

Approximating inverse cumulative distribution functions to produce approximate random variables

MICHAEL GILES and OLIVER SHERIDAN-METHVEN, Mathematical Institute, Oxford University, UK

For random variables produced through the inverse transform method, approximate random variables are introduced, which are produced by approximations to a distribution's inverse cumulative distribution function. These approximations are designed to be computationally inexpensive, and much cheaper than library functions which are exact to within machine precision, and thus highly suitable for use in Monte Carlo simulations. The approximation errors they introduce can then be eliminated through use of the multilevel Monte Carlo method. Two approximations are presented for the Gaussian distribution: a piecewise constant on equally spaced intervals, and a piecewise linear using geometrically decaying intervals. The errors of the approximations are bounded and the convergence demonstrated, and the computational savings measured for C and C++ implementations. Implementations tailored for Intel and Arm hardware are inspected, alongside hardware agnostic implementations built using OpenMP. The savings are incorporated into a nested multilevel Monte Carlo framework with the Euler-Maruyama scheme to exploit the speed ups without losing accuracy, offering speed ups by a factor of 5–7. These ideas are empirically extended to the Milstein scheme, and the non central χ^2 distribution for the Cox-Ingersoll-Ross process, offering speed-ups of a factor of 250 or more.

CCS Concepts: • **Mathematics of computing** → **Numerical analysis**; *Mathematical software performance*; **Statistical software**; Approximation; Continuous functions; • **Computing methodologies** → Vector / streaming algorithms; *Uncertainty quantification*; *Uncertainty quantification*; *Massively parallel and high-performance simulations*.

Additional Key Words and Phrases: approximations, random variables, inverse cumulative distribution functions, random number generation, the Gaussian distribution, geometric Brownian motion, the Cox-Ingersoll-Ross process, the non central χ^2 distribution, multilevel Monte Carlo, the Euler-Maruyama scheme, the Milstein scheme, and high performance computing.

ACM Reference Format:

Michael Giles and Oliver Sheridan-Methven. 2023. Approximating inverse cumulative distribution functions to produce approximate random variables. 1, 1 (May 2023), 29 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Random number generation is a fundamental building block of a wide range of computational tasks, in particular financial simulations [7, 33, 44, 76]. A frequent performance bottleneck is the generation of random numbers from a specific statistical distribution such as in Monte Carlo simulations. While computers have many excellent and fast implementations of random number generators for the uniform distribution on the interval (0, 1), sampling random variables from a generic distribution is often computationally much more expensive.

We consider random variables produced using the inverse transform method [33], which enables sampling from any uni-variate distribution, and thus is very widely applicable. Additionally, the inverse transform method is crucial for

Authors' address: Michael Giles, mike.giles@maths.ox.ac.uk; Oliver Sheridan-Methven, oliver.sheridan-methven@hotmail.co.uk, Mathematical Institute, Oxford University, Oxford, UK.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

quasi-Monte Carlo simulations [32, 48] as no samples are rejected and the low-discrepancy property is preserved [70]; such applications are common in financial simulations [44, 76]. Furthermore, we demonstrate how the inverse transform method is particularly well suited to analysis, and that it naturally provides a coupling mechanism for multilevel Monte Carlo applications for a range of stochastic processes, statistical distributions, and numerical schemes.

Our analysis focuses on the Gaussian distribution (a.k.a. the Normal distribution), and the motivations are threefold. Firstly, the distribution is representative of several continuous distributions, and due to the central limit theorem is often the limiting case. Secondly, it is analytically tractable, often admitting exact results or approximations. Lastly, it is ubiquitous in both academic analysis and scientific computation, with its role cemented within Itô calculus and financial simulations.

To produce Gaussian random variables will require the Gaussian distribution's inverse cumulative distribution function. Constructing approximations accurate to machine precision has long been investigated by the scientific community [8, 28, 36, 51, 57, 73], where the *de facto* routine implemented in most libraries is by Wichura [73]. While some applications may require such accurate approximations, for many applications such accuracy is excessive and unnecessarily costly, as is typically the case in Monte Carlo simulations.

Using the inverse cumulative distribution function to generate random samples is commonly known as the *inverse transform method* [33, 63], and is applicable for any probability distribution. However, specific to the Gaussian distribution, there are other alternative approaches for generating Gaussian random variables, with popular alternatives including the Ziggurat [50], Box-Muller [12], and Marsaglia [49] methods. The Ziggurat and Marsaglia methods are unsuitable for the framework we develop as they are rejection based methods, prohibiting use in quasi-Monte Carlo applications, and also adversely affecting efficient implementations on vectorised hardware. The Box-Muller method (and the related Marsaglia method) suffers from the same pitfalls as the inverse transform method, namely the evaluation of expensive intermediate functions when computing tail values, and hence provides no advantage over the inverse transform method.

To alleviate the cost of exactly sampling from the Gaussian distribution, a popular circumvention is to substitute these samples with random variables with similar statistics. The bulk of such schemes follow a moment matching procedure, where the most well known is to use Rademacher random variables (which take the values ± 1 with equal probability [45, page XXXII], giving rise to the weak Euler-Maruyama scheme), matching the mean and variance. Another is to sum twelve uniform random variables and subtract the mean [55, page 500], which also matches the mean and variance, and is still computationally cheap. The most recent work in this direction is by Müller et al. [54], who produce either a three or four point distribution, where the probability mass is positioned so the resulting distribution's moments match the lowest few moments of the Gaussian distribution; as in this paper, they combine this with use of the multilevel Monte Carlo (MLMC) method but in a way which is not directly comparable.

The direction we follow in this paper is more closely aligned to the work by Giles et al. [29, 30], whose analysis proposes a cost model for producing the individual random bits constituting a uniform random number. They truncate their uniforms to a fixed number of bits and then add a small offset before using the inverse transform method. The nett result from this is to produce a piecewise constant approximation, where the intervals are all of equal width, and the values are the midpoint values of the respective intervals.

The work we present directly replaces random variables produced from the inverse transform method using the exact inverse cumulative distribution function, with those produced using an approximation to the inverse cumulative distribution function. While this is primarily motivated by computational savings for applications dominated by the

cost of sampling random variables, our framework encompasses the distributions produced from the various moment matching schemes and the truncated bit schemes.

Having a framework capable of producing various such distributions has several benefits. The first is that by refining our approximation, we can construct distributions resembling the exact distribution to an arbitrary fidelity. This allows for a trade-off between computational savings, and a lower degree of variance between the exact distribution and its approximation. This naturally introduces two tiers of simulations: those using a cheap but approximate distribution, and those using an expensive but near-exact distribution. This immediately facilitates the multilevel Monte Carlo setting of Giles [27], where fidelity and cost are balanced to minimise the computational time. As an example, we will see in section 4 that Rademacher random variables, while very cheap, are too crude to exploit any savings possible with multilevel Monte Carlo, whereas our higher fidelity approximations can fully exploit the possible savings.

The second benefit of our approach is that while the approximations are specified mathematically, their implementations are left unspecified. This flexibility facilitates constructing approximations which can be tailored to a specific hardware or architecture. We will present two approximations, whose implementations can gain speed by transitioning the work load from primarily using the floating point processing units to instead exploiting the cache hierarchy. Further to this, our approximations are designed with vector hardware in mind, and are non branching, and thus suitable for implementation using single instruction multiple data (SIMD) instructions, (including Arm’s new scalable vector extension (SVE) instruction set). Furthermore, for hardware with very large vectors, such as the 512 bit wide vectors on Intel’s AVX-512 and Fujitsu’s Arm-based A64FX (such as those in the new Fugaku supercomputer), we demonstrate implementations unrivalled in their computational performance on the latest hardware. Previous work using low precision bit wise approximations targetted at reconfigurable field programmable gate arrays has previously motivated the related works by Brugger et al. [15], Omland et al. [58], and Cheung et al. [19].

We primarily focus our attention on the analysis and implementation of two approximations: a piecewise constant approximation using equally sized intervals, and a piecewise linear approximation on a geometric sequence of intervals. The former is an extension of the work by Giles et al. [30, theorem 1] to higher moments, while the latter is a novel analysis, capable of both the highest speeds and fidelities. Although we will demonstrate how these are incorporated into a multilevel Monte Carlo framework, the subsequent analysis is performed by Giles and Sheridan-Methven [31, 67] and omitted from this work.

Having outlined our approximation framework and the incorporation of approximate random variables within a nested multilevel Monte Carlo scheme, we quantify the savings a practitioner can expect. For a geometric Brownian motion process from the Black-Scholes model [11] using Gaussian random variables, savings of a factor of 5–7 are possible. Furthermore, for a Cox-Ingersoll-Ross process [20] using non-central χ^2 random variables, savings of a factor of 250 or more are possible. These savings are benchmarked against the highly optimised Intel C/C++ Maths Kernel Library for the Gaussian distribution, and against the CDFLIB library [14, 16] in C and Boost library [3] in C++ for the non-central χ^2 distribution.

Section 2 introduces and analyses our two approximations, providing bounds on the error of the approximations. Section 3 discusses the production of high performance implementations, the code for which is collected into a central repository maintained by Sheridan-Methven [65, 66]. Section 4 introduces multilevel Monte Carlo as a natural application for our approximations, and demonstrates the possible computational savings. Section 5 extends our approximations to the non central χ^2 distribution, representing a very expensive parametrised distribution which arises in simulations of the Cox-Ingersoll-Ross process. Lastly, section 6 presents the conclusions from this work.

2 APPROXIMATE GAUSSIAN RANDOM VARIABLES

The inverse transform method [33, 2.2.1] produces univariate random variable samples from a desired distribution by first generating a uniform random variable U on the unit interval $(0, 1)$, and then computing $C^{-1}(U)$, where C^{-1} is the inverse of the distribution's cumulative distribution function (sometimes called the percentile or quantile function). We will focus on sampling from the Gaussian distribution, whose inverse cumulative distribution function we denote by $\Phi^{-1}: (0, 1) \rightarrow \mathbb{R}$ (some authors use N^{-1}), and similarly whose cumulative distribution function and probability density function we denoted by Φ and ϕ respectively.

Our key proposal is to use the inverse transform method with an approximation to the inverse cumulative distribution function. As the resulting distribution will not exactly match the desired distribution, we call random variables produced in this way *approximate random variables*, and those without the approximation as *exact random variables* for added clarity. In general, we will denote exact Gaussian random variables by Z and approximate Gaussian random variables by \tilde{Z} . The key motivation for introducing approximate random variables is that they are computationally cheaper to generate than exact random variables. Consequently, our key motivating criterion in forming approximations will be a simple mathematical construction, hoping this fosters fast implementations. As such there is a trade-off between simplicity and fidelity, where we will primarily be targetting simplicity.

In this section, we present two approximation types: a piecewise constant, and a piecewise linear. For both we will bound the L^p error, focusing on their mathematical constructions and analyses. Their implementations and utilisation will be detailed in sections 3 and 4. Both analyses will share and frequently use approximations for moments and tail values of the Gaussian distribution, which we gather together in section 2.1. Thereafter, the piecewise constant and linear approximations are analysed in sections 2.2 and 2.3.

2.1 Approximating tail values and high order moments

We require bounds on the approximations of tail values and high order moments. We use the notation of Giles et al. [30] that $f(z) \approx g(z)$ denotes $\lim_{z \rightarrow \infty} \frac{f(z)}{g(z)} = 1$. Our key results are lemmas 2.1 and 2.2 which bound the tail values and high order moments. Lemma 2.1 is an extension of a similar result by Giles et al. [30, lemma 7], extended to give enclosing bounds. Similarly, lemma 2.2 is partly an extension of a result by Giles et al. [30, lemma 9], but extended to arbitrarily high moments rather than just the second. As such, neither of these lemmas are particularly noteworthy in themselves and their proofs resemble work by Giles et al. [30, appendix A]. Similarly, our resulting error bounds for the piecewise constant approximation in section 2.2 will closely resemble a related result by Giles et al. [30, theorem 1]. However, our main result for the piecewise linear approximation in section 2.3 is novel and will require these results, and thus we include them here primarily for completeness.

LEMMA 2.1. Defining $z_q := \Phi^{-1}(1-2^{-q})$, then for $q \geq 2$ we have the bounds $1 < \frac{2^q \phi(z_q)}{z_q} < (1 - z_q^{-2})^{-1}$ and for $q \geq 10$ $\sqrt{q \log 4 - \log(q\pi(81/32) \log 4)} < z_q < \sqrt{q \log 4}$.

PROOF. $2^{-q} = 1 - \Phi(z_q) = \int_{z_q}^{\infty} \phi(s) ds$. Integrating by parts using $\phi'(z) = -z\phi(z)$ gives

$$\int_{z_q}^{\infty} \phi(s) ds = \phi(z_q) \frac{1}{z_q} - \int_{z_q}^{\infty} \frac{1}{s^2} \phi(s) ds = \phi(z_q) \left(\frac{1}{z_q} - \frac{1}{z_q^3} \right) + \int_{z_q}^{\infty} \frac{1}{4s^4} \phi(s) ds,$$

which proves that $\phi(z_q) \frac{1}{z_q} > 2^{-q} > \phi(z_q) \left(\frac{1}{z_q} - \frac{1}{z_q^3} \right)$ and hence we obtain the first inequality.

For $q \geq 10$, $z_q > 3$ and so $(1 - z_q^{-2})^{-1} < \frac{9}{8}$. The inequality $\frac{9}{8} 2^{-q} z_q > \phi(z_q) > z_q$ then gives

$$\sqrt{q \log 4 - \log\left(2\pi \frac{81}{64} z_q^2\right)} < z_q < \sqrt{q \log 4 - \log\left(2\pi z_q^2\right)}.$$

Inserting the upper bound $z_q < \sqrt{q \log 4}$ into the left hand term, and the lower bound $z_q^2 > 1/2\pi$ in the right hand term, gives the desired second inequality. \square

LEMMA 2.2. For integers $p \geq 2$ we have $\int_0^z \phi(s)^{1-p} ds \approx \frac{\phi(z)^{1-p}}{(p-1)z}$ and $\int_z^\infty (s-z)^p \phi(s) ds \approx \frac{p! \phi(z)}{z^{p+1}}$ as $z \rightarrow \infty$.

