Random numbers from *continuous* probability distributions

It is often useful to generate random numbers which have a known distribution. In many cases this might be the Normal distribution, but specific scenarios may be better modelled by something else such as a *Rayleigh*, *Exponential*, *Gamma* or *Weibull* form.

As long as we are able to **generate uniformly distributed random numbers** in the range [0,1], *and* we are able to evaluate the **inverse cumulative distribution function**, we can generate random numbers from *any* continuous probability distribution.

$$\int_{-\infty}^{\infty} p(x)dx = 1$$ Defining property of *any* distribution *p(x)*

$$E[x] = \mu = \int_{-\infty}^{\infty} xp(x)dx$$ Mean

$$V[x] = E\left[x^2\right] - \mu^2 = \sigma^2 = \int_{-\infty}^{\infty} x^2 p(x)dx - \mu^2$$ Variance

Note means, variances (and higher *moments* relating to 'skewness' and 'kurtosis') can often be more easily determined using the **Moment Generating Function** (MGF) of a distribution.

$$M_x(t) = E[e^{xt}]$$

$$E[x^n] = \left.\frac{\partial^n M}{\partial t^n}\right|_{t=0}$$

The MGF is useful for *both* continuous and discrete distributions. However, the concept of **Probability Generating Functions** (PGF) is also useful for discrete distributions.

How to generate the random numbers x

$$F(x) = P(u \le x) = \int_{-\infty}^{x} p(u)du$$ Cumulative distribution function

$$r \sim \mathrm{U}(0,1)$$

$$\therefore x = F^{-1}(r)$$

This will generate random numbers with probability density *p(x)*, since the range of the cumulative distribution function is:

$$\boxed{0 \le F(r) \le 1}$$

The generation of uniformly distributed random numbers is a function in many computer languages such as Python, C or MATLAB.
`x = rand( a, b)` is the code in MATLAB. This will generate an array of size *ab* with random numbers in the range [0,1].

These functions are typically pseudorandom, i.e. based upon an *algorithm*. For true randomness, physical 'noise' is required. You can download samples of atmospheric noise from
https://www.random.org

Uniform distribution

$$x \sim \mathrm{U}(a,b)$$

$$p(x) = \begin{cases} \frac{1}{b-a} & a \le x \le b \\ 0 & \text{otherwise} \end{cases}$$

$$F(x) = \begin{cases} 0 & x \le a \\ \frac{x-a}{b-a} & a \le x \le b \\ 1 & x \ge b \end{cases}$$

$$F^{-1}(x) = a + (b-a)x$$

$$M_x(t) = \begin{cases} \dfrac{e^{tb} - e^{ta}}{t(b-a)} & t \ne 0 \\ 1 & t = 0 \end{cases}$$

$$\mu = \tfrac{1}{2}(a+b)$$

$$\sigma^2 = \tfrac{1}{12}(b-a)^2$$

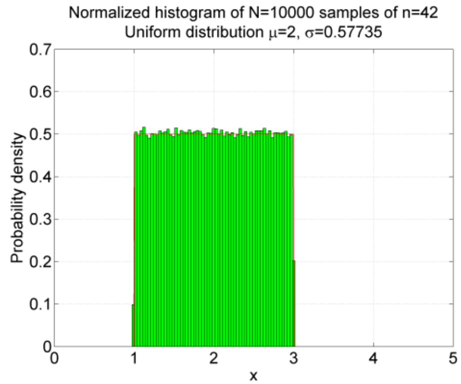

A **normalized histogram** of samples should match the predicted distribution. A normalized histogram is one where the total area of the bars is unity.

