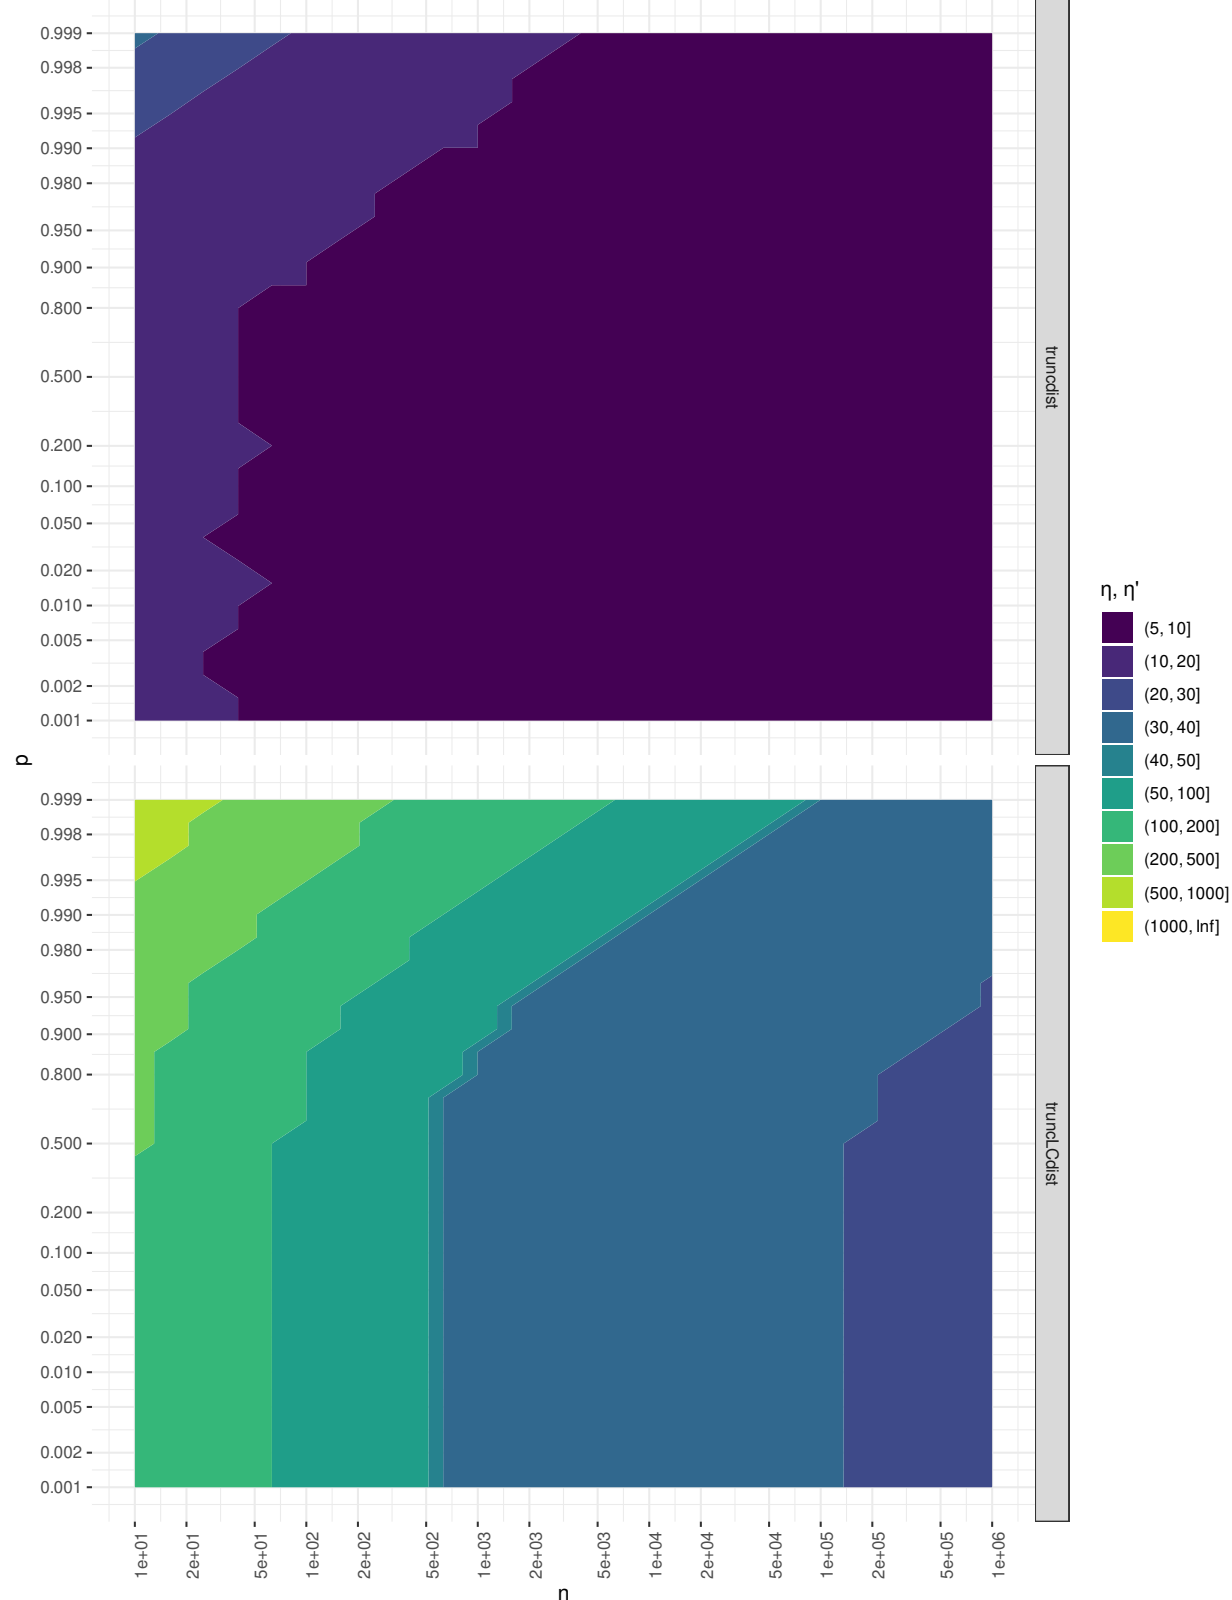
Figure 6: Safety in the inverse Gaussian(μ, ϕ) case. Both