PROOF. Applying L'Hôpital's rule gives


$$\lim_{z \rightarrow \infty} \frac{\int_0^z \phi(s)^{1-p} ds}{\left(\frac{\phi(z)^{1-p}}{(p-1)z}\right)} = \lim_{z \rightarrow \infty} \frac{\phi(z)^{1-p}}{\phi(z)^{1-p} \left(1 - \frac{1}{(p-1)z^2}\right)} = 1.$$

Similarly, applying L'Hôpital's rule $p+1$ times gives

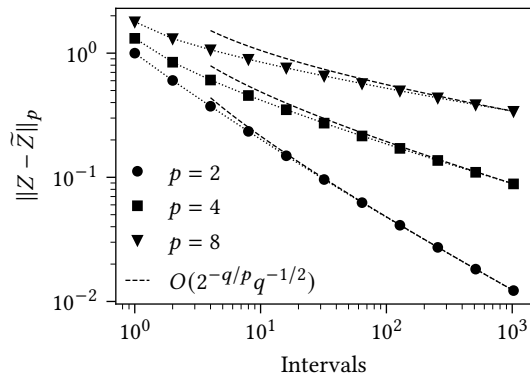
$$\lim_{z \rightarrow \infty} \frac{\int_z^\infty (s-z)^p \phi(s) ds}{\left(\frac{p! \phi(z)}{z^{p+1}}\right)} = \lim_{z \rightarrow \infty} \frac{(-1)^p p! \phi(z)}{(-1)^p p! \phi(z) (1 + O(1/z^2))} = 1. \quad \square$$

2.2 Piecewise constant approximations on equal intervals

Mathematically it is straightforward to motivate a piecewise constant approximation as the simplest possible approximation to use, especially using equally spaced intervals. As the range of values will go from a continuous to a discrete set, we say the distribution has become *quantised* and denote our approximation as $Q \approx \Phi^{-1}$, where for a uniform random variable $U \sim \mathcal{U}(0, 1)$ we have $Z := \Phi^{-1}(U)$ and $\tilde{Z} := Q(U)$. A preview of such an approximation is shown in figure 1(a), where the error will be measured using the L^p norm $\|f\|_p := (\int_0^1 |f(u)|^p du)^{1/p}$.



(a) A piecewise constant approximation using 8 intervals, and the corresponding QQ-plot using 10^3 samples.



(b) The L^p error and the bound from theorem 2.3.

Fig. 1. The piecewise constant approximation $\tilde{Z}_k^{L^1}$ from corollary 2.4 with equally spaced intervals, and its error.

Before presenting theorem 2.3, we can briefly comment on the error seen in figure 1(b). Specifically looking at the root mean squared error (RMSE), corresponding to the L^2 norm, we can see that increasing the number of intervals from 2 to 10^3 gives a drop of 10^2 in the RMSE. For our multilevel Monte Carlo applications in section 4, having approximately 1000 intervals gives a very good fidelity, whereas Rademacher random variables have a very low fidelity. Being able to achieve a reasonable fidelity from our approximation ensures that we achieve the largest portion of the possible computational savings offered by our approximations.

As we have already mentioned, the piecewise constant approximation is closely related to the resulting approximation produced by Giles et al. [30], whose approximation arises from considering uniform random variables truncated to a finite number of bits of precision. Thus our main result from this section, theorem 2.3, closely resembles a related result by Giles et al. [30, theorem 1]. To put our extension into context, we paraphrase the similar result from Giles et al. [30], which is that for a piecewise constant approximation using 2^q intervals, for some integer $q \gg 1$, then for constant values equal to each interval's midpoint value they have $\|Z - \tilde{Z}\|_2^2 = O(2^{-q}q^{-1})$. Our result from theorem 2.3 extends this to $\|Z - \tilde{Z}\|_p^p = O(2^{-q}q^{-p/2})$ for $p \geq 2$, and numerous other possible constant values other than the midpoint's. Our result enables us to increase the order of the error to arbitrarily high norms and is interesting in its own right. It shows that as the intervals become increasing small (corresponding to $q \rightarrow \infty$), the dominant term effecting the L^p error is the geometric decay $O(2^{-q/p})$ due to the error in the two end intervals, and thus the convergence exists but is slower in higher norms, with the polynomial $O(q^{-1/2})$ term being comparatively negligible (as we can see in figure 1(b)). Additionally, in the related analysis incorporating approximate random variables into a nested multilevel Monte Carlo framework by Giles and Sheridan-Methven [31, 67], their bounds on the variance of the multilevel Monte Carlo correction term (discussed more in section 4) rely on the existence of the L^p error for $p > 2$. Hence, while this strengthening of the result may appear only slight, it is crucial for nested multilevel Monte Carlo.

We can now present our main result concerning piecewise constant approximations, namely theorem 2.3. In this we will leave the interval values largely unspecified, and later demonstrate in corollary 2.4 several choices fit within the scope of theorem 2.3.

THEOREM 2.3. *Let a piecewise constant approximation $Q \approx \Phi^{-1}$ use 2^q equally spaced intervals for some integer $q > 1$. Denote the intervals $I_k := (k2^{-q}, (k+1)2^{-q}) \equiv (u_k, u_{k+1})$ for $k \in \{0, 1, 2, \dots, K\}$ where $K \equiv 2^q - 1$. On each interval the approximation constant is $Q_k := Q(u)$ for $u \in I_k$. We assume there exists a constant C independent of q such that:*

- (1) $Q_k = -Q_{K-k}$ for $k \in \{0, 1, 2, \dots, K\}$.
- (2) $\Phi^{-1}(u_k) \leq Q_k \leq \Phi^{-1}(u_{k+1})$ for $k \in \{1, 2, \dots, K-1\}$.
- (3) $\Phi^{-1}(u_K) \leq Q_K \leq \Phi^{-1}(u_K) + Cq^{-1/2}$.

Then for any even integer $p \geq 2$ we have $\|Q - \Phi^{-1}\|_p^p = O(2^{-q}q^{-p/2})$ as $q \rightarrow \infty$.

PROOF. Defining $A_k := \int_{u_k}^{u_{k+1}} |\Phi^{-1}(u) - Q_k|^p du$ for $k \in \{0, 1, \dots, K\}$, then we have $\|Q - \Phi^{-1}\|_p^p = \sum_{k=0}^K A_k$. We use condition (1) to reduce our considerations to the domain $(\frac{1}{2}, 1)$, where Φ^{-1} is strictly positive and convex, and thus obtain $\|Q - \Phi^{-1}\|_p^p = 2 \sum_{k=2^{q-1}}^{K-1} A_k + 2A_K$. As Φ^{-1} is convex, then from the intermediate value theorem there exists a $\xi_k \in I_k$ such that $Q_k = \Phi^{-1}(\xi_k)$. Noting that $\frac{d}{dz} \Phi^{-1}(z) = \frac{1}{\phi(\Phi^{-1}(z))}$, then from the mean value theorem for any $u \in [u_k, \xi_k]$ there exists an $\eta_k \in [u_k, \xi_k]$ such that $\Phi^{-1}(u) - \Phi^{-1}(\xi_k) = \frac{u - \xi_k}{\phi(\Phi^{-1}(\eta_k))}$. Furthermore, as ϕ is monotonically decreasing in $(\frac{1}{2}, 1)$, then introducing $z_k := \Phi^{-1}(u_k)$ we have $|\Phi^{-1}(u) - Q_k| \leq \frac{2^{-q}}{\phi(z_{k+1})}$. An identical argument follows

for any $u \in [\xi_k, u_{k+1})$ giving the same bound. Using this to bound $\sum_{k=2^{q-1}}^{K-1} A_k$ in our expression for $\|Q - \Phi^{-1}\|_p^p$ gives

$$\sum_{k=2^{q-1}}^{K-1} A_k \leq \sum_{k=2^{q-1}}^{K-1} \left(\frac{2^{-q}}{\phi(z_{k+1})} \right)^p 2^{-q} \leq 2^{-qp-q} \sum_{k=2^{q-1}}^{K-1} \phi(z_{k+1})^{-p} \leq 2^{-pq} \int_{\frac{1}{2}}^{1-2^{-q}} \phi(\Phi^{-1}(u))^{-p} du + 2^{-q(p+1)} \phi(z_K)^{-p},$$

where the last bound comes from considering a translated integral. Changing integration variables the integral becomes $\int_0^{z_K} \phi(z)^{1-p} dz$, from which we can use lemma 2.2 to give

$$\sum_{k=2^{q-1}}^{K-1} A_k \leq \frac{2^{-q} z_K^{-p}}{p-1} \left(\frac{2^q \phi(z_K)}{z_K} \right)^{1-p} + 2^{-q} z_K^{-p} \left(\frac{2^q \phi(z_K)}{z_K} \right)^{-p} \leq 2^{-q} \left(\frac{p}{p-1} \right) z_K^{-p},$$

where the last bound follows from lemma 2.1. Turning our attention to the final interval's contribution A_K , then using Jensen's inequality, condition (3), and lemmas 2.1 and 2.2 we obtain

$$A_K \leq 2^{p-1} \int_{z_K}^{\infty} |Q_K - z|^p \phi(z) dz + 2^{p-1} \int_{u_K}^1 |Q_K - z_K|^p du \leq 2^{p-q-1} p! z_K^{-p} + 2^{p-q-1} C^p q^{-p/2}.$$

Combining our two bounds for $\sum_{k=2^{q-1}}^{K-1} A_k$ and A_K into our expression for $\|Q - \Phi^{-1}\|_p^p$ we obtain

$$\|Q - \Phi^{-1}\|_p^p \leq 2^{-q+1} \left(\frac{p}{p-1} \right) z_K^{-p} + 2^{p-q} p! z_K^{-p} + 2^{p-q} C^p q^{-p/2} \leq O(2^{-q} q^{-p/2}),$$

where the coefficients inside the O -notation are only a function of p and not of q . \square

COROLLARY 2.4. *For approximate Gaussian random variables \tilde{Z} using constants \tilde{Z}_k on intervals I_k constructed as either*

$$\tilde{Z}_k^E := \mathbb{E}(Z \mid \Phi(Z) \in I_k), \quad \tilde{Z}_k^C := \Phi^{-1} \left(\frac{\max I_k + \min I_k}{2} \right), \quad \text{or} \quad \tilde{Z}_k^I := \begin{cases} \Phi^{-1}(\min I_k) & \text{if } \min I_k \geq 0.5 \\ \Phi^{-1}(\max I_k) & \text{if } \max I_k < 0.5, \end{cases}$$

then all finite moments of $|\tilde{Z}|$ are uniformly bounded, and conditions (1–3) of theorem 2.3 are satisfied.

PROOF. As all three approximations are anti-symmetric about $u = 1/2$, and therefore all satisfy conditions (1, we consider the domain $(\frac{1}{2}, 1)$ where Z and \tilde{Z} are both positive. Letting the interval I_k lie in this domain, then from Jensen's inequality and the law of iterated expectations we obtain

$$\mathbb{E}((\tilde{Z}^E)^n) = \mathbb{E}(\mathbb{E}_{I_k}(Z \mid \Phi(Z) \in I_k))^n \leq \mathbb{E}(\mathbb{E}_{I_k}(Z^n \mid \Phi(Z) \in I_k)) = \mathbb{E}(Z^n) < \infty,$$

for any $1 \leq n < \infty$, where the last inequality is a standard result. Mirroring this result to the full domain $(0, 1)$, we can directly see that the moments of $|\tilde{Z}^E|$ are uniformly bounded. Furthermore, as Φ^{-1} is increasing and convex in $[\frac{1}{2}, 1)$, then the value \tilde{Z}_k^E is an upper bound on \tilde{Z}_k^C and \tilde{Z}_k^I , and so these too have uniformly bounded moments.

It is immediately clear that all three choices satisfy condition 2) and the lower bound in 3). As \tilde{Z}_k^E is an upper bound for the other two choices, it will suffice to show that this satisfies the upper bound in 3). Inspecting the difference between the smallest value in the final I_K interval, namely z_K , and the average value $Q_K = \mathbb{E}(Z \mid \Phi(z) \in I_K)$, we obtain

$$Q_K - z_K = 2^q \int_{u_K}^1 (\Phi^{-1}(u) - \Phi^{-1}(u_K)) du = 2^q \int_{z_K}^{\infty} (z - z_K) \phi(z) dz < 2^q \frac{\phi(z_K)}{z_K^2} < \frac{1}{z_K} (1 - z_K^{-2})^{-1} \leq \frac{9}{8z_K},$$


where the first inequality follows from lemma 2.2, the second from lemma 2.1, and the last uses $(1 - z_K^{-2})^{-1} \leq \frac{9}{8}$ for $q \geq 10$. We bound the final $\frac{1}{z_K}$ term using lemma 2.1 for $q \geq 10$ to give $\frac{1}{z_K} < 1.1 q^{-1/2}$ and therefore $Q_K \leq z_K + 1.25 q^{-1/2}$. \square

We remark that the \widetilde{Z}_k^C construction from corollary 2.4 is the same as the value used in the analysis of Giles et al. [30, (4)]. In any piecewise constant approximations we use, such as that in figure 1(a), we will use the constants defined by the \widetilde{Z}_k^E construction from corollary 2.4. Furthermore, we can see that our bound from theorem 2.3 appears tight in figure 1(b).

2.3 Piecewise linear approximations on geometric intervals

Looking at the piecewise constant approximation in figure 1(a), it is clear there are two immediate improvements that can be made. The first is to use a piecewise linear approximation, which is considerably more appropriate for the central region. Secondly, the intervals should not be of equal sizes, but denser near the singularities at either end. We will make both these modifications in a single step, where we will construct a piecewise linear approximation with geometrically decreasing intervals which are dense near the singularities. For brevity we will denote this just as the piecewise linear approximation. An example piecewise linear approximation using 8 intervals is shown in figure 2(a). The precise nature of the interval widths, and how the linear functions are fitted will be detailed shortly, but by direct comparison against figure 1(a) it is clear that the fidelity of a piecewise linear approximation is much better than the piecewise constant.