To **normalize**, scale the probability density of the histogram by the **sum of the probability densities** x **the bar width.**

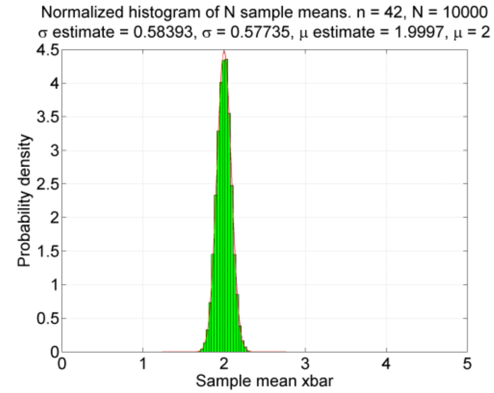

The **Central Limit Theorem** states that if the sample size n is large enough, the distribution of sample means tends to a Normal distribution, regardless of the *population* distribution.

$$\bar{x} \sim \mathrm{N}\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

The distribution mean is the population mean μ and standard deviation is the population standard deviation σ divided by the square root of the sample size n.

An example non-normal distribution: a "parabolic distribution"

$$p(x) = \begin{cases} \frac{3}{2}(x-1)^2 & 0 \le x \le 2 \\ 0 & \text{otherwise} \end{cases} \qquad F(x) = \begin{cases} 0 & x < 0 \\ \frac{1}{2}\left((x-1)^3 + 1\right) & 0 \le x \le 2 \\ 1 & x \ge 2 \end{cases}$$

$$F^{-1}(x) = \sqrt[3]{2x-1} + 1 \qquad \mu = 1 \qquad \sigma^2 = \tfrac{3}{5}$$

The distribution of sample means still tends to a Normal distribution!

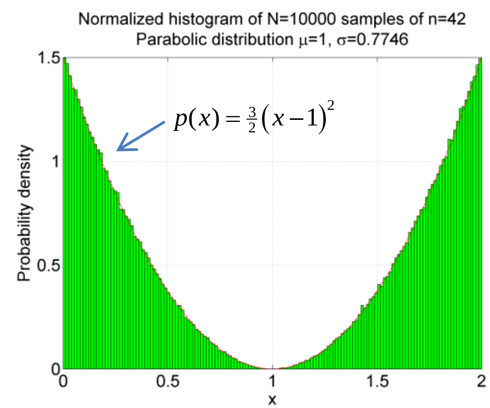

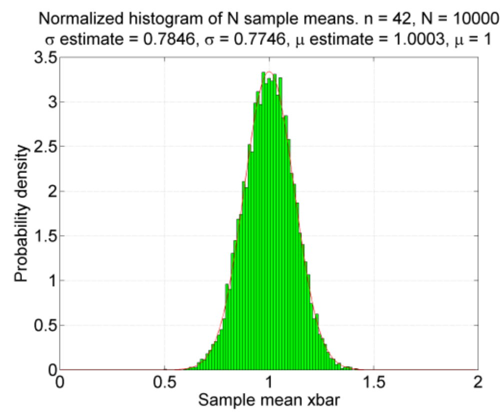

Normal Distribution

The classic 'bell shaped' curve. Also known as the *Gaussian* distribution. Mean μ corresponds to the peak of the distribution and standard deviation σ scales the width. The *Central Limit Theorem* states the sample means of *any* distribution will tend to a Normal distribution if the sample size is large enough.

$$x \sim \mathrm{N}(\mu,\sigma)$$
$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$
$$F(x) = \tfrac{1}{2} + \tfrac{1}{2}\mathrm{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)$$
$$F^{-1}(x) = \mu + \sigma\sqrt{2}\mathrm{erf}^{-1}(2x-1)$$
$$M_x(t) = e^{\mu t + \frac{1}{2}\sigma^2 t^2}$$

$$z = \frac{x-\mu}{\sigma}$$
$$z \sim \mathrm{N}(0,1)$$
$$p(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2}$$
$$\Phi(z) = \tfrac{1}{2} + \tfrac{1}{2}\mathrm{erf}\left(\frac{z}{\sqrt{2}}\right)$$
$$\Phi^{-1}(z) = \sqrt{2}\mathrm{erf}^{-1}(2z-1)$$