x- and y- axis are on a logarithmic scale.
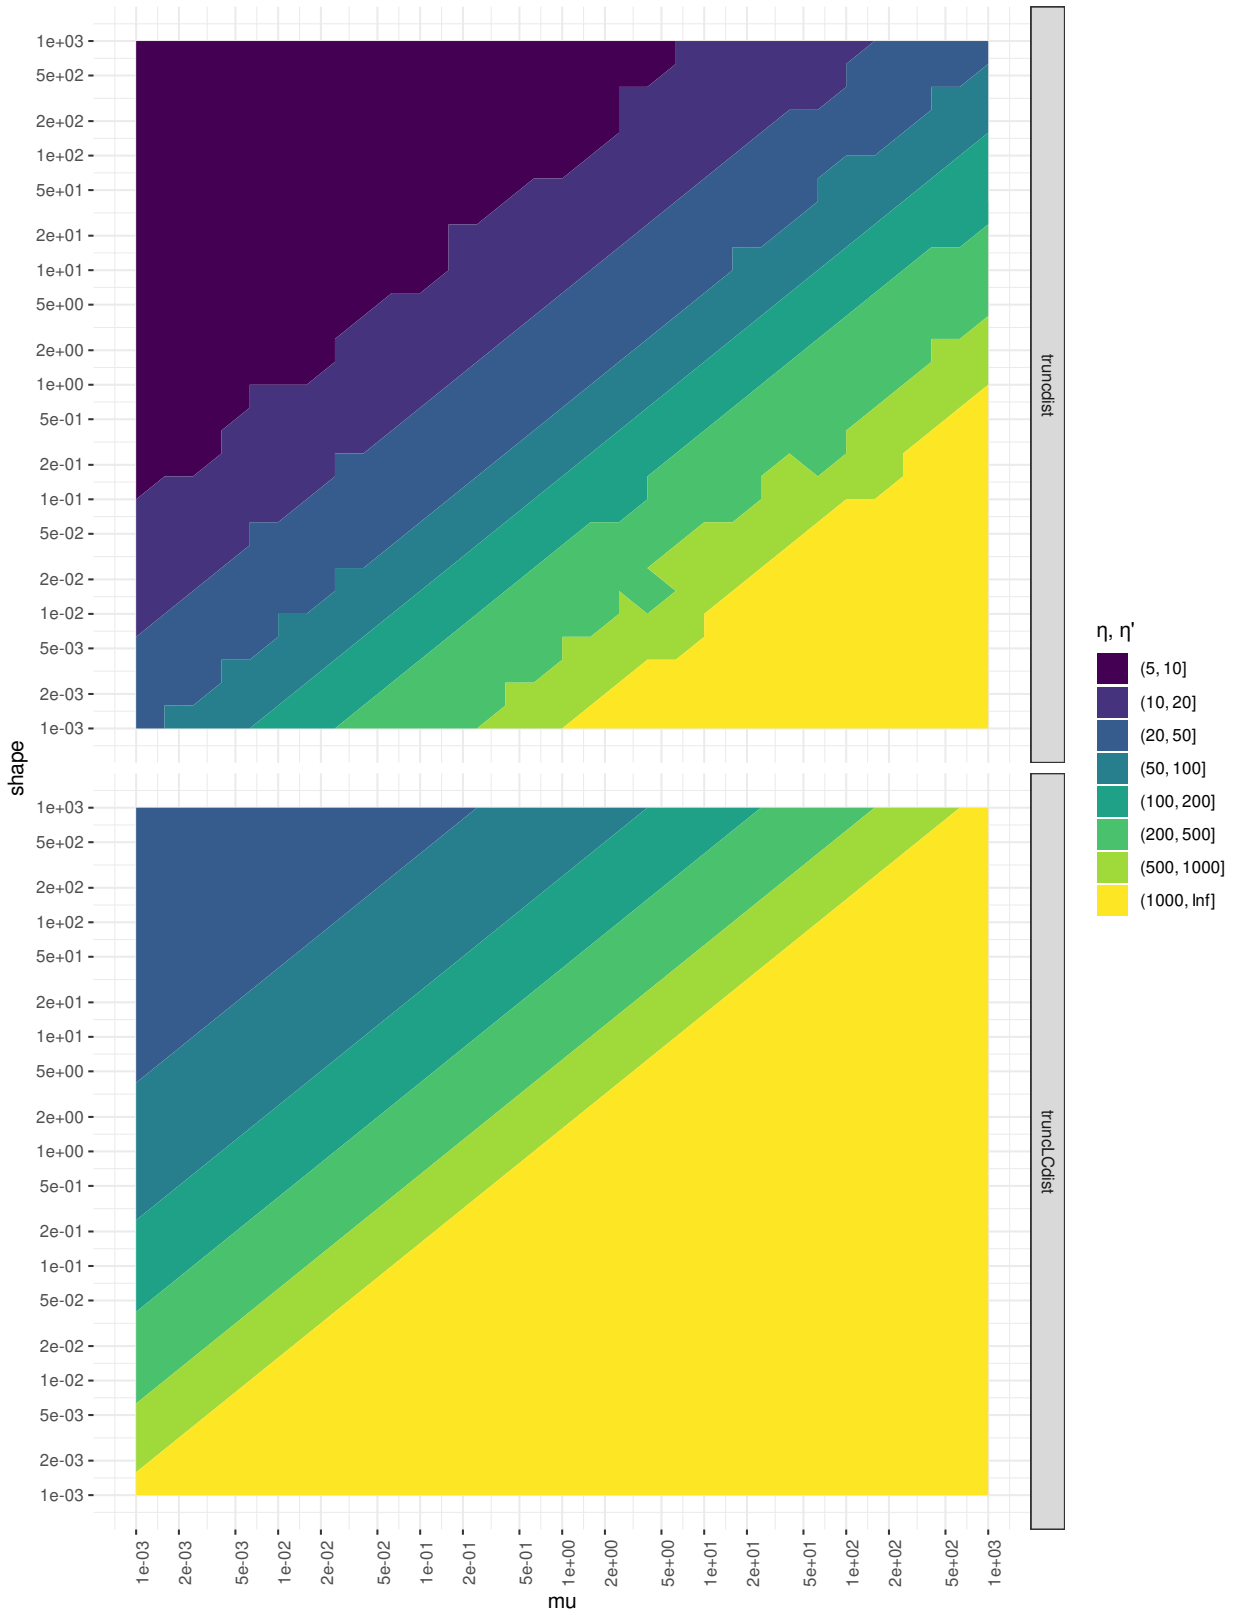