The main result from this section will be theorem 2.5, which will bound the L^p error of our piecewise linear approximation. The proof will proceed in a similar fashion the proof of theorem 2.3, where we will bound the sum of the central intervals and the end intervals separately. For the central intervals we will use the Peano kernel theorem to bound the point wise error, and in the end intervals several results will be a mixture of exact results and bounds from lemmas 2.1 and 2.2.



(a) A piecewise linear approximation using 8 intervals, and the corresponding QQ-plot using 10^3 samples. (b) The L^2 error for various polynomial orders, with the number of intervals in $(0, \frac{1}{2})$ labeled.

Fig. 2. The piecewise linear approximation and its error from theorem 2.5 with geometric intervals using $r = \frac{1}{2}$.

THEOREM 2.5. *For an approximation $D \approx \Phi^{-1}$ which is anti-symmetric about $u = 1/2$, with K intervals in $(0, \frac{1}{2})$, we define the k -th interval $I_k := [\frac{r^k}{2}, \frac{r^{k-1}}{2})$ for $k \in \{1, 2, \dots, K-1\}$ and $I_K := (0, \frac{r^{K-1}}{2})$ for some decay rate $r \in (0, 1)$. Each interval uses a piecewise linear approximation $D_k(u) \equiv D(u)$ for any $u \in I_k$. The gradient and intercept in each interval is set by the L^2 minimisation $D_k := \operatorname{argmin}_{D' \in \mathcal{P}_1} \int_{I_k} |\Phi^{-1}(u) - D'(u)|^2 du$ where \mathcal{P}_1 is the set of all 1-st order polynomials.*

Then we have for any $2 \leq p < \infty$

$$\|D - \Phi^{-1}\|_p^p = O((1-r)^{2p}) + O(r^{K-1} \log^{-p/2}(r^{1-K} \sqrt{2/\pi})) = O((1-r)^{2p}) + o(r^{K-1}).$$

PROOF. Considering the domain $(0, \frac{1}{2})$, we split the contribution into those from the intervals without the singularity, and that from the final interval with the singularity, where

$$\|D - \Phi^{-1}\|_p^p = 2 \sum_{k=1}^{K-1} \int_{I_k} |D_k(u) - \Phi^{-1}(u)|^p du + 2 \int_{I_K} |D_K(u) - \Phi^{-1}(u)|^p du,$$

where the factors of 2 correct for us only considering the lower half of the domain $(0, 1)$.

Beginning with the non-singular intervals, we express the pointwise error using the Peano kernel theorem [42, 60], which we will later bound. For notational simplicity, we denote the approximated function as f , where $f \equiv \Phi^{-1}$, and a given interval as $[a, b] \equiv I_k$. The L^2 optimal linear approximation is $\alpha(f) + \beta(f)u$ for $u \in [a, b]$ where α and β are functionals. The point wise error $f(u) - \alpha(f) - \beta(f)u$ is a linear mapping L acting on f where $L(f)(u) := f(u) - \alpha(f) - \beta(f)u$. By construction L annihilates linear functions, so the Peano kernel is $k(\xi; u) := (u - \xi)^+ - \alpha((\cdot - \xi)^+) - \beta((\cdot - \xi)^+)u \equiv (u - \xi)^+ - \bar{\alpha}(\xi) - \bar{\beta}(\xi)u$ for $\xi \in [a, b]$ where we defined $\bar{\alpha}(\xi) := \alpha((\cdot - \xi)^+)$ and similarly $\bar{\beta}(\xi)$. The pointwise error is $\varepsilon(u) := L(f)(u) = \int_a^b k(\xi; u) f''(\xi) d\xi$.

To determine the intercept and gradient, we use that they are L^2 optimal, and so the functional derivatives of $\int_a^b |\varepsilon(u)|^2 du$ with respect to α and β are zero, giving the simultaneous equations

$$\alpha(f)(b-a) + \beta(f) \left(\frac{b^2 - a^2}{2} \right) = \int_a^b f(u) du \quad \text{and} \quad \alpha(f) \left(\frac{b^2 - a^2}{2} \right) + \beta(f) \left(\frac{b^3 - a^3}{3} \right) = \int_a^b u f(u) du.$$

It is important to notice that because we chose the L^2 norm, these are a set of linear simultaneous equations, and thus α and β are linear functionals, thus showing that L is linear, (a requirement of the Peano kernel theorem). Evaluating these for the kernel function ($f \rightarrow (\cdot - \xi)^+$) gives

$$\bar{\alpha}(\xi) = -\frac{(b-\xi)^2((b+a)\xi - 2a^2)}{(b-a)^3} \quad \text{and} \quad \bar{\beta}(\xi) = \frac{(b-\xi)^2(2\xi + b - 3a)}{(b-a)^3}.$$

Thus, the pointwise error is

$$\varepsilon(u) = \int_a^b ((u-\xi)^+ - \bar{\alpha}(\xi) - \bar{\beta}(\xi)u) f''(\xi) d\xi = (b-a)^2 \int_0^1 \left((\tilde{u} - \tilde{\xi})^+ - (1 - \tilde{\xi}^2)(\tilde{\xi} + \tilde{u}) \right) f''((b-a)\tilde{\xi} + a) d\tilde{\xi},$$

where to achieve the last equality we rescaled our interval $[a, b] \rightarrow [0, 1]$ and variables $\eta \rightarrow \tilde{\eta}$ where $\tilde{\eta} := \frac{\eta-a}{b-a}$. Taking the absolute value and applying Jensen's inequality immediately gives

$$|\varepsilon(u)| \leq -6(b-a)^2 \frac{d^2 \Phi^{-1}}{du^2}(a) = -6 \left(\frac{1-r}{r} \right)^2 a^2 \frac{d^2 \Phi^{-1}}{du^2}(a) \leq 2.59 \left(\frac{1-r}{r} \right)^2,$$

where for the first inequality we used that $\frac{d^2}{du^2} \Phi^{-1}$ is maximal at the lower boundary, and for the second inequality we bound this by the maximum with respect to a (which is at $a \approx 0.177$). Using this expression for the pointwise error in our summation of the non-singular intervals gives (as $r \rightarrow 1$)

$$\sum_{k=1}^{K-1} \int_{I_k} |D_k(u) - \Phi^{-1}(u)|^p du = \sum_{k=1}^{K-1} \int_{I_k} |\varepsilon(u)|^p du \leq 2.59^p \left(\frac{1-r}{r} \right)^{2p} \int_0^{\frac{1}{2}} du = O((1-r)^{2p}).$$

We now consider the interval $[0, b] \equiv I_K$ containing the singularity at 0, which has intercept and gradient $\alpha = \frac{1}{b} \int_0^b \Phi^{-1}(u) du - \frac{b\beta}{2}$ and $\beta = \frac{12}{b^3} \int_0^b (u - \frac{b}{2}) \Phi^{-1}(u) du$. These integrals can be calculated exactly (with a change of

variables), where denoting $z_b := \Phi^{-1}(b)$ we obtain

$$\alpha = \frac{2\phi(z_b)}{b} - \frac{3\Phi(\sqrt{2}z_b)}{b^2\sqrt{\pi}} \quad \text{and} \quad \beta = \frac{6}{b^3} \left(\frac{\Phi(\sqrt{2}z_b)}{\sqrt{\pi}} - b\phi(z_b) \right).$$

For the gradient, as the interval becomes ever smaller and $b \rightarrow 0$, we can use lemma 2.1 to give

$$\beta \approx -\frac{6}{b^3} \left(\phi^2(z_b) \left(\frac{1}{z_b} - \frac{1}{2z_b^3} \right) + b\phi(z_b) \right) \approx -\frac{6z_b}{b} \left(\frac{1}{2z_b^2} + O(z_b^{-4}) \right) \approx \frac{-3}{bz_b},$$

where in the first approximation we used $\phi(\sqrt{2}z) \equiv \sqrt{2\pi}\phi^2(z)$. Interestingly, this means the range of values $\beta b \approx -\frac{3}{z_b} \rightarrow 0$, and our approximation “flattens” relative to the interval $[0, b]$ as $b \rightarrow 0$.

With the intercept and gradient in the singular interval I_K known exactly, we define the two points u_- and u_+ where the error is zero, where $0 < u_- < u_+ < b$, and there are two as Φ^{-1} is concave in $(0, \frac{1}{2})$. Corresponding to these we define $z_- := \Phi^{-1}(u_-)$ and $z_+ := \Phi^{-1}(u_+)$, where $-\infty < z_- < z_+ < z_b < 0$. Thus in the singular interval we have

$$\int_{I_K} |D_K(u) - \Phi^{-1}(u)|^p du = \int_{-\infty}^{z_-} |z - \alpha - \beta\Phi(z)|^p \phi(z) dz + \int_{z_-}^{z_b} |z - \alpha - \beta\Phi(z)|^p \phi(z) dz.$$

Using lemmas 2.1 and 2.2, then for the first of these integrals we obtain

$$\int_{-\infty}^{z_-} |z - \alpha - \beta\Phi(z)|^p \phi(z) dz \leq \int_{-\infty}^{z_-} |z - z_-|^p \phi(z) dz \leq \int_{-\infty}^{z_b} |z - z_b|^p \phi(z) dz \approx \frac{p! \phi(z_b)}{|z_b|^{p+1}} = O\left(\frac{b}{|z_b|^p}\right),$$

and for the second integral we similarly obtain

$$\int_{z_-}^{z_b} |z - \alpha - \beta\Phi(z)|^p \phi(z) dz \leq \int_{-\infty}^{z_b} |z_+ - \alpha - \beta\Phi(z)|^p \phi(z) dz \leq |\beta b|^p \int_0^b du \approx \frac{3^p}{|z_b|^p} b = O\left(\frac{b}{|z_b|^p}\right).$$

Combining the results for the central intervals and the singular interval we obtain

$$\|D - \Phi^{-1}\|_p^p = O((1-r)^{2p}) + O\left(\frac{b}{|z_b|^p}\right) = O((1-r)^{2p}) + O(r^{K-1}(\log(r^{1-K}\sqrt{2/\pi}))^{-p/2}) = O((1-r)^{2p}) + o(r^{K-1}),$$

where for the second equality we used lemma 2.1 in the limit $r^{K-1} \rightarrow 0$. \square

We can see from theorem 2.5 that the $O((1-r)^{2p})$ term comes from the central regions, and is reduced by taking $r \rightarrow 1$, and the $o(r^{K-1})$ term is from the singular interval, and is reduced by taking $r^{K-1} \rightarrow 0$. The key point of interest with this result is that the error from the central regions and the singular region is decoupled. In order to decrease the overall error, it is not sufficient to only increase the number of intervals ($K \rightarrow \infty$), which would only improve the error from the singular interval, but the decay rate must also be decreased ($r \rightarrow 1$). The independence and interplay of these two errors is important for balancing the fidelity between the central and edge regions.

It is possible to generalise this construction and analysis to piecewise polynomial approximations, constructing the best L^2 approximation on each interval. We could also require the approximations to be continuous over the entire interval $(0, 1)$ by turning our previous minimisation into a constrained minimisation, with a coupled set of linear constraints. While such a continuous approximation may be more aesthetically pleasing, it is of no practical consequence for the inverse transform method, and by definition will have a worse overall L^2 error than the discontinuous approximation from theorem 2.5, and thus we choose not to do this.

The errors from piecewise polynomial approximations for various polynomial orders and interval numbers are shown in figure 2(b), where we have set the decay rate $r = \frac{1}{2}$. As we increase the number of intervals used, then for the piecewise linear function, the L^2 error plateaus to approximately 10^{-2} for 16 intervals, which is approximately equal to the error

for the piecewise constant approximation using 10^3 intervals. Thus we can appreciate the vast increase in fidelity that the piecewise linear approximation naturally offers over the piecewise constant approximation. Furthermore, this plateau demonstrates that the central regions are limiting the accuracy of the approximation. Inspecting the approximation using 16 intervals in $(0, \frac{1}{2})$, we see that as we increase the polynomial order from linear to cubic there is a considerable drop in the error approximately equal to a factor of 10^2 , where the piecewise cubic approximation is achieving L^2 errors of around 10^{-4} . However, we see that increasing the polynomial order any further does not lead to any significant reduction in the L^2 error, indicating that in this regime it is the error in the singular interval which is dominating the overall error.

3 HIGH PERFORMANCE IMPLEMENTATIONS

Our motivation for presenting the piecewise constant and linear approximations was to speed up the inverse transform method, where the approximations were inherently simple in their constructions, with the suggestion that any implementations would likely be simple and performant. However, it is not obvious that the implementations from various software libraries, especially heavily optimised commercial libraries such as Intel’s Maths Kernel Library, should be slow. In this section, we will mention how most libraries implement these routines, and why many are inherently ill posed for modern vector hardware. Our two approximations capitalise on the features where most libraries stumble, avoiding division and conditional branching, which will be the key to their success. We give a brief overview of their implementations in C and showcase their superior speed across Intel and Arm hardware. While we will briefly outline the implementations here in moderate detail, a more comprehensive suite of implementations and experiments, including Intel AVX-512 and Arm SVE specialisations, are hosted in a centralised repository by Sheridan-Methven [66].

3.1 The shortcomings in standard library implementations

Several libraries providers, both open source and commercial, offer implementations of the inverse Gaussian cumulative distribution function, (Intel, NAG, Nvidia, MathWorks, Cephes, GNU, Boost, etc.). The focus of these implementations is primarily ensuring near machine precision is achieved for all possible valid input values, and handling errors and edge cases appropriately. While some offer single precision implementations, most offer (or default to) double precision, and to achieve these precisions algorithms have improved over the years [8, 28, 36, 51, 57, 73]. The most widely used is by Wichura [73], which is used in GNU’s scientific library (GSL) [26], the Cephes library [53] (used by Python’s SciPy package [72]), and the NAG library [56] (NAG uses a slight modification).