Standard normal distribution using substitution z

Exponential distribution

Describes the time between independent events which occur at a given rate i.e. a *Poisson* process.

$$x \sim \mathrm{Exp}(\lambda)$$
$$p(x) = \lambda e^{-\lambda x}$$
$$F(x) = 1 - e^{-\lambda x}$$
$$F^{-1}(x) = \frac{-\ln(1-x)}{\lambda}$$
$$M_x(t) = \frac{\lambda}{\lambda - t} \quad ; \quad t < \lambda$$
$$\mu = \tfrac{1}{\lambda}$$
$$\sigma = \tfrac{1}{\lambda}$$

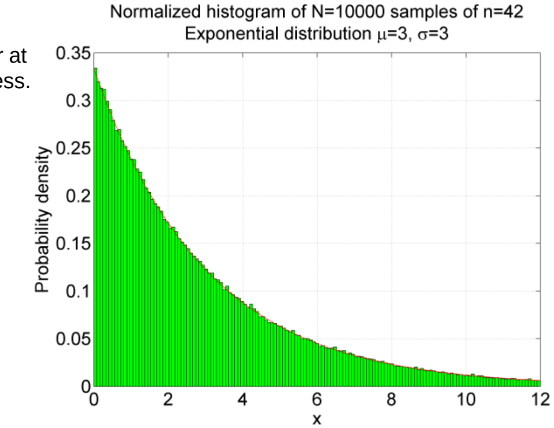

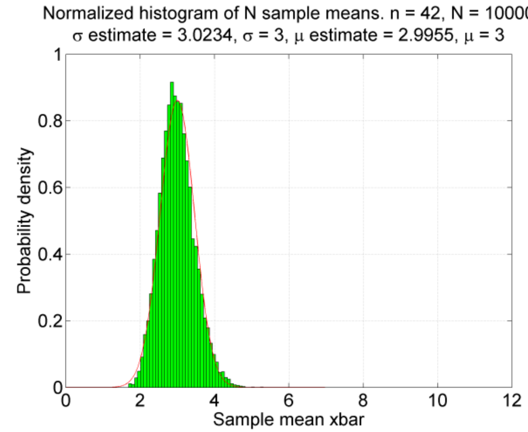

Rayleigh distribution

When orthogonal x,y components of a vector are Normally distributed, the magnitude of the vector is Rayleigh distributed.

$$x \sim \mathrm{Ra}(k)$$
$$p(x) = \frac{x}{k^2} e^{-\frac{x^2}{2k^2}}$$
$$F(x) = 1 - e^{-\frac{x^2}{2k^2}}$$
$$F^{-1}(x) = k\sqrt{-2\ln(1-x)}$$
$$M_x(t) = 1 + \sigma t e^{\frac{1}{2}\sigma^2 t^2}\sqrt{\tfrac{1}{2}\pi}\left(\mathrm{erf}\left(\frac{\sigma t}{\sqrt{2}+1}\right)\right)$$
$$\mu = \sqrt{\tfrac{\pi}{2}}k \qquad \sigma = \sqrt{\tfrac{4-\pi}{2}}k$$

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}dt$$ Error function

$$\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt$$ Gamma function

$$\gamma(x,a) = \frac{1}{\Gamma(a)}\int_0^x t^{a-1}e^{-t}dt$$ Lower incomplete gamma function