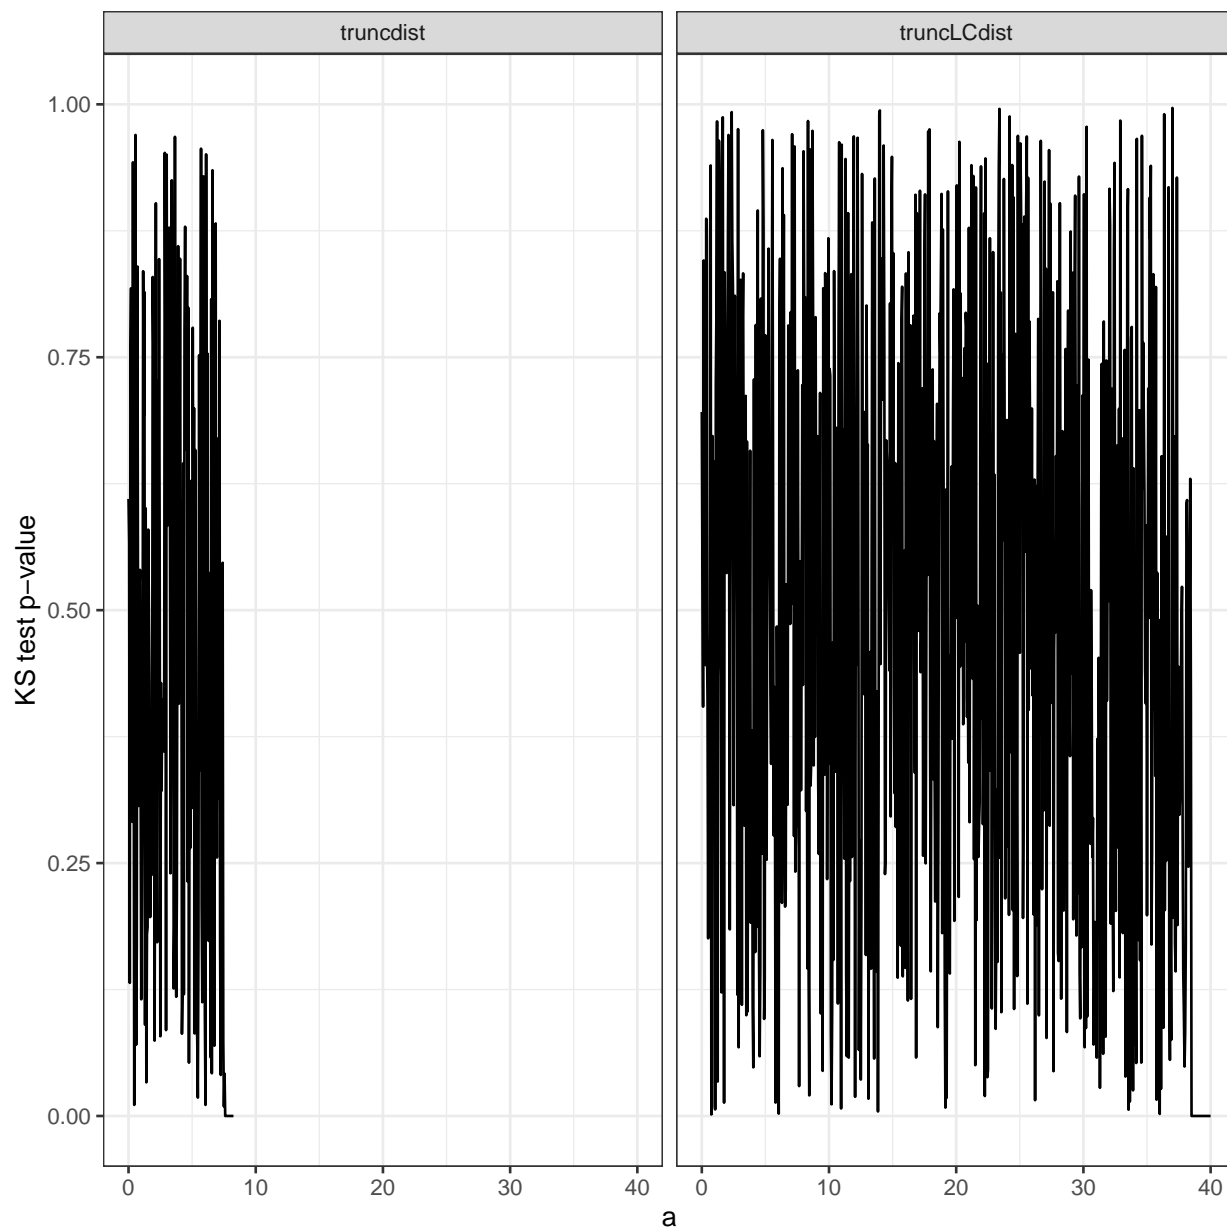Figure 7: Z statistic for the truncated normal mean, $n = 100000$ samples.