These implementations split the function into two regions: an easy central region, and the difficult tail regions. Inside the central region a rational Padé approximation is used with seventh order polynomials [73]. In the tail regions, the square root of the logarithm is computed, and then a similar rational approximation performed on this. While these are very precise routines, they have several shortcomings with respect to performance.

The first shortcoming is the division operation involved in computing the rational approximation. Division is considerably more expensive than addition or multiplication, and with a much higher latency [25, 75]. Thus, for the very highest performance, avoiding division is preferable. (The implementation by Giles [28] already capitalises on this). Similarly, taking logarithms and square roots is expensive, making the tail regions very expensive, and so are best avoided too.

The second short coming is that the routine branches: requesting very expensive calculations for infrequent tail values, and less expensive calculations for more frequent central values. Branching is caused by using “if-else” conditions, but unfortunately this often inhibits vectorisation, resulting in compilers being unable to issue single

instruction multiple data (SIMD) operations [71]. Even if a vectorised SIMD implementation is produced, typically the results from both branches are computed, and the correct result selected *a posteriori* by predication/masking. Thus, irrespective of the input, both the very expensive and less expensive calculations may be performed, resulting in the relatively infrequent tail values producing a disproportionately expensive overall calculation. This effect is commonly known as *warp divergence* in GPU programs, and in our setting we will call this *vector divergence*. The wider the vector and the smaller the data type, (and thus the greater the amount of vector parallelisation), the worse the impact of vector divergence. Noting that half precision uses only 16 bits, both Intel’s AVX-512 and Fujitsu’s A64FX are 512 bits wide, and Arm’s SVE vectors can be up to 2048 bits [59, 69], the degree of vector parallelisation can be substantial, and the impact of conditional branching becomes evermore crippling as vector parallelisation increases.

In light of these two shortcomings, we will see our approximations can be implemented in a vectorisation friendly manner which avoids conditional branching and are homogenous in their calculations. Furthermore, their piecewise constant and polynomial designs avoid division operations, and only require addition, multiplication, and integer bit manipulations, and thus result in extremely fast executables. We will see that our piecewise constant approximation will rely on the high speed of querying the cache. Furthermore, for the piecewise linear approximation using a decay rate $r = \frac{1}{2}$ and 16 intervals, then in single precision all the coefficients can be held in 512 bit wide vector registers, bypassing the need to even query the cache, and thus are extremely fast provided the vector widths are sufficiently large.

3.1.1 Alternatives specific to the Gaussian distribution. Specifically for the Gaussian distribution, in section 1 we mentioned alternatives to the inverse transform method, namely the Ziggurat [50], Box-Muller [12], and Marsaglia [49] methods, which are all heavily used in practice. While we don’t detail these algorithms nor their implementations here, they also suffer from the same shortcomings as the inverse Gaussian cumulative distribution function. The Box-Muller and Marsaglia methods require the evaluation of expensive logarithmic, square-root, and trigonometric functions, where furthermore the logarithmic function is divergent at the origin (corresponding to the tails of the Gaussian distribution). Additionally, the rejection mechanism of the Ziggurat and Marsaglia methods immediately introduces conditional code execution.

Consequently, the Ziggurat and Marsaglia methods are ill positioned to capitalise on the benefits offered by vectorised SIMD hardware. For the remaining Box-Muller method, this would require producing at least 2 approximating functions (the product of the two functions giving the resultant approximate Gaussian random variable), compared to just the one when using the inverse transform method, without any obvious accuracy improvements. Consequently, none of these alternatives are expected to have a performance exceeding our approximations’ for the inverse Gaussian cumulative distribution function.

Weighing the drawbacks of these alternatives against the wider applicability of the inverse transform method to any target distribution, the supporting analysis presented herein, and the accompanying multilevel Monte Carlo framework [31], makes the inverse transform method the preferred approach.

3.2 Implementing the piecewise constant approximation

Splitting the domain $(0, 1)$ into the N intervals $[\frac{m}{N}, \frac{m+1}{N})$ zero indexed by $m \in \{0, 1, 2, \dots, N - 1\}$, the values for each interval are easily computed *a priori*, and stored in a lookup table. An example of how this looks in C is shown in code 1, where we use OpenMP (`omp.h`) to signal to the compiler that the for loop is suitable for vectorisation. This relies on typecasting to an integer, where any fractional part is removed. The benefit to such an implementation is that on modern chips a copy of the lookup table can readily fit within either the L1 or L2 caches, where 1024 values stored

in 64 bit double precision consume 8 kB. As an example, an Intel Skylake Xeon Gold 6140 CPU has L1 and L2 caches which are 32 kB and 2 MB respectively, and thus the lookup table can exploit the fast speeds of the L1 cache, which typically has latencies of 2–5 clock cycles.

```
#define LOOKUP_TABLE_SIZE 1024
const float lookup_table[LOOKUP_TABLE_SIZE] = {-3.3, -2.9, -2.8, ..., 2.8, 2.9, 3.3};

void piecewise_constant_approximation(const unsigned int n_samples,
                                     const float * restrict input,
                                     float * restrict output) {
    #pragma omp simd
    for (unsigned int n = 0; n < n_samples; n++)
        output[n] = lookup_table[(unsigned int) (LOOKUP_TABLE_SIZE * input[n])];
}
```

Code 1. C implementation of the piecewise constant approximation.

3.3 Implementing the piecewise linear approximation

The piecewise linear approximation is a bit more involved than the piecewise constant approximation. Not only do we have a polynomial to evaluate (albeit only linear), but the varying widths of the intervals means identifying which interval a given value corresponds to is more involved. Once the appropriate interval's polynomial's coefficients are found, evaluating the polynomial is trivial, where the linear polynomial can be evaluated using a fused multiply and add (FMA) instruction. Higher order polynomials can be similarly computed (e.g. with Horner's rule).

The primary challenge is determining which interval a given value corresponds to based on the construction from theorem 2.5 for a given decay rate r . As floating point numbers are stored in binary using their sign, exponent, and mantissa, the natural choice most amenable for computation is $r = \frac{1}{2}$, which we call the *dyadic* rate, producing the dyadic intervals shown in table 1. From table 1, we notice the intervals are only dense near the singularity at 0, but not at 1. This is not problematic, as we said in theorem 2.5 that the approximation is anti-symmetric about $u = 1/2$, and thus we can use the N intervals in $(0, \frac{1}{2})$, and if our input value is within $(\frac{1}{2}, 1)$, it is straight forward to simply negate the value computed for the input reflected about $\frac{1}{2}$ (equivalent to using the inverse complementary cumulative distribution function).

Table 1. Dyadic intervals and their corresponding array indices.

Index	0	1	2	...	n	...	N
Interval	$[\frac{1}{2}, 1)$	$[\frac{1}{4}, \frac{1}{2})$	$[\frac{1}{8}, \frac{1}{4})$...	$[\frac{1}{2^{n+1}}, \frac{1}{2^n})$...	$(0, \frac{1}{2^N})$

Such an implementation can handle any value in $(0, \frac{1}{2}) \cup (\frac{1}{2}, 1)$, but unfortunately $\frac{1}{2}$ is not included within this. While a mathematician may say the individual value $\frac{1}{2}$ has zero measure, from a computational perspective such a value is perfectly feasible input, and quite likely to appear in any software tests. Thus our implementation will correctly handle this value. The reader will notice that when we reflect an input value $x > \frac{1}{2}$ about $\frac{1}{2}$ by using $x \rightarrow 1 - x$, then the value $\frac{1}{2}$ will not be in any of the intervals indexed between 1 and N in table 1, but remains in the interval indexed by 0. Thus in our later implementations, the arrays of coefficients will always hold as their first entry (index 0) zero values ($\Phi^{-1}(\frac{1}{2}) = 0$) so $\frac{1}{2}$ is correctly handled. We found during the development of the implementation, being able

to correctly handle this value is of considerable practical importance, so we encourage practitioners to also correctly handle this value.

If we are using the dyadic decay rate, then the interval an input x belongs to is $\lceil -\log_2(x) \rceil - 1$. However, to compute this we need not take any logarithms, which are expensive to compute. We can obtain the same result by reading off the exponent bits in the floating point representation, treating these as an integer, and then correcting for the exponent bias. This only involves simple bit manipulations, and interpreting the bits in a floating point representation as an integer. In C this is a technique called type punning, and can be achieved either by pointer aliasing, using a union, or using specialised intrinsic functions. Of these, pointer aliasing technically breaks the strict aliasing rule in C (and C++) [41, 6.5.2.3] [68, pages 163–164]. Similarly, type punning with a union in C89 is implementation defined, whereas in C11 the bits are re-interpreted as desired. In the presented implementation we will leave this detail undefined and just use a macro to indicate the type punning. (The implementation approximating the Gaussian distribution uses pointer aliasing, whereas the later approximation to the non central χ^2 distribution in section 5 uses a union, to demonstrate both possibilities [66]).

Input: Floating point uniform random variable $U \in (0, 1)$.

Output: Floating point approximate Gaussian random variable \tilde{Z} .

Form predicate using $U > \frac{1}{2}$.

Reflect about $\frac{1}{2}$ to obtain $U \in (0, \frac{1}{2}]$.

Interpret U as an unsigned integer.

Read the exponent bits using bit wise AND.

Right shift away the mantissa's bits.

Obtain an array index by correcting for exponent bias.

Cap the array index to avoid overflow.

Read the polynomial's coefficients.

Re-interpret U as a float.

Form the polynomial approximation \tilde{Z} .

Correct sign of approximation based on the predicate.

Algorithm 1. Piecewise polynomial approximation using dyadic intervals.

Overall then, the general construction from theorem 2.5 is given in algorithm 1. Furthermore, a single precision C implementation is shown in code 2, where we assume the single precision floats are stored in IEEE 32 bit floating point format [40].

The reason why we decide to implement the approximation in single precision using coefficient arrays of 16 entries is because each requires only 512 bits (16×32 bits) to store all the possible values for a given monomial's coefficient. The significance of 512 bits cannot be overstated, as it is the width of an AVX-512 and A64FX vector register. Thus, instead of querying the cache to retrieve the coefficients, they can be held in vector registers, bypassing the cache entirely, and achieving extremely fast speeds. Recognising the coefficients can be stored in a single coalesced vector register is currently largely beyond most compilers' capabilities using OpenMP directives and compiler flags alone. However, in the repository by Sheridan-Methven [66], specialised implementations using Intel vector intrinsics and Arm inline assembly code achieve this, obtaining the ultimate in performance.

```

729 typedef unsigned int uint32; // Assuming 32 bit floats.
730 typedef float float32;      // Assuming 32 bit integers.
731
732 #define TABLE_SIZE 16
733 #define N_MANTISSA_32 23      // For IEEE 754
734 #define FLOAT32_EXPONENT_BIAS 127 // For IEEE 754
735 #define FLOAT32_EXPONENT_BIAS_TABLE_OFFSET (FLOAT32_EXPONENT_BIAS - 1)
736 #define TABLE_MAX_INDEX (TABLE_SIZE - 1) // Zero indexing.
737 #define FLOAT32_AS_UINT32(x) (...) // Type punning.
738
739 const float32 poly_coef_0[TABLE_SIZE] = {0.0, -1.3, -1.6, ..., -4.0, -4.1, -4.5};
740 const float32 poly_coef_1[TABLE_SIZE] = {0.0, 2.6, 3.7, ..., 2800.0, 5300.0, 21000.0};
741
742 #pragma omp declare simd
743 static inline float32 polynomial_approximation(const float32 u, const uint32 index) {
744     return poly_coef_0[index] + poly_coef_1[index] * u;
745 }
746
747 #pragma omp declare simd
748 static inline uint32 get_table_index_from_float_format(const float32 u) {
749     uint32 index = FLOAT32_AS_UINT32(u) >> N_MANTISSA_32; // Remove the mantissa.
750     index = FLOAT32_EXPONENT_BIAS_TABLE_OFFSET - index; // Get the index.
751     return index > TABLE_MAX_INDEX ? TABLE_MAX_INDEX : index; // Avoid overflow.
752 }
753
754 void piecewise_polynomial_approximation(const unsigned int n_samples,
755                                         const float32 * restrict input,
756                                         float32 * restrict output) {
757     #pragma omp simd
758     for (unsigned int i = 0; i < n_samples; i++) {
759         float32 u = input[i];
760         bool predicate = u < 0.5;
761         u = predicate ? u : 1.0 - u;
762         uint32 index = get_table_index_from_float_format(u);
763         float32 z = polynomial_approximation(u, index);
764         z = predicate ? z : -z;
765         output[i] = z;
766     }
767 }

```

Code 2. C implementation of the piecewise linear approximation.

3.4 Performance of the implementations

Both the piecewise constant and linear implementations in codes 1 and 2 are non branching, vector capable, and use only basic arithmetic and bit wise operations, and thus we anticipate their performance should be exceptionally good. Indeed, their performance, along with several other implementations', is shown in table 2, with experiments performed on an Intel Skylake Xeon Gold CPU and an Arm based Cavium ThunderX2 [66].

Looking at the results from table 2, of all of these, the piecewise linear implementation using Intel vector intrinsics achieves the very fastest speeds, closely approaching the maximum speed of just reading and writing. Unsurprisingly, the freely available implementations from Cephess and GSL are not competitive with the commercial offerings from Intel. Nonetheless, even in single precision, our approximations, on Intel hardware, consistently beat the performance achieved

Table 2. Performance of various approximations and implementations of the inverse Gaussian cumulative distribution function, with performance measured as clock cycles per random number.