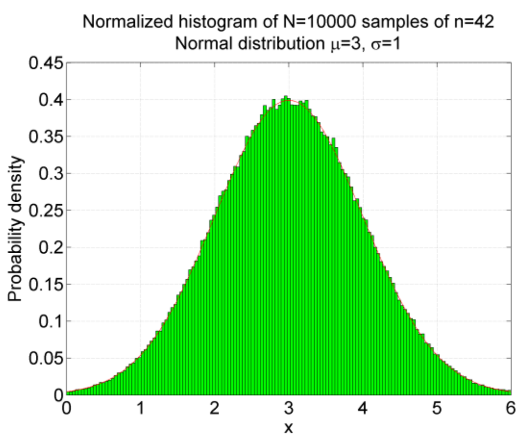

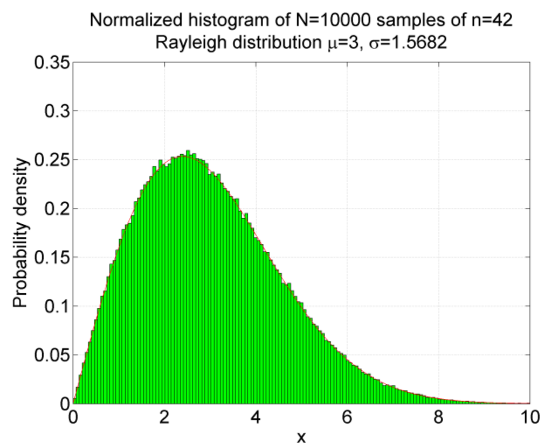

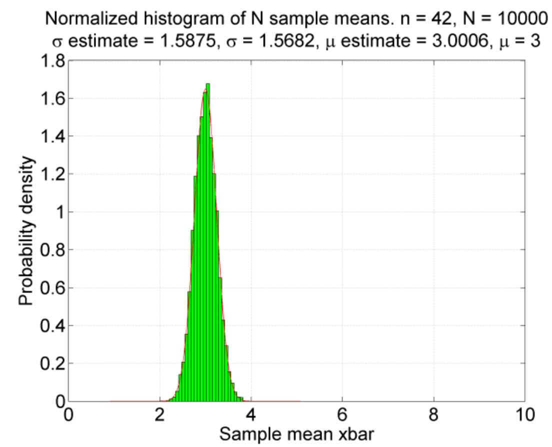

Planck distribution of black-body radiation of wavelength λ

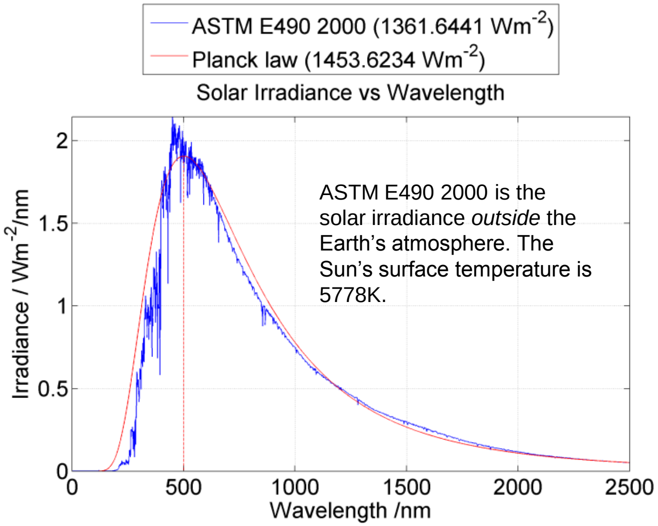

$$\boxed{B(\lambda) = \frac{2\pi hc^2}{\lambda^5}\frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1}}$$

Red	620-750nm
Yellow	570-590nm
Green	495-570nm
Blue	450-495nm

$k_B = 1.381 \times 10^{-23}\,\mathrm{JK}^{-1}$	Boltzmann's constant
$h = 6.626 \times 10^{-34}\,\mathrm{m^2kgs^{-1}}$	Planck's constant
$c = 2.998 \times 10^{8}\,\mathrm{ms^{-1}}$	Speed of light

$$I = \int_0^\infty B(\lambda, T)\,d\lambda = \sigma T^4$$ Total radiation intensity (Wm-2)

$$\sigma = \frac{2\pi^5 k_B^4}{15c^2h^3} \approx 5.67 \times 10^{-8}\,\mathrm{Wm^{-2}K^{-4}}$$ Stefan-Boltzmann constant