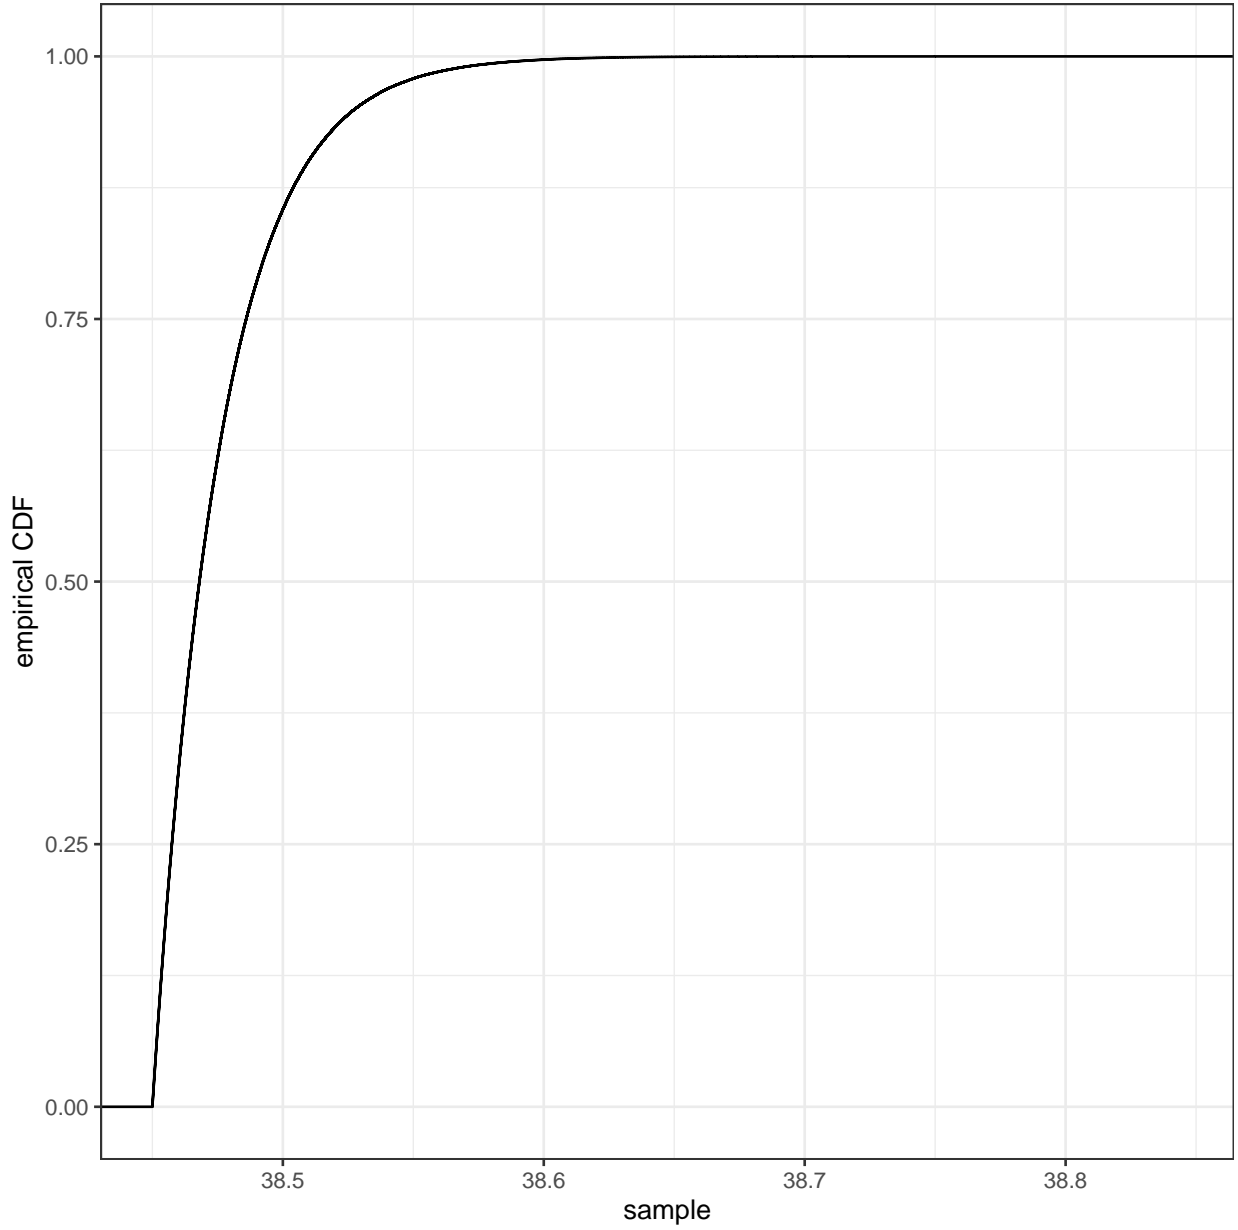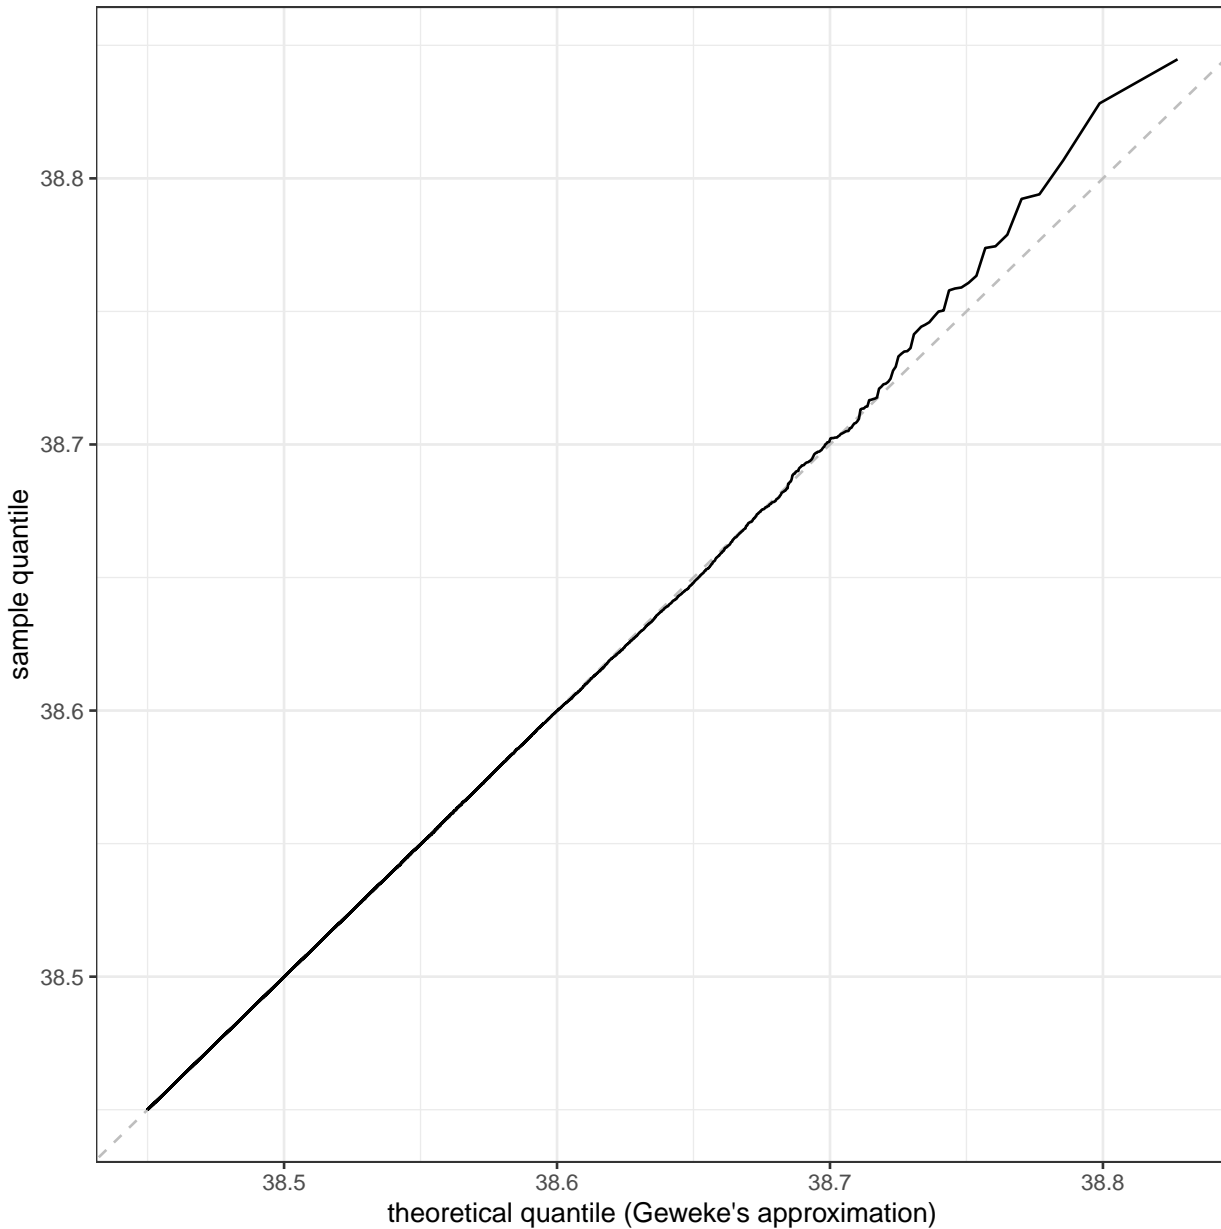Figure 8: Truncated normal, $a = 38.45$, 10^6 samples via `truncLCdist`.