Description	Implementation	Hardware	Compiler	Precision	Clock cycles
Cephes [53]	ndtri	Intel	icc	Double	60 ± 1
GNU GSL	gsl_cdf_ugaussian_Pinv	Intel	icc	Double	52 ± 10
ASA241 [17, 73]	r4_normal_01_cdf_inverse	Intel	icc	Single	47 ± 1
Giles [28]	—	Intel	icc	Single	46 ± 2
Intel (HA)	vdCdfNormInv	Intel	icc	Double	9 ± 0.5
Intel (LA)	vdCdfNormInv	Intel	icc	Double	7 ± 0.5
Intel (HA)	vsCdfNormInv	Intel	icc	Single	3.4 ± 0.1
Intel (LA)	vsCdfNormInv	Intel	icc	Single	2.6 ± 0.1
Piecewise constant	OpenMP	Arm	armclang	Double	4.0 ± 0.5
Piecewise constant	OpenMP	Intel	icc	Double	1.5 ± 0.3
Piecewise cubic	OpenMP	Intel	icc	Single	0.9 ± 0.1
Piecewise cubic	Intrinsics	Intel	icc	Single	0.7 ± 0.1
Piecewise linear	Intrinsics	Intel	icc	Single	0.5 ± 0.1
Read and write	—	Intel	icc	Single	0.4 ± 0.1

from the Intel high accuracy (HA) or low accuracy (LA) offerings. Comparing the high accuracy Intel offering and the piecewise linear implementation, there stands to be a speed up by a factor of seven by switching to our approximation. These results vindicate our efforts, and that our simple approximations offer considerable speed improvements. It is also needless to say, that compared to the freely available open source offerings, the savings become even more significant.

4 MULTILEVEL MONTE CARLO

One of the core use cases for our high speed approximate random variables is in Monte Carlo applications. Frequently, Monte Carlo is used to estimate expectations of the form $\mathbb{E}(P)$ of functionals P which act on solutions X of stochastic differential equations of the form $dX_t = a(t, X_t) dt + b(t, X_t) dW_t$ for given drift and diffusion processes a and b . The underlying stochastic process X_t is itself usually approximated by some \hat{X} using a numerical method, such as the Euler-Maruyama or Milstein schemes [7, 45, 46]. These approximation schemes simulate the stochastic process from time $t = 0$ to $t = T$ over N time steps of size $\Delta t = \delta = \frac{T}{N}$, where the incremental update at the n -th iteration requires a Gaussian random variable Z_n to simulate the increment to the underlying Wiener process W_t , where $\Delta W_n = \sqrt{\delta} Z_n$. Such types of Monte Carlo simulations are widespread, with the most famous application being to price financial options.

Approximate random variables come into this picture by substituting the exact random variable samples in the numerical scheme with approximate ones. This facilitates running faster simulations, at the detriment of introducing error. However, using the multilevel Monte Carlo method [27], this error can be compensated for with negligible cost. Thus, the speed improvements offered by switching to approximate random variables can be largely recovered, and the original accuracy can be maintained. A detailed inspection of the error introduced from incorporating approximate Gaussian random variables into the Euler-Maruyama scheme and the associated multilevel Monte Carlo analysis is presented by Giles and Sheridan-Methven [31, 67]. As such, we will only briefly review the key points of the setup, and focus on detailing the resultant computational savings that can be expected from using approximate random variables.

For the Euler-Maruyama scheme, the unmodified version using exact Gaussian random variables Z produces an approximation \hat{X} , whereas the modified scheme using approximate random variables \tilde{Z} produces an approximation \tilde{X} ,

where the two schemes are respectively

$$\widehat{X}_{n+1} = \widehat{X}_n + a(t_n, \widehat{X}_n)\delta + b(t_n, \widehat{X}_n)\sqrt{\delta}Z_n \quad \text{and} \quad \widetilde{X}_{n+1} = \widetilde{X}_n + a(t_n, \widetilde{X}_n)\delta + b(t_n, \widetilde{X}_n)\sqrt{\delta}\widetilde{Z}_n,$$

where $t_n := n\delta$.

The regular multilevel Monte Carlo construction varies the discretisation between two levels, producing *fine* and *coarse* simulations. The functional P would act on each path simulation, producing the fine and coarse approximations P^f and P^c respectively. If the path simulation uses the exact random variables we denote these as \widehat{P}^f and \widehat{P}^c , and alternatively if it uses approximate random variables as \widetilde{P}^f and \widetilde{P}^c . In general there may be multiple tiers of fine and coarse levels, so we use l to index these, where increasing values of l correspond to finer path simulations. Thus for a given l we have $\widehat{P}_l \equiv \widehat{P}^f$ and $\widehat{P}_{l-1} \equiv \widehat{P}^c$, and similarly $\widetilde{P}_l \equiv \widetilde{P}^f$ and $\widetilde{P}_{l-1} \equiv \widetilde{P}^c$. If we have levels $l \in \{0, 1, 2, \dots, L\}$ and use the convention $\widehat{P}_{-1} := \widetilde{P}_{-1} := 0$, then Giles and Sheridan-Methven [31] suggest the nested multilevel Monte Carlo framework

$$E(P) \approx E(\widehat{P}_L) = \sum_{l=0}^L E(\widehat{P}_l - \widehat{P}_{l-1}) = \sum_{l=0}^L E(\widetilde{P}_l - \widetilde{P}_{l-1}) + E(\widehat{P}_l - \widehat{P}_{l-1} - \widetilde{P}_l + \widetilde{P}_{l-1}),$$

where the first approximation is the regular Monte Carlo procedure [33], the first equality is the usual multilevel Monte Carlo decomposition [27], and the final equality is the nested multilevel Monte Carlo framework [31, 67]. Giles and Sheridan-Methven [31, 67] show that the two way differences in the regular and nested multilevel Monte Carlo settings behave almost identically, and the main result of their analysis is determining the behaviour of the final four way difference's variance [31, lemmas 4.10 and 4.11] [67, corollaries 6.2.6.2 and 6.2.6.3]. They find that for Lipschitz continuous and differentiable functionals

$$\|\widehat{P}^f - \widehat{P}^c - \widetilde{P}^f + \widetilde{P}^c\|_p \leq O(\delta^{1/2}\|Z - \widetilde{Z}\|_{p'})$$

for any p' such that $2 \leq p < p' < \infty$, and that for Lipschitz continuous but non-differentiable functionals that

$$\|\widehat{P}^f - \widehat{P}^c - \widetilde{P}^f + \widetilde{P}^c\|_p \leq O(\min\{\delta^{1/2}\|Z - \widetilde{Z}\|_{p'}^{p'(1-\epsilon)/(p'+1)}, \delta^{(1-\epsilon)/2p-1/2p'}\|Z - \widetilde{Z}\|_{p'}\})$$

for any $\epsilon > 0$. We can see that in all circumstances covered by their analysis that there is a dependence on the approximation error $\|Z - \widetilde{Z}\|_{p'}$ for some $L^{p'}$ norm where $p' > 2$.

4.1 Expected time savings

The regular multilevel estimator $\hat{\theta}$ is

$$\hat{\theta} := \sum_{l=0}^L \frac{1}{\widehat{m}_l} \sum_{i=1}^{\widehat{m}_l} (\widehat{P}_l^{(i,l)} - \widehat{P}_{l-1}^{(i,l)}),$$

and the nested multilevel estimator $\tilde{\theta}$ is

$$\tilde{\theta} := \sum_{l=0}^L \frac{1}{\widetilde{m}_l} \sum_{i=1}^{\widetilde{m}_l} (\widetilde{P}_l^{(i,l)} - \widetilde{P}_{l-1}^{(i,l)}) + \frac{1}{\widetilde{M}_l} \sum_{i=1}^{\widetilde{M}_l} (\widehat{P}_l^{(i,l+L+1)} - \widehat{P}_{l-1}^{(i,l+L+1)} - \widetilde{P}_l^{(i,l+L+1)} + \widetilde{P}_{l-1}^{(i,l+L+1)}),$$

where \widehat{m}_l , \widetilde{m}_l , and \widetilde{M}_l are the number of paths generated, each with a computational time cost of \hat{c}_l , \tilde{c}_l , and \widetilde{C}_l , and variance \hat{v}_l , \tilde{v}_l , and \widetilde{V}_l respectively, and all terms with the same suffix (i, l) are computed using the same uniform random numbers. Each of these estimators will have an error due to the finite number of paths used and the approximation scheme employed. The total error arising from these two factors is commonly referred to as the variance bias trade-off, where the mean squared error (MSE) of an estimator $\theta \in \{\hat{\theta}, \tilde{\theta}\}$ is given by $\text{MSE}(\theta) = \mathbb{V}(\theta) + \text{Bias}^2(\theta)$ [33, page 16].

Setting the desired MSE to ε^2 and choosing the maximum simulation fidelity such that we just satisfy $\text{Bias}^2(\theta) \leq \frac{\varepsilon^2}{2}$, then we can derive an expression for the total computational time T . Forcing T to be minimal is achieved by performing a constrained minimisation with an objective function \mathcal{F} . Considering the estimator $\hat{\theta}$, the corresponding objective function is $\hat{\mathcal{F}} := \sum_{l=0}^L \hat{m}_l \hat{c}_l + \mu (\sum_{l=0}^L \frac{\hat{\theta}_l}{\hat{m}_l} - \frac{\varepsilon^2}{2})$, where μ is a Lagrange multiplier enforcing the constraint $\mathbb{V}(\hat{\theta}) = \frac{\varepsilon^2}{2}$. Similarly the estimator $\tilde{\theta}$ would have the objective function $\tilde{\mathcal{F}} := \sum_{l=0}^L \tilde{m}_l \tilde{c}_l + \tilde{M}_l \tilde{C}_l + \lambda (\sum_{l=0}^L \frac{\tilde{\theta}_l}{\tilde{m}_l} + \frac{\tilde{V}_l}{\tilde{M}_l} - \frac{\varepsilon^2}{2})$, with λ being another Lagrange multiplier. Treating the number of paths as continuous variables, these are readily minimised to give

$$\hat{T} = 2\varepsilon^{-2} \left(\sum_{l=0}^L \sqrt{\hat{\theta}_l \hat{c}_l} \right)^2 \quad \text{and} \quad \tilde{T} = 2\varepsilon^{-2} \left(\sum_{l=0}^L \sqrt{\tilde{\theta}_l \tilde{c}_l} + \sqrt{\tilde{V}_l \tilde{C}_l} \right)^2,$$

where the minimal number of paths required are given by

$$\hat{m}_l = \varepsilon^{-1} \sqrt{\frac{2\hat{T}\hat{\theta}_l}{\hat{c}_l}}, \quad \tilde{m}_l = \varepsilon^{-1} \sqrt{\frac{2\tilde{T}\tilde{\theta}_l}{\tilde{c}_l}}, \quad \text{and} \quad \tilde{M}_l = \varepsilon^{-1} \sqrt{\frac{2\tilde{T}\tilde{V}_l}{\tilde{C}_l}},$$

and hence an overall saving of

$$\tilde{T} \approx 2\varepsilon^{-2} \left(\sum_{l=0}^L \sqrt{\tilde{\theta}_l \tilde{c}_l} \left(\sqrt{\frac{\tilde{\theta}_l \tilde{c}_l}{\hat{\theta}_l \hat{c}_l}} + \sqrt{\frac{\tilde{V}_l \tilde{C}_l}{\hat{\theta}_l \hat{c}_l}} \right) \right)^2 \leq \hat{T} \max_{l \leq L} \left\{ \frac{\tilde{\theta}_l \tilde{c}_l}{\hat{\theta}_l \hat{c}_l} \left(1 + \sqrt{\frac{\tilde{V}_l \tilde{C}_l}{\hat{\theta}_l \hat{c}_l}} \right)^2 \right\}. \quad (1)$$

For modest fidelity approximations where $\tilde{\theta}_l \approx \hat{\theta}_l$, the term $\frac{\tilde{\theta}_l \tilde{c}_l}{\hat{\theta}_l \hat{c}_l} \approx \frac{\tilde{c}_l}{\hat{c}_l}$ in (1) measures the potential time savings, and the term $(1 + (\tilde{V}_l \tilde{C}_l / \hat{\theta}_l \hat{c}_l)^{1/2})^2$ assesses the efficiency of realising these savings. A balance needs to be achieved, and the approximations should be sufficiently fast so there is the potential for large savings, but of a sufficient fidelity so the variance of the expensive four way difference is considerably lower than the variance of the cheaper two way difference.

Although we have estimates for the costs of each estimator from table 2, the variance will depend on the stochastic process being simulated, the numerical method being used, and the approximation employed. To make our estimates more concrete and quantify the possible variance reductions, we consider a geometric Brownian motion where $a(t, X_t) := \mu X_t$ and $b(t, X_t) := \sigma X_t$ for two strictly positive constants μ and σ , where we take $\mu = 0.05$, $\sigma = 0.2$, $X_0 = 1$, and $T = 1$ [27, 6.1]. The coarse level's Weiner increments are formed by pair wise summing the fine level's (which uses the fine time increment δ^f). Importantly, the uniform random variable samples producing the exact Gaussian random variables will be the same as those used for producing the approximate Gaussian random variables, ensuring a tight coupling. Thus, for some approximation $\tilde{\Phi}^{-1} \approx \Phi^{-1}$, such as those from section 2, for a single uniform sample $U_n \sim \mathcal{U}(0, 1)$, then $Z_n := \Phi^{-1}(U_n)$ and $\tilde{Z}_n := \tilde{\Phi}^{-1}(U_n)$, where the U_n is the same for both of these.

The variance for the various multilevel terms for different time increments for the Euler-Maruyama and Milstein schemes are shown in figure 3 for various approximations. We consider the underlying process itself, corresponding to the functional $P(X) = X$. The piecewise constant approximation uses 1024 intervals, which from figure 1(b) has an RMSE of approximately 10^{-2} . The piecewise linear approximation uses 15 dyadic intervals (16 coefficients stored to correctly handle $\frac{1}{2}$), which from figure 2(b) has a similar RMSE of approximately 10^{-2} . The piecewise cubic approximation uses the same number of intervals, which from figure 2(b) has an RMSE of approximately 10^{-4} . We also include the approximation using Rademacher random variables for comparison.