Note: $\int_0^\infty \frac{x^3}{e^x - 1}dx = \frac{\pi^4}{15}$

Maxwell-Boltzmann distribution of molecular speeds v in an *ideal gas*.
m is the molecular mass and T the absolute temperature.

$$p(v)dv = 4\pi\left(\frac{m}{2\pi k_B T}\right)^{\frac{3}{2}} v^2 e^{-\frac{\frac{1}{2}mv^2}{k_B T}}\,dv$$

$$\overline{v^2} = E\left[v^2\right] = \frac{3k_B T}{m} \quad \therefore v_{rms} = \sqrt{\overline{v^2}} = \sqrt{\frac{3k_B T}{m}}$$

$$p(\varepsilon)\,d\varepsilon = \frac{2}{\sqrt{\pi}}(k_B T)^{-\frac{3}{2}}\sqrt{\varepsilon}\,e^{-\frac{\varepsilon}{k_B T}}\,d\varepsilon$$

$$\boxed{E\left[\tfrac{1}{2}mv^2\right] = \tfrac{3}{2}k_B T}$$

$$\boxed{E[\varepsilon] = \text{degrees of freedom} \times \tfrac{1}{2}k_B T}$$

Useful standard integrals

$$I_n = \int_0^\infty x^n e^{-ax^2}dx$$

$$I_n = \frac{n-1}{2a}I_{n-2}$$

$$I_0 = \tfrac{1}{2}\sqrt{\tfrac{\pi}{a}}$$

$$I_1 = \tfrac{1}{2a}$$

$$I_2 = \frac{\sqrt{\pi}}{4a^{\frac{3}{2}}}$$

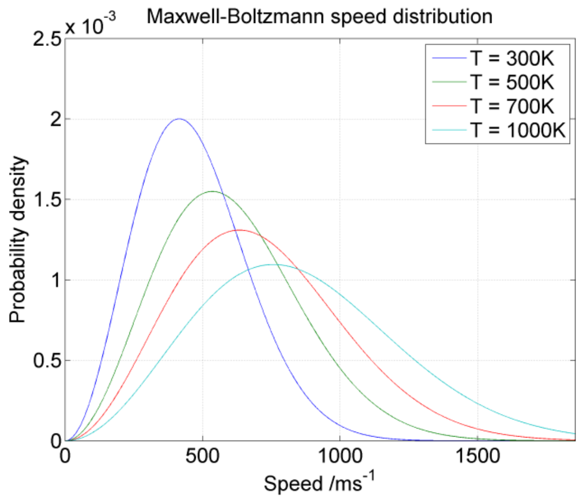

Weibull distribution

$$x \sim \mathrm{W}(k,\lambda)$$

$$p(x) = \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-\left(\frac{x}{\lambda}\right)^k} \quad ; \quad 0 \le x < \infty$$

$$F(x) = 1 - e^{-\left(\frac{x}{\lambda}\right)^k}$$

$$F^{-1}(x) = \lambda\left(-\ln(1-x)\right)^{\frac{1}{k}}$$

$$M_x(t) = \sum_{n=0}^{\infty} \frac{t^n \lambda^n}{n!} \Gamma\left(1+\frac{n}{k}\right) \ ; \ k \ge 1$$