Figure 9: Truncated normal, $a = 38.45$, Q-Q plot for an exponential(a) null hypothesis with 10^6 samples via `truncLCdist`.

We now focus on the behavior of DS as $a = \eta'$. When truncating the standard normal over $\mathcal{I} =]\eta', +\infty]$, the empirical cumulative distribution function with $n = 1000000$ samples can be as in Figure 8. This distribution may not be assessed out of simulations as even the R package `truncnorm` does not represent it accurately in terms of quantiles.

It can be shown easily that, provided a is "large", for $X \sim \mathrm{normal}(0, 1)$ it holds [49]

$$X - a|X > a \approx \mathrm{exponential}(a).$$

This exponential approximation seems to hold in Figure 9, where we report the Q-Q plot for sample quantiles versus the approximate theoretical ones.

6.2 Poisson

As a discrete case, we consider the Poisson distribution. In the normal case it was possible to set the scale and location parameters arbitrarily without loss off generality, while Poisson distribution is non-trivially affected by a shape parameter $\lambda > 0$. We then probe a grid of values (λ, z), and we set $a = \lambda + z\sqrt{\lambda}$. For each configuration, we compute the mean and standard deviation on 100000 draws from $X \sim \text{Poisson}(\lambda)$ truncated over $]a, +\infty]$. We use these data to test the samples for Poisson truncated mean, which is defined as follows, see for instance [55].

$$\lambda(a) = \text{E}(X - a | X > a) = \lambda \cdot \frac{1 - F_X(a - 1)}{1 - F_X(a)} .$$

The sample mean and standard deviation are used to compute a Z statistic in testing $\lambda(a)$, which is reported versus λ and z, in Figure 10. A contour plot is used to show the Z value. This statistic seems well behaved, as it appears to be not affected by the configuration in a systematic way. Moreover, 16 values of λ and 100 values of z were probed, making it likely to satisfy $|\text{Z}| < 3.5$ when the samplers are not troubled.

This last example based on the Poisson distribution should be telling as to why the validation of diagnostic tools in this paper is so limited. Truncation makes statistical higher dimensional and thus more difficult to analyze. Even with a one dimensional base parameter, the truncated distribution would required a rather complicated output to be accurately assessed.

7 Discussion

We have shown that for log-concave distributions it is safer to use Devroye's sampling method rather than the inverse transform. This approach is expected to improve the reliability of statistical analyses that involve both simulation and generalized linear models.

ITS has often been taken for granted to be a safe simulation technique. It is, at least, a very compute-efficient strategy, when $F(x)$ and $q(p)$ tractable. It also offers, with its unmatched generality, an automatic solution that is fast to deploy in practice. DS, instead, looks to be used mostly sparsely, potentially as a patch in cases where no other RS is easy to implement. DS actually represents the automatic counterpart of ITS among log-concave densities, which is why we ship an R package entirely based on it. In RS, a 20-25% acceptance rate can be improved upon, but not in terms of orders of magnitude, so it is worth considering even if suboptimal, at least from a programming standpoint, since it can address a large set of models in few lines of code.

Our proposal exploits an under-considered aspect of log-concavity under truncation, that is, its closure, which makes DS a scalable sampler even under extreme truncation, wherein other envelopes in RS may work inefficiently. Not all famous log-concave distributions are implemented in the package, but the user may easily add any on their own, though we implemented all famous Tweedie distributions for regression purposes.

Our illustrative examples show that simulation of truncated distributions is currently, at best, a perilous task, since it pushes computation to its limits. We implemented just few distributions, while planning to add more of them in the future, defaults having been tested where present. All the other untested log-concave distributions are left to the user's care who, in the light of above results, uses them at a risk.