We can see from figure 3 that the two way differences exhibit the usual strong convergence order $\frac{1}{2}$ and 1 of the Euler-Maruyama and Milstein schemes, as expected [45]. Furthermore, as the functional is differentiable and Lipschitz continuous, this strong convergence rate is preserved for the four way differences [31, 67]. Giles and Sheridan-Methven

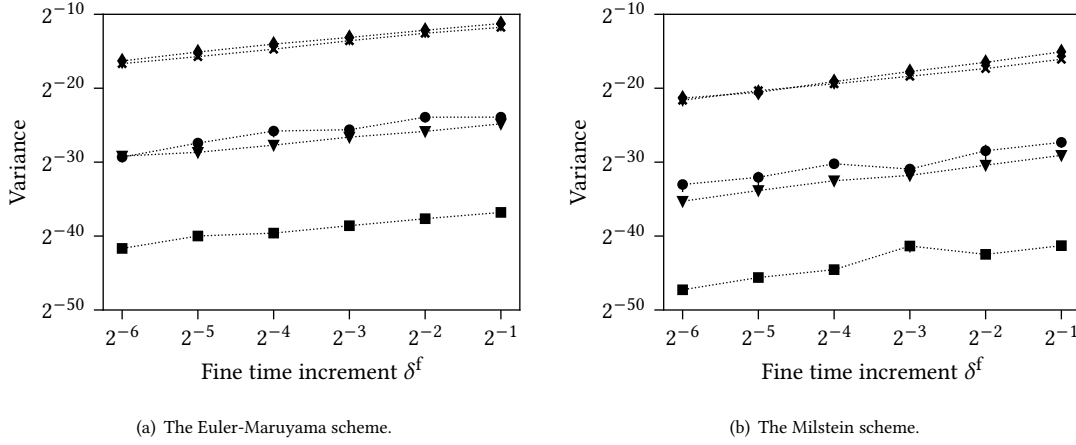


Fig. 3. The variances from using the Euler-Maruyama and Milstein schemes for a geometric Brownian motion for the functional $P(X) = X$. The (\diamond) -marker is the two way difference $\hat{P}^f - \hat{P}^c$, and the remaining markers $(\times, \bullet, \nabla, \blacksquare)$ are the four way difference $\hat{P}^f - \hat{P}^c - \hat{P}^f + \hat{P}^c$ for various approximations. (\times) Rademacher. (\bullet) Piecewise constant. (∇) Piecewise linear. (\blacksquare) Piecewise cubic.

[31, 67] derive this result for the Euler-Maruyama scheme, but the analysis for the Milstein scheme remains an open problem. Aside from this, the key point to note from figure 3 is the substantial drop in the variance between the two way and four way variances using our approximations. This drop in the variance is large for the piecewise constant and linear approximations, and huge for the piecewise cubic approximation. The approximation using Rademacher random variables has only a very small drop in variance.

Estimating the best cost savings seen compared to Intel (HA) from table 2, the variance reductions from figure 3(a), and using the simplifications $\tilde{C}_l \approx \hat{c}_l + \tilde{c}_l$ and $\tilde{v}_l \approx \hat{v}_l$, then the estimated speed ups and their efficiencies are shown in table 3. For the Rademacher random variables we use the optimistic approximations that this achieves the maximum speed set by reading and writing and that it offers a variance reduction of 2^{-1} . We can see from this that the piecewise linear approximation would give the largest savings, although the savings from the piecewise constant and cubic approximations are quite similar. Notice that while the piecewise cubic is achieving near perfect efficiency, its cost savings are not substantial enough to beat the marginally less efficient piecewise linear approximation which offers the larger potential savings. For all of our approximations the four way difference simulations are required very infrequently. Lastly, while the Rademacher random variables may offer the best time savings, the absence of any considerable variance reduction caused by the extremely low fidelity approximation results in a very inefficient multilevel scheme, demonstrating the balance required between speed and fidelity.

These cost savings are idealised in the respect that we have attributed the cost entirely to the generation of the random numbers. While this is quite a common assumption, the validity of this assumption will diminish the faster the approximations become, as the basic cost of the other arithmetic operations becomes significant. Thus, while a practitioner should have their ambitions set to achieve these savings, they should set more modest expectations.

Table 3. The cost savings, variance reductions, and possible speed ups from (1) by approximating the Gaussian distribution.

Approximation	$\log_2 \left(\frac{\tilde{V}_I}{\tilde{\sigma}_I} \right)$	$\frac{\hat{c}_I}{\tilde{c}_I}$	Speed up $\lim_{\tilde{\sigma}_I \rightarrow \hat{\sigma}_I} \left(\frac{\hat{T}}{\tilde{T}} \right)$	Efficiency $\left(1 + \sqrt{\tilde{V}_I \tilde{C}_I / \tilde{\sigma}_I \tilde{c}_I} \right)^{-2}$	$\frac{\tilde{m}_I}{\tilde{M}_I}$	$\frac{\tilde{m}_I}{\tilde{M}_I}$
Rademacher	-1	9	0.86	(9.5%)	3.24	1.4
Piecewise constant	-13	6	5.66	(94.4%)	1.03	240
Piecewise linear	-14	7	6.70	(95.7%)	1.02	360
Piecewise cubic	-25	5	5.00	(99.9%)	1.00	14000

5 THE NON CENTRAL χ^2 DISTRIBUTION

A second distribution of considerable practical and theoretical interest is the non-central χ^2 distribution, which regularly arises from the Cox-Ingersoll-Ross (CIR) interest rate model [20] (and the Heston model [37]). The distribution is parametrised as $\chi_\nu^2(\lambda)$, where $\nu > 0$ denotes the degrees of freedom and $\lambda \geq 0$ the non-centrality parameter, where we denote the inverse cumulative distribution function as $C_\nu^{-1}(\cdot; \lambda)$. Having a parametrised distribution naturally increases the complexity of any implementation, exact or approximate, and makes the distribution considerably more expensive to compute.

To gauge the procedure used, the associated function in Python's SciPy package (`ncx2.ppf`) calls the Fortran routine `CDFCHN` from the `CDFLIB` library by Brown et al. [14] (C implementations available [16]). This computes the value by root finding [18, algorithm R] on the offset cumulative distribution function $C_\nu(\cdot; \lambda)$, where $C_\nu(\cdot; \lambda)$ is itself computed by a complicated series expansion [2, (26.4.25)] involving the cumulative distribution function for the central χ^2 distribution. Overall, there are many stages involved, and as remarked by Burkardt [16, `cdflib.c`]: “*Very large values of $[\lambda]$ can consume immense computer resources*”. The analogous function `ncx2inv` in MATLAB, from its statistics and machine learning toolbox, appears to follow a similar approach based on its description [52, page 4301]. To indicate the costs, on an Intel core i7-4870HQ CPU the ratios between sampling from the non-central χ^2 distribution and the Gaussian distribution (`norm.ppf` and `norminv` in Python and MATLAB respectively) are shown in table 4, from which it is clear that the non-central χ^2 distribution can be vastly more expensive than the Gaussian distribution.

5.1 Approximating the non-central χ^2 distribution

There has been considerable research effort into quickly sampling from the non-central χ^2 distribution by various approximating distributions [1, 10, 39, 43, 61, 62, 64, 74], with the most notable being an approximation using Gaussian random variables by Abdel-Aty [1], (and a similar type of approximation by Sankaran [62]). With the exception of Abdel-Aty [1] and Sankaran [62], the remaining methods do not directly approximate the inverse cumulative distribution function, so do not give rise to an obvious coupling mechanism for multilevel Monte Carlo simulations, whereas our framework does. The schemes by Abdel-Aty [1] and Sankaran [62], both of which require evaluating Φ^{-1} , are appropriate for comparison to our scheme. Hence, our approximation scheme, while modest in its sophistication, appears novel research in this direction.

For approximating the non-central χ^2 distribution, we simplify our considerations by taking ν to be fixed, and thus the distribution then only has the parameter λ varying. While this may appear a gross simplification, in the financial applications ν is a model constant independent of the numerical method's parameters, and thus this simplification is appropriate.


We define the function $P_x(\cdot; y)$ for $x > 0$ and $0 < y < 1$ as

$$P_x(U; y) := \sqrt{\frac{x}{4y}} \left(\frac{y}{x} C_x^{-1} \left(U; \frac{(1-y)x}{y} \right) - 1 \right) \quad \text{so that} \quad C_v^{-1}(U; \lambda) \equiv \lambda + v + 2\sqrt{\lambda + v} P_v \left(U; \frac{v}{\lambda + v} \right).$$

This is because $P_x(\cdot; y)$ is better scaled than $C_v^{-1}(\cdot; \lambda)$ for the range of possible parameters, with the limits

$$P_x(U; y) \xrightarrow{y \rightarrow 0} \Phi^{-1}(U) \quad \text{and} \quad P_x(U; y) \xrightarrow{y \rightarrow 1} \frac{C_x^{-1}(U)}{2\sqrt{x}} - \frac{\sqrt{x}}{2},$$

where $C_v^{-1}(\cdot)$, without the parameter λ , is the inverse cumulative distribution function of the central χ^2 distribution. Thus to construct our approximation $\tilde{C}_v^{-1}(\cdot; \lambda) \approx C_v^{-1}(\cdot; \lambda)$, we first construct the piecewise polynomial approximation $\tilde{P}_x(\cdot; y) \approx P_x(\cdot; y)$, and then define $\tilde{C}_v^{-1}(\cdot; \lambda) := \lambda + v + 2\sqrt{\lambda + v} \tilde{P}_v(\cdot; \frac{v}{\lambda + v})$. An example piecewise linear approximation of $\tilde{C}_v^{-1}(\cdot; \lambda)$ using 8 intervals is shown in figure 4(a) for various values of the non-centrality parameter λ with the degrees of freedom fixed at $v = 1$. We can see from figure 4(a) that the fidelity of the approximation appears quite high across a range of parameter values.



(a) Piecewise linear approximations to the non-central χ^2 distribution with $v = 1$ and $\lambda \in \{1, 10, 20\}$ using 8 intervals and 16 interpolation values.

(b) The variance of the underlying CIR process and various two way differences (using the same time increment): (●) X_T , (▼) $X_T - \tilde{X}_T^{\text{EM}}$, (■) $X_T - \tilde{X}_T^{\text{linear}}$, and (◆) $X_T - \tilde{X}_T^{\text{cubic}}$.

Fig. 4. The non-central χ^2 distribution from the CIR process, and the variance of path simulations and possible multilevel corrections.

There are only two difficulties with constructing such an approximation. The first is that the distribution is no longer anti-symmetric about $u = 1/2$, which is easily remedied by constructing two approximations: one for $(0, \frac{1}{2}]$ and a second for $(\frac{1}{2}, 1)$. The second difficulty is the parametrisation. Noting that $y \in (0, 1)$ is dependent on the non-centrality parameter, we construct several approximations for various values of y (knowing the limiting cases) and linearly interpolate. A good choice of knot points are equally spaced values of \sqrt{y} , where figure 4(a) uses 16 interpolation values.

The performance of C and C++ implementations are shown in table 5 on Intel Skylake hardware compiled with g++. Neither Intel, NAG, nor GSL offer an exact implementation of the inverse cumulative distribution function for the non-central χ^2 distribution. Consequently, for our C implementation we compare against the `cdfchn` function from CDFLIB [14, 16], and for our C++ implementation the quantile function from the Boost library [3] acting on a `non_central_chi_squared` object (`boost/math/distributions`). Like the other implementations, Boost computes the

inverse by a numerical inversion of the cumulative distribution function. We can see from table 5 that our approximation is orders of magnitude faster than the exact functions in both languages.

The fidelity of our non central χ^2 approximations $\tilde{C}_v^{-1}(\cdot; \lambda)$ across the range of parameter values are quantified by the errors shown in table 6. A piecewise linear approximation achieves a consistently good RMSE, as seen in table 6(a). Similarly, the piecewise cubic approximation demonstrates the same behaviour in table 6(b), and has an RMSE which is typically an order of magnitude less than the piecewise linear approximation's.

The approximations from Abdel-Aty [1] and Sankaran [62] are also shown in tables 6(c) and 6(d) respectively. We can see from table 6(c) that the approximation from Abdel-Aty, while improving for increasing v , performs relatively poorly for large λ . Conversely, the approximation by Sankaran in table 6(d) is notably better than that by Abdel-Aty and often comparable to our piecewise cubic approximation, although demonstrating a much less consistent performance across the range of parameter values explored.

Having compared the error coming from the approximation by Sankaran [62], we can inspect the computational complexity of the approximation. The approximation can be expressed as

$$\tilde{C}_v^{-1}(U; \lambda) := (k + \lambda) \left(1 + c_1 c_2 \left(c_1 - 1 + \frac{1}{2} (2 - c_1) c_2 c_3 \right) + c_1 \sqrt{2c_3} \left(1 + \frac{1}{2} c_2 c_3 \right) \Phi^{-1}(U) \right)^{1/c_1},$$

where $c_1 := 1 - \frac{2}{3}(k + \lambda)(k + 3\lambda)(k + 2\lambda)^{-2}$, $c_2 := (k + 2\lambda)(k + \lambda)^{-2}$, and $c_3 := (c_1 - 1)(1 - 3c_1)$. Just as in our polynomial approximations, it contains an identical mix of basic arithmetic operations, but a Φ^{-1} evaluation. This can either be calculated exactly, or for improved speed it can be further approximated itself, but this then compounds a second level of approximation on top of the original. Furthermore, the expression contains an exponentiation by a factor of $1/c_1$, which will not in general be an integer. It can be expected that such a non-integer exponentiation operation will be quite expensive. Consequently, we anticipate this would be measurably slower than our polynomial approximations.

Contrasting the approximations between Sankaran [62] and our piecewise cubic, the piecewise cubic demonstrates a more consistent RMSE alongside a simpler computational composition, makes our piecewise polynomial the more attractive of the two.