$$\mu = \lambda \times \Gamma\left(1+\tfrac{1}{k}\right)$$

$$\sigma^2 = \lambda^2 \Gamma\left(1+\tfrac{2}{k}\right) - \mu^2$$

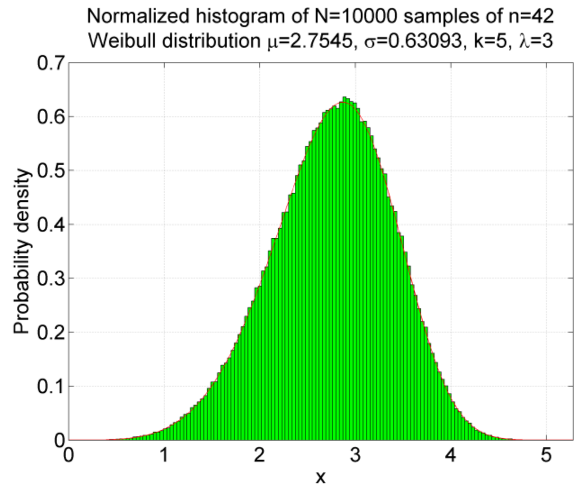

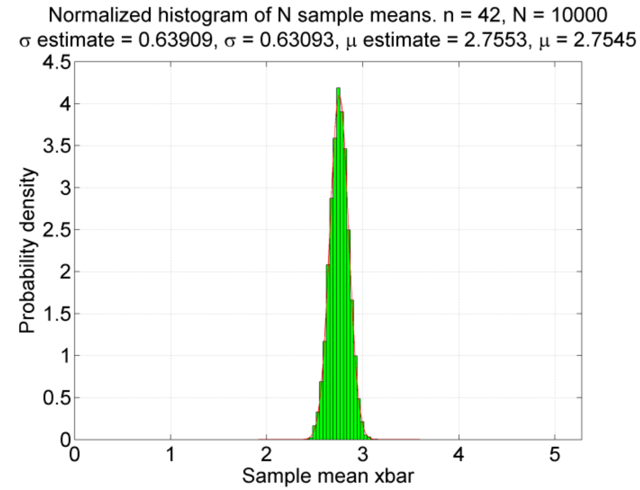

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$$

Gamma function

$$\gamma(x,a) = \frac{1}{\Gamma(a)} \int_0^x t^{a-1} e^{-t} dt$$

Lower incomplete gamma function

Gamma distribution

Note this is the distribution of a *sum* of independent, Exponentially distributed random variables