There are some open problems, like implementing a sampler for truncated distributions that can work even with un-normalized distributions. This would be much of an improvement, as DS requires $f_{\mathcal{I}}(x)$ to integrate or summate to unit, so it needs to evaluate the normalization constant involving $F(a)$ and $F(b)$. Other samplers do not require such a knowledge, for instance, [56] propose a piecewise-exponential envelope for log-concave distributions that is tailored down to any specific $f(x)$ in an adaptive fashion. This algorithm exploits the fact that a log-concave distribution can not only be upper-bounded, but also lower-bounded, by means of two piecewise-exponential functions, and both help controlling the acceptance rate.

The adaptive sampler can be seen as a more robust alternative to DS in the sense we outlined, as it requires evaluating just $f(x)$ and not $F(x)$, though it is more elaborate. Indeed, DS involves an envelope that is ruled by a location parameter, m and a scale parameter $\sigma = 1/f(m)$, which allow for a simple generation of proposed variates and subsequent rejection test. The adaptive sampler relies instead on entirely different envelopes from one target $f(x)$ to another. This latter sampling strategy deserves consideration, but it will be explored further in the future of our `truncLCdist` package. For now, a rather automatic sampler is provided, whose execution time is bounded in probability, wherein the truncation is tractable, that is, within the safety margin summarized by η in the examples.

Possible extensions in the future should definitely target all non-log-concave distributions [57], which we have not covered in our application, due to the lack of a solution of comparable generality. In the case of failure of DS, likely

Figure 10: Z statistic for the truncated Poisson mean, $n = 100000$ samples for each configuration of (λ, z). A conjectured asymptote at $z = 38$ is also shown.
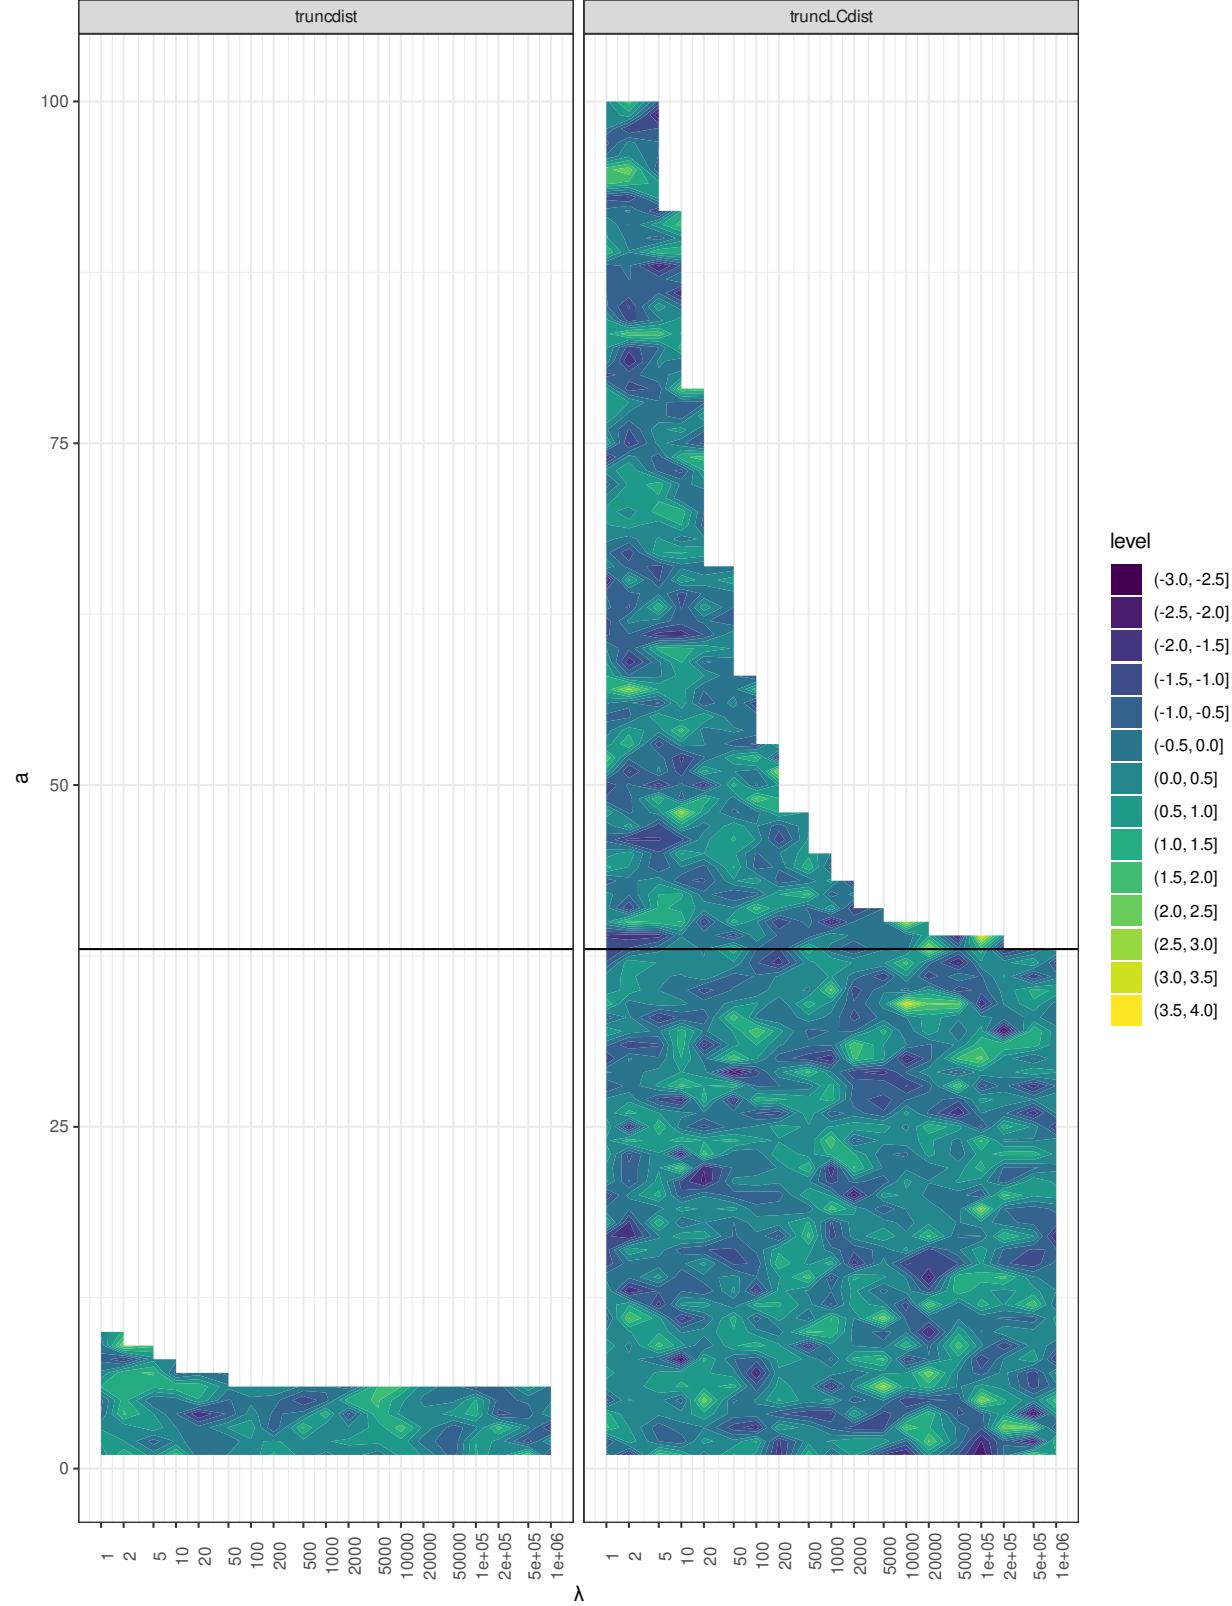
due to numerical instability of $F(x)$, it might be possible to resort to some adaptive version of rejection sampling [58, 56]. In this case, more acumen in required in coding than it was the focus of this paper. DS needs little tuning, so the comparison with ITS is less viced by the lack of wit in coding. As far as the issues highlighted in this paper look like subtleties, it should be understood why we are careful not to overestimate our own coding skills. Exponential approximations of the normalization constant are also possible in line with Geweke [49] at least for the normal, but also for other log-concave distributions based on heuristic grounds.

