The Density of the Fractional Part of a Normal Distribution

Arif Zaman
Lahore University of Management Sciences

June 28, 2004

Keywords: Normal Distribution, Approximation, Fourier series
AMS Subject Classification: Primary 62E17, Secondary 42A10

Abstract

The distribution of the fractional part of a normal distribution with any mean and variance is often of interest when looking at cyclic phenomena. Some useful formulas, which allow for very accurate computations are developed here. A surprising fact is noted and explained, that for variances larger that 1, the distribution of the fractional part is remarkably uniform.

1 Introduction

A common way to provide randomization is to spin a large spinner rapidly, and then wait for it to stop by friction, and note the number that comes up under a fixed arrow. In many game shows and games of chance, it is allowed to make your prediction of the outcome after the spinner has been spun. In such cases, it appears that knowing the velocity and location of the spinner should allow a person to obtain some advantage.

We obtain a very rapidly converging series for the density of the final loacation of a ball, and show that to obtain any advantage the standard deviation of the final location of the ball needs to be much less than a revolution.

2 Definitions

We number the circumference of the spinner continuously with numbers from 0 to 1. Assume that the initial position is 0 and initial velocity is known. Using the initial velocity, and a large sample of observations, or some detailed mechanical model of the frictional forces, we can represent the total distance Z, that the spinner is going to spin before coming to rest as a normal random variable with some known mean μ and variance σ^2. The stopping point of the spinner will then be $Y = \frac{(}{Z})$ the fractional part of Z.

We define the normal density with mean $\mu = 0$ and variance σ^2, and its sum over a unit spaced lattice by the following two equations

$$\begin{aligned} \phi_\sigma(x) &= (2\pi\sigma^2)^{-1/2} \exp(-x^2/(2\sigma^2)) \\ f_\sigma(x) &= \sum_{k=-\infty}^{\infty} \phi_\sigma(x+k) \end{aligned} \tag{1}$$

The density of Y, the stopping point of the spinner, can then be written simply as $f_\sigma(x-\mu)$, when $0 \le x \le 1$.

3 A few surprises

With the above definitions, and with confidence in the magic of statistics to be able to pull out an advantage from noisy and partial observations, as a test case I set out to see what advantage we could obtain when $\sigma = 1$. The first surprise was the graph of the density function, $f_1(x)$, as plotted by Maple, which is shown in figure 1. The first surprise is how close it is to an exact uniform distribution. The next surprise is that this difference looks almost exactly like a cosine function.

Clearly we should be able to better if the standard deviation was much smaller. As can be seen from figure 2, for values of the standard deviation σ near 1, the most likely point has a density that is different from the least likely point by about 10^{-8}. Thus a crude measure of the stopping distance gives almost no information at all, and σ has to be close to 0.6 before even a 1% edge is obtained. Obviously, for very small σ, we know almost exactly where the spinner is going to stop, and hence there is large difference between the most and least probable point.

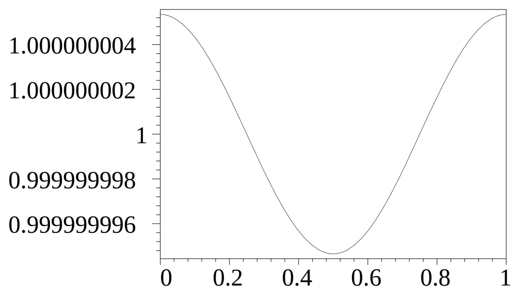

Figure 1. The density of the fractional part of a standard normal density $f_1(x)$ vs. x

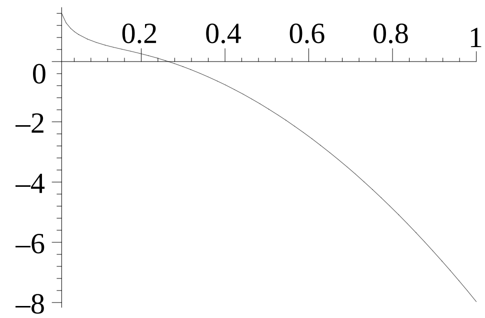

Figure 2. A graph of $\log_{10}(f_\sigma(0) - \log_{10}(f_\sigma(1/2)$ vs. σ. The log of difference between the most probable and the least probable points.

4 An explanation of the first surprise

One way to explain the fact that the sum of the normal density over a lattice of points is nearly 1 is to say that summing is approximate integration, and any density integrates to 1. This may be true but it begs the question of why we should get 8 digits of accuracy with such a crude sum. In fact, if we were to take the fractional part of an exponential density with variance 1, its, sum would just be markedly non uniform, with a density proportional to e^{-x} on the unit interval. For this density, the ratio between the most likely and the least likely densities would be e.