5.2 Simulating the Cox-Ingersoll-Ross process

We mentioned the non-central χ^2 distribution arises from the Cox-Ingersoll-Ross (CIR) process: $dX_t = \kappa(\theta - X_t) dt + \sigma\sqrt{X_t} dW_t$ for strictly positive parameters κ , θ and σ . The distribution of X_T is a scaled non-central χ^2 distribution with $v = \frac{4\kappa\theta}{\sigma^2}$ and $\lambda = \frac{4\kappa X_0 \exp(-\kappa T)}{\sigma^2(1 - \exp(-\kappa T))}$ [20] [55, pages 67–68]. Simulating this with the Euler-Maruyama scheme forcibly approximates the non-central χ^2 distribution as a Gaussian distribution, and can require very finely resolved path simulations to bring the bias down to an acceptable level, as explored by Broadie and Kaya [13]. Similarly the Euler-Maruyama scheme is numerically ill posed due to the $\sqrt{X_t}$ term, and several adaptations exist to handle this appropriately [4–6, 9, 21–23, 35, 38, 47], especially when the Feller condition $2\kappa\theta \geq \sigma^2$ is not satisfied [24, 34]. The Feller condition is significant because [20, page 391]: “[X_t] can reach zero if $\sigma^2 > 2\kappa\theta$. If $2\kappa\theta \geq \sigma^2$, the upward drift is sufficiently large to make the origin inaccessible [see Feller [24]]”.

Rather than approximating the non-central χ^2 distribution with an exact Gaussian distribution when using the Euler-Maruyama scheme, we propose using approximate non-central χ^2 random variables, such as those from a piecewise linear approximation. This has the benefit of offering vast time savings, whilst introducing far less bias than the Euler-Maruyama scheme. The piecewise linear approximation is on average hundreds of times faster than the exact function, as was seen in table 4, giving vast savings, while still achieving a high fidelity, as was seen from table 6.

We can generate path simulations of the CIR process using the exact non-central χ^2 distribution, piecewise linear and cubic approximations, and the truncated Euler-Maruyama scheme from Higham et al. [38] (where $\sqrt{X_n} \rightarrow \sqrt{|X_n|}$), where we set $\kappa = \frac{1}{2}$ and $\theta = \sigma = T = X_0 = 1$ (just satisfying the Feller condition¹). These respectively give rise to the exact underlying process X_T and the three approximations $\tilde{X}_T^{\text{linear}}$, $\tilde{X}_T^{\text{cubic}}$, and \tilde{X}_T^{EM} . The variances of the underlying process and the differences between this and the approximations are shown in figure 4(b). Note that the difference terms are evaluated using the same time increments, and there is no mixing of coarse and fine paths in these two-way differences. The results of figure 4(b) do not change if we instead satisfy the Feller condition with strict inequality by setting e.g. $\kappa = \frac{3}{2}$.

From figure 4(b) we can see the exact underlying process' variance does not vary with the discretisation, which is to be expected as the simulation is from the exact distribution, and thus has no bias. We can see that the Euler-Maruyama approximation to the process exhibits a strong convergence order $\frac{1}{2}$, as expected [34, 38]. For the piecewise linear approximation, as the fidelity will not vary with the discretisation, then subsequently the variance also does not vary. Furthermore, given the high fidelity seen in table 6(a), the drop in the variance is approximately 2^{-15} . Similarly, the even higher fidelity cubic approximation demonstrates the same behaviour in table 6(b), but with a drop in variance of 2^{-25} . These findings demonstrate that the most substantial variance reductions come from using our piecewise polynomial approximations to the non central χ^2 distribution's inverse cumulative distribution function, rather than using the Gaussian distribution with the Euler-Maruyama scheme.

We can repeat these same variance comparisons with a different functional of the terminal value which is path dependent. A common example in financial simulations is the arithmetic average of the path, (an *Asian option* in mathematical finance [33]). Doing so produces an exactly identical variance structure as seen in figure 4(b). From this we conclude that our findings seen and explained thus far appear to carry over immediately to simulations requiring an entire path simulation, not only a terminal value, justifying the use of our approximations in time marching schemes.

Quantifying the savings that can be expected, we conservatively approximate the cost savings for both the linear and cubic approximations as a factor of 300 using table 5. Coupling this with the variance reductions seen in figure 4(b), the anticipated cost savings are shown in table 7. We can see from this that both offer impressive time savings, and the piecewise cubic again achieves a near perfect efficiency. However, the piecewise linear approximation's efficiency, while good, is losing an appreciable fraction, making the piecewise cubic the preferred choice. Ultimately, both offer vast time savings by factors of 250 or higher.

6 CONCLUSIONS

The expense of sampling random variables can be significant. In the work presented, we proposed, developed, analysed, and implemented approximate random variables as a means of circumventing this for a variety of modern hardware. By incorporating these into a nested multilevel Monte Carlo framework we showcased how the full speed improvements offered can be recovered with near perfect efficiency without losing any accuracy. With a detailed treatment, we showed that even for basic simulations of geometric Brownian motions requiring Gaussian random variables, speedups of a factor 5 or more can be expected. The same framework was also applied to the more difficult Cox-Ingersoll-Ross process and its non-central χ^2 distribution, offering the potential for vast speedups of a factor of 250 or more.

For sampling from a wide class of univariate distributions, we based our work on the inverse transform method. Unfortunately, this method is expensive as it relies on evaluating a distribution's inverse cumulative distribution function.

¹In this case $\nu = 2$ for which there is a finite and non zero probability density at zero.

We showed that many implementations for the Gaussian distribution, while accurate to near machine precision, are ill-suited in several respects for modern vectorised hardware. To address this, we introduced a generalised notion of approximate random variables, produced using approximations to a distribution’s inverse cumulative distribution function. The two major classes of approximations we introduced were: piecewise constant approximations using equally spaced intervals, and piecewise linear approximations using geometrically small intervals dense near the distribution’s tails. These cover a wide class of possible approximations, and notably recover as special cases: Rademacher random variables and the weak Euler-Maruyama scheme [33], moment matching schemes by Müller et al. [54], and truncated bit approximations by Giles et al. [30]. For piecewise constant and linear approximations, we analysed and bounded the errors for a range of possible norms. The significance of these bounds is that they are valid for arbitrarily high moments, which extends the results from Giles et al. [30], and crucially is necessary for the nested multilevel Monte Carlo analysis by Giles and Sheridan-Methven [31]. Lastly, these approximations were able to achieve very high fidelities, with the piecewise linear approximation using a geometric sequence of intervals providing high resolution of singularities.

With the approximations detailed from a mathematical perspective, we highlighted two possible implementations in C [66]. The benefit of these approximations is that they are by design ideally suited for modern vector hardware and achieve the highest possible computational speeds. They can be readily vectorised using OpenMP SIMD directives, have no conditional branching, avoid division and expensive function evaluations, and only require simple additions, multiplications, and bit manipulations. The piecewise constant and linear implementations were orders of magnitude faster than most freely available open source libraries, and typically a factor of 5–7 times faster than the proprietary Intel Maths Kernel Library, achieving close to the maximum speed of reading and writing to memory. This speed comes from the simplicity of their operations, and heavy capitalisation on the fast speeds of querying the cache and vector registers.

Incorporating approximate random variables into the nested multilevel Monte Carlo framework by Giles and Sheridan-Methven [31], the low errors and fast speeds of the approximations can be exploited to obtain their full speed benefits without losing accuracy. Inspecting the appropriate multilevel variance reductions, we demonstrated how practitioners can expect to obtain speed improvements of a factor of 5–7 by using approximate Gaussian random variables, where the fastest approximation was the piecewise linear approximation. This appears to be the case when using either the Euler-Maruyama or Milstein schemes.

The Cox-Ingersoll-Ross process’ [20] exact solution follows a non-central χ^2 distribution, and is an example of a parametrised distribution which is extremely expensive to sample from. We applied our approximation framework to this to produce a parametrised approximation. The error of using our approximate random variables was orders of magnitude lower than approximating paths using the Euler-Maruyama scheme, showing our approximate random variables are considerably more suitable for generating path simulations than the Euler-Maruyama scheme. This circumvents the problem of the Euler-Maruyama scheme having a very large bias for the Cox-Ingersoll-Ross process [13]. Furthermore, our implementation was substantially quicker than those by CDFLIB and Boost, offering a speed improvement of 250 times or higher.

All of the code to produce these figures, tables, and approximations is freely available and hosted in repositories by Sheridan-Methven [65, 66], along with an introductory guide aimed at practitioners [66].

7 ACKNOWLEDGEMENTS

We would like to acknowledge and thank those who have financially sponsored this work. This includes the Engineering and Physical Sciences Research Council (EPSRC) and Oxford University’s centre for doctoral training in Industrially

1249 Focused Mathematical Modelling (InFoMM), with the EP/L015803/1 funding grant. Furthermore, this research stems
 1250 from a PhD project [67] which was funded by Arm and NAG. Funding was also provided by the EPSRC ICONIC
 1251 programme grant EP/P020720/1, the Hong Kong Innovation and Technology Commission (InnoHK Project CIMDA),
 1252 and by Mansfield College, Oxford.

1255 REFERENCES

- 1257 [1] S.H. Abdel-Aty. 1954. Approximate formulae for the percentage points and the probability integral of the non-central χ^2 distribution. *Biometrika*
 1258 41, 3/4 (1954), 538–540.
- 1259 [2] Milton Abramowitz and Irene A. Stegun. 1948. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*. Vol. 55. US
 1260 government printing office. (6 printing, November 1967).
- 1261 [3] Nikhar Agrawal, Anton Bikineev, Paul A. Bristow, Marco Guazzone, Christopher Kormanyos, Hubert Holin, Bruno Lalande, John Maddock, Jeremy
 1262 Murphy, Matthew Pulver, Johan Råde, Gautam Sewani, Benjamin Sobotta, Nicholas Thompson, Thijs van den Berg, Daryle Walker, Xiaogang Zhang,
 1263 et al. 2020. The Boost C++ library. <https://www.boost.org/> Version 1.74.0.
- 1264 [4] Aurélien Alfonsi. 2005. On the discretization schemes for the CIR (and Bessel squared) processes. *Monte Carlo methods and applications* 11, 4 (2005),
 1265 355–384.
- 1266 [5] Aurélien Alfonsi. 2008. A second-order discretization scheme for the CIR process: application to the Heston model. *Preprint CERMICS hal-00143723*
 1267 14 (2008).
- 1268 [6] Aurélien Alfonsi. 2010. High order discretization schemes for the CIR process: application to affine term structure and Heston models. *Mathematics*
 1269 *of computation* 79, 269 (2010), 209–237.
- 1270 [7] Søren Asmussen and Peter W. Glynn. 2007. *Stochastic simulation: algorithms and analysis*. Vol. 57. Springer science & business media.
- 1271 [8] J.D. Beasley and S.G. Springer. 1977. Algorithm AS 111: the percentage points of the normal distribution. *Journal of the royal statistical society. Series*
 1272 *C (applied statistics)* 26, 1 (March 1977), 118–121.
- 1273 [9] Abdel Berkaoui, Mireille Bossy, and Awa Diop. 2008. Euler scheme for SDEs with non-Lipschitz diffusion coefficient: strong convergence. *ESAIM:*
 1274 *probability and statistics* 12 (2008), 1–11.
- 1275 [10] D.J. Best and D.E. Roberts. 1975. Algorithm AS 91: The percentage points of the χ^2 distribution. *Journal of the royal statistical society. Series C*
 1276 *(applied statistics)* 24, 3 (1975), 385–388.
- 1277 [11] Fischer Black and Myron Scholes. 1973. The pricing of options and corporate liabilities. *Journal of political economy* 81, 3 (1973), 637–654.
- 1278 [12] George E.P. Box and Mervin E. Muller. 1958. A note on the generation of random normal deviates. *The annals of mathematical statistics* 29 (1958),
 1279 610–611.
- 1280 [13] Mark Broadie and Özgür Kaya. 2006. Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations research* 54, 2
 1281 (2006), 217–231.
- 1282 [14] Barry W. Brown, James Lovato, and Kathy Russell. 1994. CDFLIB: library of Fortran routines for cumulative distribution functions, inverses, and
 1283 other parameters.
- 1284 [15] Christian Brugger, Christian de Schryver, Norbert Wehn, Steffen Omland, Mario Hefter, Klaus Ritter, Anton Kostiuik, and Ralf Korn. 2014. Mixed
 1285 precision multilevel Monte Carlo on hybrid computing systems. In *2104 IEEE conference on computational intelligence for financial engineering &*
 1286 *economics (CIFER)*. IEEE, 215–222.
- 1287 [16] John Burkardt. 2019. C source codes for CDFLIB. https://people.sc.fsu.edu/~jburkardt/c_src/cdflib/cdflib.html Accessed Wednesday 2 September
 1288 2020.
- 1289 [17] John Burkardt. 2020. MATLAB source codes for ASA241 and ASA111. https://people.sc.fsu.edu/~jburkardt/m_src/m_src.html Accessed Thursday
 1290 18 June 2020.
- 1291 [18] Jacques C.P. Bus and Theodorus Jozef Dekker. 1975. Two efficient algorithms with guaranteed convergence for finding a zero of a function. *ACM*
 1292 *transactions on mathematical software (TOMS)* 1, 4 (1975), 330–345.
- 1293 [19] Ray C. C. Cheung, Dong-U Lee, Wayne Luk, and John D. Villasenor. 2007. Hardware Generation of Arbitrary Random Number Distributions From
 1294 Uniform Distributions Via the Inversion Method. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 15, 8 (2007), 952–962.
- 1295 [20] John C. Cox, Jonathan E. Ingersoll Jr, and Stephen A. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53, 2 (March 1985),
 1296 385–408–164.
- 1297 [21] Andrei Cozma and Christoph Reisinger. 2018. Strong order 1/2 convergence of full truncation Euler approximations to the Cox–Ingersoll–Ross
 1298 process. *IMA journal of numerical analysis* 40, 1 (October 2018), 358–376.
- 1299 [22] Griselda Deelstra and Freddy Delbaen. 1998. Convergence of discretized stochastic (interest rate) processes with stochastic drift term. *Applied*
 1300 *stochastic models and data analysis* 14, 1 (1998), 77–84.
- [23] Steffen Dereich, Andreas Neuenkirch, and Lukasz Szpruch. 2012. An Euler-type method for the strong approximation of the Cox–Ingersoll–Ross
 process. *Proceedings of the royal society A: mathematical, physical and engineering sciences* 468, 2140 (2012), 1105–1115.
- [24] William Feller. 1951. Two singular diffusion problems. *Annals of mathematics* (1951), 173–182.