$$x \sim \mathrm{Gamma}(k,\theta) \ ; \ k,\theta > 0$$

$$p(x) = \frac{x^{k-1}}{\Gamma(k)\theta^k} e^{-\frac{x}{\theta}}$$

$$F(x) = \gamma\left(\frac{x}{\theta}, k\right)$$

$$F^{-1}(x) = \theta\gamma^{-1}(x,k)$$

$$M_x(t) = (1-\theta t)^{-k} \quad ; \quad t < \frac{1}{\theta}$$

$$\mu = k\theta$$

$$\sigma^2 = k\theta^2$$

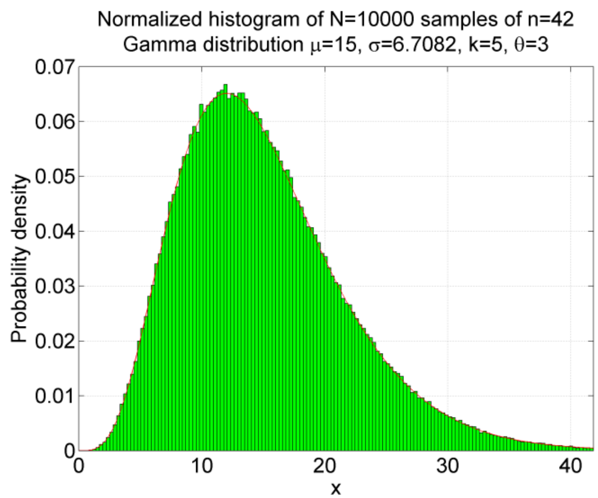

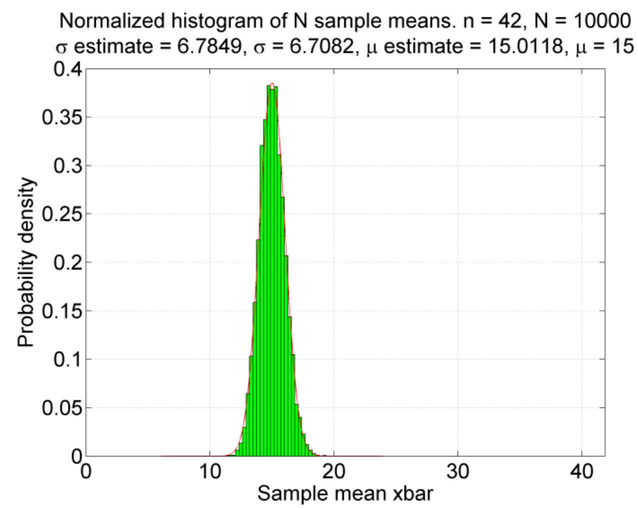

χ^2 distribution

The distribution of sum of squares of independent random variables *which are themselves distributed* by a Normal distribution with mean 0 and standard deviation 1.

$$z_i \sim \mathrm{N}(0,1)$$
$$x = \sum_{i=1}^{k} z_i^2$$
$$x \sim \chi^2(k)$$
$$p(x) = \frac{x^{\frac{1}{2}k-1}e^{-\frac{1}{2}x}}{2^{\frac{1}{2}k}\Gamma\left(\frac{1}{2}k\right)}$$
$$F(x,k) = \gamma\left(\tfrac{1}{2}x, \tfrac{1}{2}k\right)$$
$$F^{-1}(x,k) = 2\gamma^{-1}\left\{x, \tfrac{1}{2}k\right\}$$
$$M_x(t) = (1-2t)^{-\frac{1}{2}k} \quad ; \quad t < \tfrac{1}{2}$$
$$\mu = k$$
$$\sigma^2 = 2k$$

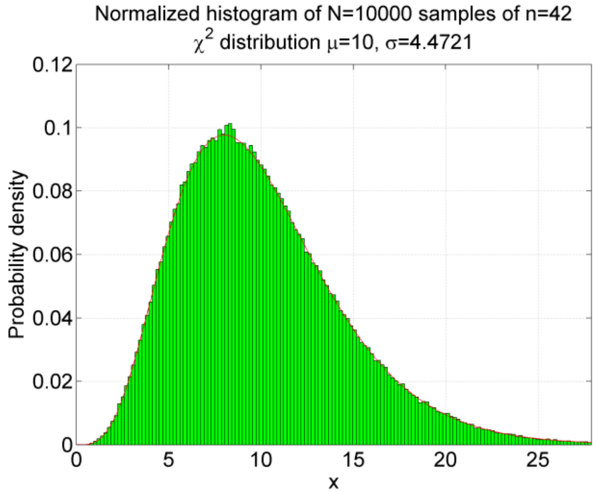

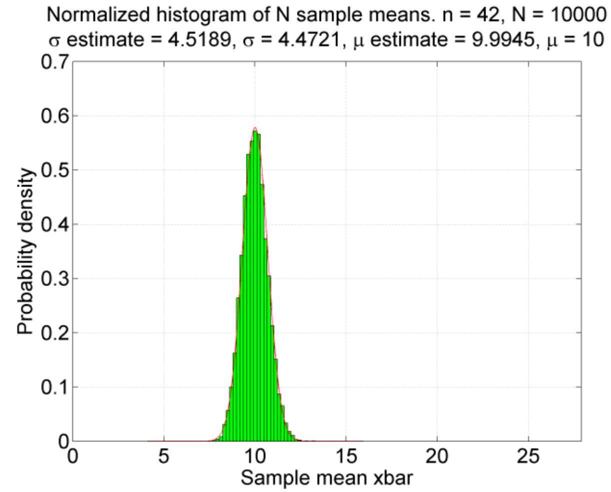

$$\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt$$

Gamma function

$$\gamma(x,a) = \frac{1}{\Gamma(a)}\int_0^x t^{a-1}e^{-t}dt$$

Lower incomplete gamma function

Student's *t*-distribution

If a sample size n is small and the sample is thought to derive from a normal distribution, the t-statistic will assess the possible range of values that the population mean will occur, given a sample mean and (unbiased) sample standard deviation.