A better explanation relates to a well-known quick and dirty way of generating normal variates on a computer by simply summing 12 uniforms and subtracting 6. Since the variance of a uniform is $1/12$, by the central limit theorem, this procedure should have a density that is very close to a standard normal. But the distribution of the fractional part of the sum of any number of uniforms is *exactly* uniform. The conclusion can only be that the fractional part of a standard normal must be very close to uniform.

5 An explanation of the second surprise

The deviation from uniformity appeared to be almost perfectly a cosine function. To try and understand why this happens, let us try to take the Fourier transform of the density. Since the density of Y is only on the $(0, 1)$ interval, we can take the discrete Fourier transform to get the individual Fourier coefficients

$$\begin{aligned}
c_n &= \int_0^1 f_\sigma(x)\exp(2\pi i n x)\,dx \\
&= \int_0^1 \sum_{k=-\infty}^{\infty} \phi_\sigma(x+k)\exp(2\pi i n x)\,dx \\
&= \sum_{k=-\infty}^{\infty} \int_0^1 \phi_\sigma(x+k)\exp(2\pi i n x)\,dx \\
&= \sum_{k=-\infty}^{\infty} \int_k^{k+1} \phi_\sigma(x)\exp(2\pi i n x)\,dx \\
&= \int_{-\infty}^{\infty} \phi_\sigma(x)\exp(2\pi i n x)\,dx \\
&= \int_{-\infty}^{\infty} \phi_\sigma(x)\cos(2\pi n x)\,dx \\
&= \exp(-(2\pi n\sigma)^2/2)
\end{aligned}$$

Since ϕ is an even function, we are able to replace the complex exponential with just a cosine in the penultimate step above, giving us a discrete cosine transform, instead of the Fourier transform. Adding up all the Fourier terms then gives the result

$$\begin{aligned}
f_\sigma(x) &= \sum_{k=-\infty}^{\infty} c_k\cos(2\pi k x) = \sum_{k=-\infty}^{\infty} \exp(-(2\pi k\sigma)^2/2)\cos(2\pi k x) \\
&= 1 + 2\sum_{k=1}^{\infty} \exp(-2(\pi k\sigma)^2)\cos(2\pi k x) \qquad (2)
\end{aligned}$$

Thus, especially for large values of σ we have a uniform distribution, with a small cosine perturbation, and very much smaller higher order terms.

6 Computing the density

Equations 1 and 2, allow for very efficient computation of the density, because they are complimentary. While they are both infinite sums, equation 1 converges rapidly when σ is small, while equation 2 converges faster when

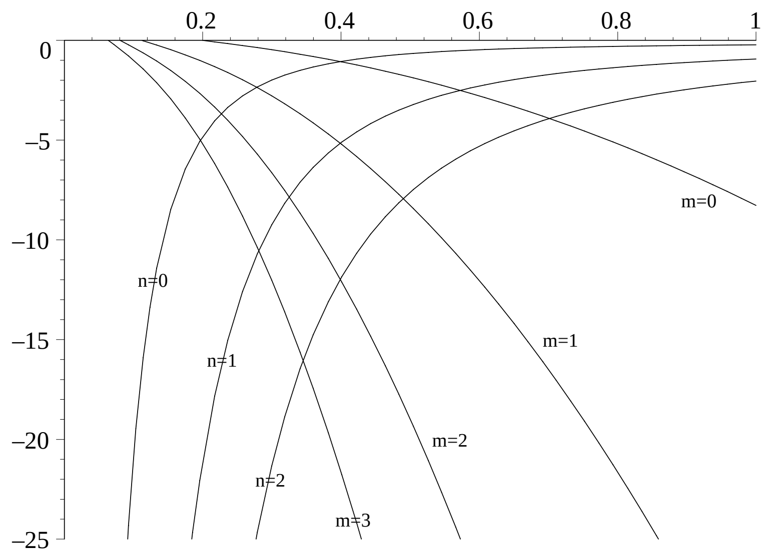

Figure 3. $\log_{10}$ of the errors in the truncated series approximations of equations 3 and 4

σ is large. If we replace the infinite sums by finite approximations, we get the following two approximations

$$f_\sigma(x) \approx \sum_{k=-n}^{n} \phi_\sigma(x+k) \tag{3}$$

$$f_\sigma(x) \approx 1 + 2\sum_{k=1}^{m} \exp(-2(\pi k \sigma)^2)\cos(2\pi k x) \tag{4}$$

Figure 3 shows the number of digits of accuracy in the maximum error using the approximations of equations 3 and 4. Starting from the left, the 3 rising curves are the one, two and three term approximations using equation 3. As can be seen they are best (most negative) for small values of σ. Continuing to the right we have the three, two and then one term approximations using equation 4. Finally the last line which gets to -8 when $\sigma = 1$ is the approximation to the density by the constant 1 (or the zero term approximation using equation 4). Thus with a three term computation, we can easily get better than 15 digit accuracy for the density.