- [25] Agner Fog. 2018. Instruction tables: lists of instruction latencies, throughputs and micro-operation breakdowns for Intel, AMD and VIA CPUs. https://www.agner.org/optimize/instruction_tables.pdf Updated 9 April 2018, Copenhagen University college of engineering.
- [26] Mark Galassi, Jim Davies, James Theiler, Brian Gough, Gerard Jungman, Patrick Alken, Michael Booth, Fabrice Rossi, Rhys Ulerich, et al. 2017. GNU scientific library 2.4.
- [27] Michael B. Giles. 2008. Multilevel Monte Carlo path simulation. *Operations research* 56, 3 (2008), 607–617.
- [28] Michael B. Giles. 2011. Approximating the $\text{erf} \operatorname{inv}$ function. In *GPU computing gems, jade edition*. Vol. 2. Elsevier, 109–116.
- [29] Michael B. Giles, Mario Hefter, Lukas Mayer, and Klaus Ritter. 2019. Random bit multilevel algorithms for stochastic differential equations. *Journal of complexity* (2019).
- [30] Michael B. Giles, Mario Hefter, Lukas Mayer, and Klaus Ritter. 2019. Random bit quadrature and approximation of distributions on Hilbert spaces. *Foundations of computational mathematics* 19, 1 (2019), 205–238.
- [31] Michael B. Giles and Oliver Sheridan-Methven. 2022. Analysis of nested multilevel Monte Carlo using approximate normal random variables. *SIAM/ASA Journal on Uncertainty Quantification* 10, 1 (2022), 200–226. <https://doi.org/10.1137/21M1399385>
- [32] Michael B. Giles and Benjamin J. Waterhouse. 2009. Multilevel quasi-Monte Carlo path simulation. *Advanced financial modelling, Radon series on computational and applied mathematics* 8 (2009), 165–181.
- [33] Paul Glasserman. 2013. *Monte Carlo methods in financial engineering* (1 ed.). Stochastic modelling and applied probability, Vol. 53. Springer science & business media.
- [34] István Gyöngy. 1998. A note on Euler’s approximations. *Potential analysis* 8, 3 (1998), 205–216.
- [35] István Gyöngy and Miklós Rásonyi. 2011. A note on Euler approximations for SDEs with Hölder continuous diffusion coefficients. *Stochastic processes and their applications* 121, 10 (2011), 2189–2200.
- [36] Cecil Hastings Jr, Jeanne T. Wayward, and James P. Wong Jr. 1955. *Approximations for digital computers*. Princeton University press.
- [37] Steven L. Heston. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The review of financial studies* 6, 2 (1993), 327–343.
- [38] Desmond J. Higham, Xuerong Mao, and Andrew M. Stuart. 2002. Strong convergence of Euler-type methods for nonlinear stochastic differential equations. *SIAM journal on numerical analysis* 40, 3 (2002), 1041–1063.
- [39] David C. Hoaglin. 1977. Direct approximations for chi-squared percentage points. *Journal of the American statistical association* 72, 359 (1977), 508–515.
- [40] IEEE. 2008. IEEE standard for binary floating-point arithmetic. Computer society standards committee. Working group of the microprocessor standards subcommittee.
- [41] International organization for standardization (ISO). 2012. ISO/IEC 9899:2011: programming languages—C. *ISO working group 14* (2012).
- [42] Arieh Iserles. 1996. *A first course in the numerical analysis of differential equations*. Cambridge University press.
- [43] Norman L. Johnson, Samuel Kotz, and Narayanaswamy Balakrishnan. 1995. *Continuous univariate distributions*. John Wiley & sons, Ltd.
- [44] Corwin Joy, Phelim P. Boyle, and Ken Seng Tan. 1996. Quasi-Monte Carlo methods in numerical finance. *Management science* 42, 6 (1996), 926–938.
- [45] Peter E. Kloeden and Eckhard Platen. 1999. *Numerical solution of stochastic differential equations*. Stochastic modelling and applied probability, Vol. 23. Springer. Corrected 3 printing.
- [46] Gabriel J. Lord, Catherine E. Powell, and Tony Shardlow. 2014. *An introduction to computational stochastic PDEs*. Cambridge University press.
- [47] Roger Lord, Remmert Koekoek, and Dick van Dijk. 2010. A comparison of biased simulation schemes for stochastic volatility models. *Quantitative finance* 10, 2 (2010), 177–194.
- [48] Pierre L’Ecuyer. 2016. Randomized quasi-Monte Carlo: an introduction for practitioners. In *International conference on Monte Carlo and quasi-Monte Carlo methods in scientific computing*. Springer, 29–52.
- [49] George Marsaglia and T.A. Bray. 1964. A Convenient Method for Generating Normal Variables. *SIAM Rev.* 6, 3 (1964), 260–264.
- [50] George Marsaglia and Wai Wan Tsang. 2000. The Ziggurat method for generating random variables. *Journal of statistical software* 5, 1 (2000), 1–7.
- [51] George Marsaglia, Arif Zaman, and John C.W. Marsaglia. 1994. Rapid evaluation of the inverse of the normal distribution function. *Statistics & probability letters* 19, 4 (15 March 1994), 259–266.
- [52] Matlab MathWorks. 2018. Statistics and machine learning toolbox: user’s guide. R2018b.
- [53] Stephen L. Moshier. 1992. Cephes mathematical library. <http://www.netlib.org/cephes/> Accessed Tuesday 8 September 2020.
- [54] Eike H. Müller, Rob Scheichl, and Tony Shardlow. 2015. Improving multilevel Monte Carlo for stochastic differential equations with application to the Langevin equation. *Proceedings of the royal society of London A: mathematical, physical and engineering sciences* 471, 2176 (2015).
- [55] Claus Munk. 2011. *Fixed income modelling*. Oxford University press.
- [56] NAG. 2017. NAG library manual, mark 26.
- [57] Robert E. Odeh and J.O. Evans. 1974. Algorithm AS 70: the percentage points of the normal distribution. *Journal of the royal statistical society. Series C (applied statistics)* 23, 1 (1974), 96–97.
- [58] Steffen Omland, Mario Hefter, Klaus Ritter, Christian Brugger, Christian De Schryver, Norbert Wehn, and Anton Kostiuik. 2015. Exploiting mixed-precision arithmetics in a multilevel Monte Carlo approach on FPGAs. In *FPGA based accelerators for financial applications*. Springer, 191–220.
- [59] Francesco Petrogalli. 2016. A sneak peek into SVE and VLA programming. White paper.
- [60] Michael James David Powell. 1981. *Approximation theory and methods*. Cambridge University press.

- [61] Stephen O. Rice. 1968. Uniform asymptotic expansions for saddle point integrals—application to a probability distribution occurring in noise theory. *Bell system technical journal* 47, 9 (1968), 1971–2013.
- [62] Munuswamy Sankaran. 1959. On the non-central chi-square distribution. *Biometrika* 46, 1/2 (1959), 235–237.
- [63] Rüdiger Seydel. 2006. *Tools for computational finance*. Vol. 3. Springer.
- [64] B.L. Shea. 1991. Algorithm AS R85: a remark on AS 91: the percentage points of the χ^2 distribution. *Journal of the royal statistical society. Series C (applied statistics)* 40, 1 (1991), 233–235.
- [65] Oliver Sheridan-Methven. 2020. Approximating inverse cumulative distribution functions. https://github.com/oliversheridanmethven/approximating_inverse_cumulative_distribution_functions GitHub repository.
- [66] Oliver Sheridan-Methven. 2020. Getting started with approximate random variables: a brief guide for practitioners. https://github.com/oliversheridanmethven/approximate_random_variables GitHub repository.
- [67] Oliver Sheridan-Methven. 2021. Nested multilevel Monte Carlo methods and a modified Euler-Maruyama scheme utilising approximate Gaussian random variables suitable for vectorised hardware and low-precisions. DPhil thesis, supervisor Prof. Michael B. Giles, Mathematical Institute, University of Oxford.
- [68] Richard M. Stallman and the GCC developer community. 2020. Using the GNU compiler collection. Version 10.1.0.
- [69] Nigel Stephens, Stuart Biles, Matthias Boettcher, Jacob Eapen, Mbou Eyole, Giacomo Gabrielli, Matt Horsnell, Grigoris Magklis, Alejandro Martinez, Nathanael Premillieu, Alastair Reid, Alejandro Rico, and Paul Walker. 2017. The ARM scalable vector extension. *IEEE micro* 37, 2 (2017), 26–39.
- [70] Shu Tezuka. 1995. *Uniform random numbers: theory and practice*. Kluwer.
- [71] Ruud van der Pas, Eric Stotzer, and Christian Terboven. 2017. *Using OpenMP — the next step: affinity, accelerators, tasking, and SIMD*. MIT press.
- [72] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, the SciPy 1.0 contributors, et al. 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods* 17 (2020), 261–272.
- [73] Michael J. Wichura. 1988. Algorithm AS 241: the percentage points of the normal distribution. *Journal of the royal statistical society. Series C (applied statistics)* 37, 3 (April 1988), 477–484.
- [74] Edwin B. Wilson and Margaret M. Hilferty. 1931. The distribution of chi-square. *proceedings of the national academy of sciences of the United States of America* 17, 12 (1931), 684.
- [75] Markus Wittmann, Thomas Zeiser, Georg Hager, and Gerhard Wellein. 2015. Short note on costs of floating point operations on current x86-64 architectures: Denormals, overflow, underflow, and division by zero. *arXiv preprint arXiv:1506.03997* (2015).
- [76] Linlin Xu and Giray Ökten. 2015. High-performance financial simulation using randomized quasi-Monte Carlo methods. *Quantitative finance* 15, 8 (2015), 1425–1436.

Table 4. The computing ratio between the non-central χ^2 and Gaussian distributions.

(a) Python.						(b) MATLAB.					
λ	ν					λ	ν				
	1	5	10	50	100		1	5	10	50	100
1	37	36	40	54	73	1	168	214	259	456	294
5	40	46	48	62	85	5	651	782	840	1510	2046
10	54	56	63	69	97	10	935	1086	1050	1838	2496
50	101	103	103	144	143	50	3000	2969	2562	4118	5333
100	191	190	192	189	185	100	4929	3461	5039	6046	6299
200	243	246	240	233	221	200	9456	9603	10129	11524	12766
500	465	474	465	446	416	500	22691	22713	22702	23328	26273
1000	459	458	455	471	474	1000	45872	43968	43807	44563	46780

Table 5. The computing ratio between the exact function and the piecewise linear approximation for the non-central χ^2 distribution.

(a) C implementation (compared against CDFLIB).						(b) C++ implementation (compared against Boost).					
λ	ν					λ	ν				
	1	5	10	50	100		1	5	10	50	100
1	333	412	458	666	864	1	671	1643	1534	1734	2093
5	391	447	534	701	966	5	1884	1831	1733	2037	2344
10	600	668	724	801	992	10	1924	1937	1863	2490	2490
50	1411	1424	1231	1811	1811	50	2576	2565	2876	2945	2974
100	2271	2174	2164	2207	2029	100	3238	3265	3255	3299	3354
200	2539	2624	2791	2304	2113	200	4382	4384	4373	4356	4333
500	5020	4912	4860	4908	4886	500	5260	5294	5243	5249	5224
1000	4822	4859	4866	4791	4980	1000	6101	6022	6026	6147	6093

Table 6. The RMSE of approximations $\tilde{C}_\nu^{-1}(\cdot; \lambda)$ to the non-central χ^2 distribution.

(a) Piecewise linear.						(b) Piecewise cubic.					
λ	ν					λ	ν				
	1	5	10	50	100		1	5	10	50	100
1	0.036	0.036	0.041	0.070	0.095	1	0.004	0.005	0.006	0.007	0.006
5	0.045	0.047	0.050	0.076	0.100	5	0.004	0.004	0.005	0.010	0.015
10	0.054	0.056	0.059	0.081	0.104	10	0.006	0.005	0.005	0.009	0.014
50	0.098	0.099	0.101	0.116	0.133	50	0.006	0.007	0.005	0.010	0.011
100	0.134	0.135	0.136	0.148	0.161	100	0.013	0.008	0.009	0.011	0.014
200	0.186	0.187	0.188	0.196	0.207	200	0.009	0.012	0.011	0.012	0.015

(c) Abdel-Aty.						(d) Sankaran.					
λ	ν					λ	ν				
	1	5	10	50	100		1	5	10	50	100
1	0.153	0.041	0.023	0.009	0.006	1		0.031	0.021	0.009	0.006
5	0.329	0.155	0.083	0.013	0.007	5		0.058	0.031	0.010	0.007
10	0.385	0.243	0.156	0.026	0.011	10	0.121	0.057	0.034	0.012	0.007
50	0.451	0.403	0.353	0.157	0.079	50	0.030	0.025	0.020	0.012	0.009
100	0.461	0.435	0.405	0.250	0.157	100	0.015	0.014	0.012	0.009	0.008
200	0.466	0.453	0.436	0.334	0.251	200	0.008	0.007	0.007	0.006	0.006

Table 7. The cost savings, variance reductions, and possible speed ups from (1) by approximating the non-central χ^2 distribution.

Approximation	$\log_2 \left(\frac{\tilde{V}_l}{\tilde{v}_l} \right)$	$\frac{\hat{c}_l}{\tilde{c}_l}$	Speed up	Efficiency	$\frac{\tilde{m}_l}{\hat{m}_l}$	$\frac{\tilde{m}_l}{\tilde{M}_l}$
			$\lim_{\tilde{v}_l \rightarrow \hat{v}_l} \left(\frac{\hat{T}}{\tilde{T}} \right)$	$\left(1 + \sqrt{\tilde{V}_l \tilde{C}_l / \tilde{v}_l \tilde{c}_l} \right)^{-2}$		
Piecewise linear	-15	300	249	(83.3%)	1.09	2900
Piecewise cubic	-25	300	298	(99.4%)	1.00	100000