DS can be made more safe to a wider extent than currently offered by `truncLCdist`, just by providing better implementations of $f(x)$, $F(x)$. It is indeed the evaluation of $f(x)$ and of $F(x)$ that troubles the evaluation of $f_{\mathcal{I}}(x)$. In the meantime, our package serves as a wrapper that combines together $f(x)$ and $F(x)$ automatically, in order to provide reliable truncated samples.

The R package `truncLCdist`, used in the simulations for this paper, is publicly available at `https://github.com/mlambardi/truncLCdist`.

References

[1] Purushottam W. Laud, Paul Damien, and Thomas S. Shively. Sampling some truncated distributions via rejection algorithms. *Communications in Statistics - Simulation and Computation*, 39(6):1111–1121, June 2010.

[2] Zdravko Botev and Pierre L'Ecuyer. Simulation from the normal distribution truncated to an interval in the tail. In *Proceedings of the 10th EAI International Conference on Performance Evaluation Methodologies and Tools*. ACM, 2017.

[3] Luc Devroye. *Non-Uniform Random Variate Generation*. Springer New York, 1986.

[4] Luc Devroye. A simple generator for discrete log-concave distributions. *Computing*, 39(1):87–91, March 1987.

[5] Michael Evans and Timothy Swartz. *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, 2000.

[6] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2021.

[7] Larry R. Russell. Probability distributions for hurricane effects. *Journal of the Waterways, Harbors and Coastal Engineering Division*, 97(1):139–154, February 1971.

[8] Bakhodir Ergashev, Konstantin Pavlikov, Stan Uryasev, and Evangelos Sekeris. Estimation of truncated data samples in operational risk modeling. *Journal of Risk and Insurance*, 83(3):613–640, February 2015.

[9] John P. Klein and Prem K. Goel. *Survival Analysis: State of the Art*. Springer Netherlands, 1992.

[10] Igor V. Cadez, Padhraic Smyth, Geoff J. McLachlan, and Christine E. McLaren. Maximum likelihood estimation of mixture densities for binned and truncated multivariate data. *Machine Learning*, 47(1):7–34, 2002.

[11] R. M. Nixon and S. G. Thompson. Parametric modelling of cost data in medical studies. *Statistics in Medicine*, 23(8):1311–1331, 2004.

[12] H. Smith and D. W. Grace. Applications of truncated distributions in process start-ups and inventory control. *Technometrics*, 3(3):429–432, August 1961.

[13] Rahul A. Parsa, Jay J. Kim, and Myron Katzoff. Application of the truncated distributions and copulas in masking data. In *Joint Statistical Meetings*, pages 2770–2780, 2009.

[14] M. S. Tokmachev. Modeling of truncated probability distributions. *IOP Conference Series: Materials Science and Engineering*, 441:012056, nov 2018.

[15] Dennis Aigner, C. A. Knox Lovell, and Peter Schmidt. Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics*, 6(1):21–37, July 1977.

[16] Wim Meeusen and Julien van Den Broeck. Efficiency estimation from Cobb-Douglas production functions with composed error. *International Economic Review*, 18(2):435–444, June 1977.

[17] Mark M. Pitt and Lung-Fei Lee. The measurement and sources of technical inefficiency in the Indonesian weaving industry. *Journal of Development Economics*, 9(1):43–64, August 1981.

[18] Phill Wheat, Alexander D. Stead, and William H. Greene. Robust stochastic frontier analysis: A Student's t-half normal model with application to highway maintenance costs in England. *Journal of Productivity Analysis*, 51(1):21–38, January 2019.

[19] H. R. Thompson. Truncated normal distributions. *Nature*, 165(4194):444–445, March 1950.

[20] David B. Thomas, Wayne Luk, Philip H. W. Leong, and John D. Villasenor. Gaussian random number generators. *ACM Computing Surveys*, 39(4):11:1–11:38, November 2007.

[21] Justin R. Williams, Hyung-Woo Kim, and Catherine M. Crespi. Modeling observations with a detection limit using a truncated normal distribution with censoring. *BMC Medical Research Methodology*, 20(1):1–25, June 2020.

[22] Siqi Chen and Wenhao Gui. Estimation of unknown parameters of truncated normal distribution under adaptive progressive type II censoring scheme. *Mathematics*, 9(1):49, December 2021.

[23] Seth J. Wenger and Mary C. Freeman. Estimating species occurrence, abundance, and detection probability using zero-inflated distributions. *Ecology*, 89(10):2953–2959, October 2008.

[24] Hans R. Künsch, Leonard A. Stefanski, and Raymond J. Carroll. Conditionally unbiased bounded-influence estimation in general regression models, with applications to generalized linear models. *Journal of the American Statistical Association*, 84(406):460–466, June 1989.

[25] Stephen C. Nelson and Peter J. Katzenstein. Uncertainty, risk, and the financial crisis of 2008. *International Organization*, 68(2):361–392, 2014.

[26] Reuven Y. Rubinstein and Dirk P. Kroese. *Simulation and the Monte Carlo Method*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., November 2016.

[27] Mark J. van der Laan and Jenny Bryan. Gene expression analysis with the parametric bootstrap. *Biostatistics*, 2(4):445–461, December 2001.

[28] Yili Hong, William Q. Meeker, and James D. McCalley. Prediction of remaining life of power transformers based on left truncated and right censored lifetime data. *The Annals of Applied Statistics*, 3(2), jun 2009.

[29] Daniel C. Monchuk, Zhuo Chen, and Yosef Bonaparte. Explaining production inefficiency in China's agriculture using data envelopment analysis and semi-parametric bootstrapping. *China Economic Review*, 21(2):346–354, June 2010.

[30] Frederick Novomestky and Saralees Nadarajah. *truncdist: Truncated Random Variables*, 2016. R package version 1.0-2.

[31] Christophe Dutang and Patrice Kiener. *CRAN Task View: Probability Distributions*, May 2021.

[32] Martin Maechler. *Rmpfr: R MPFR - Multiple Precision Floating-Point Reliable*, 2021. R package version 0.8-4.

[33] Micah Altman, Jeff Gill, and Michael P. McDonald. *Numerical Issues in Statistical Computing for the Social Scientist*. John Wiley & Sons, Inc., December 2003.

[34] Catherine Loader. Fast and accurate computation of binomial probabilities. Copy of a paper no longer available, accessed April 2, 2022, 2002.

[35] Adrien Saumard and Jon A. Wellner. Log-concavity and strong log-concavity: A review. *Statistics Surveys*, 8:45–114, January 2014.

[36] Anne Philippe. Simulation of right and left truncated gamma distributions by mixtures. *Statistics and Computing*, 7(3):173–181, 1997.

[37] Martin S. Ridout. Generating random numbers from a distribution specified by its Laplace transform. *Statistics and Computing*, 19(4):439–450, October 2008.

[38] B. D. McCullough. Assessing the reliability of statistical software: Part I. *The American Statistician*, 52(4):358–366, November 1998.

[39] Guy Mélard. On the accuracy of statistical procedures in Microsoft Excel 2010. *Computational Statistics*, 29(5):1095–1128, April 2014.

[40] Olaf Mersmann, Heike Trautmann, Detlef Steuer, and Björn Bornkamp. *truncnorm: Truncated Normal Distribution*, 2018. R package version 1.0-8.

[41] Stefan Wilhelm and B. G. Manjunath. tmvtnorm: A package for the truncated multivariate normal distribution. *The R Journal*, 2(1):25–29, 2010.

[42] Alan E. Gelfand, Adrian F. M. Smith, and Tai-Ming Lee. Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling. *Journal of the American Statistical Association*, 87(418):523–532, June 1992.

[43] Luc Devroye. A note on generating random variables with log-concave densities. *Statistics & Probability Letters*, 82(5):1035–1039, May 2012.

[44] Luc Devroye. Chapter 4: Nonuniform Random Variate Generation. In *Simulation*, pages 83–121. Elsevier, 2006.

[45] Luc Devroye. Random variate generation in one line of code. In *Winter Simulation Conference*, pages 265–272, Los Alamitos, CA, USA, December 1996. IEEE Computer Society.

[46] Hassan Maatouk and Xavier Bay. A new rejection sampling method for truncated multivariate Gaussian random variables restricted to convex sets. In *Springer Proceedings in Mathematics & Statistics*, pages 521–530. Springer International Publishing, 2016.

[47] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, August 2005.

[48] Kenneth Burdett. Truncated means and variances. *Economics Letters*, 52(3):263–267, September 1996.

[49] John Geweke. Exact inference in the inequality constrained normal linear regression model. *Journal of Applied Econometrics*, 1(2):127–141, April 1986.

[50] Christophe Dutang, Vincent Goulet, and Mathieu Pigeon. actuar: An R package for actuarial science. *Journal of Statistical Software*, 25(7):1–37, 2008.

[51] Hakan Demirtas, Rawan Allozi, and Ran Gao. *UnivRNG: Univariate Pseudo-Random Number Generation*, 2021. R package version 1.2.3.

[52] Tymoteusz Wolodzko. *extraDistr: Additional Univariate and Multivariate Distributions*, 2020. R package version 1.9.1.

[53] Felipe Osorio. *heavy: Robust Estimation Using Heavy-Tailed Distributions*, 2019. R package version 0.38.196.

[54] Brian Lee Yung Rowe. *futile.options: Futile Options Management*, 2018. R package version 1.0.1.

[55] Zinoviy Landsman and Emiliano A. Valdez. Tail conditional expectations for exponential dispersion models. *ASTIN Bulletin*, 35(1):189–209, May 2005.

[56] P. Wild and W. R. Gilks. Algorithm AS 287: Adaptive rejection sampling from log-concave density functions. *Applied Statistics*, 42(4):701–709, 1993.

[57] Luca Martino and Joaquín Míguez. A generalization of the adaptive rejection sampling algorithm. *Statistics and Computing*, 21(4):633–647, August 2010.

[58] W. R. Gilks and P. Wild. Adaptive rejection sampling for Gibbs sampling. *Journal of the Royal Statistical Society (Series C, Applied Statistics)*, 41(2):337–348, 1992.