$$t = \frac{\mu - \bar{x}}{\sqrt{s^2/n}}$$
$$t \sim \mathrm{tdist}(v)$$
$$p(t,v) = \frac{\Gamma\left(\frac{1}{2}(1+v)\right)}{\sqrt{v\pi}\,\Gamma\left(\frac{1}{2}v\right)}\left(1+\frac{t^2}{v}\right)^{-\frac{1}{2}(1+v)} \quad ; \quad v = n-1 \quad ; \quad n \in \mathbb{Z}^+$$
$$p(t,v) = \mathrm{tpdf}(t,v)$$
$$F(t,v) = \mathrm{tcdf}(t,v)$$
$$F^{-1}(t,v) = \mathrm{tinv}(t,v)$$
$$\mu = 0$$
$$\sigma^2 = \frac{v}{v-2}$$

These MATLAB functions require the *Statistics Toolbox*

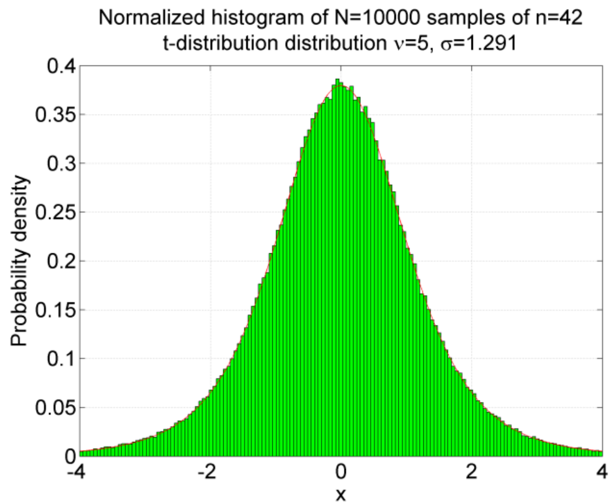

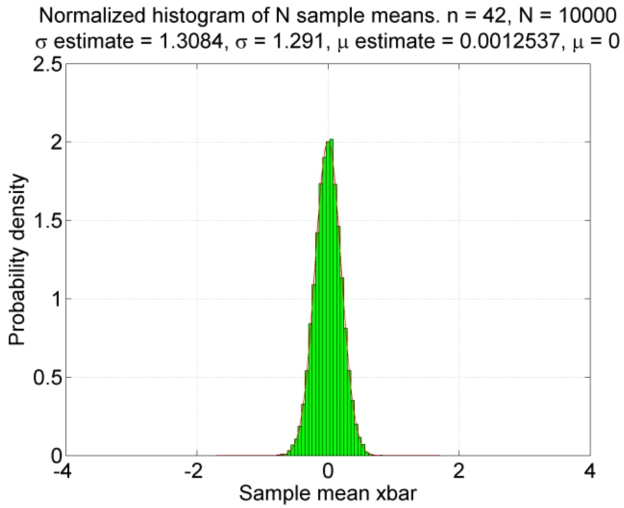

A selection of *discrete* probability distributions

The set of integers 0,1,2.... are the random variables. Each integer has a defined probability.

Geometric distribution

The random variable x is the number of *binary trials* (i.e. success or failure) up to, and including, the first success.
p is a fixed probability of success in each independent trial.

$$x \sim \mathrm{Geo}\left(p\right)$$
$$p(x,p) = (1-p)^{x-1}p$$
$$M_x(t) = \frac{pe^t}{1-(1-p)e^t} \;\; ; \; t < -\ln(1-p)$$
$$\mu = \tfrac{1}{p}$$
$$\sigma^2 = \frac{1-p}{p^2}$$

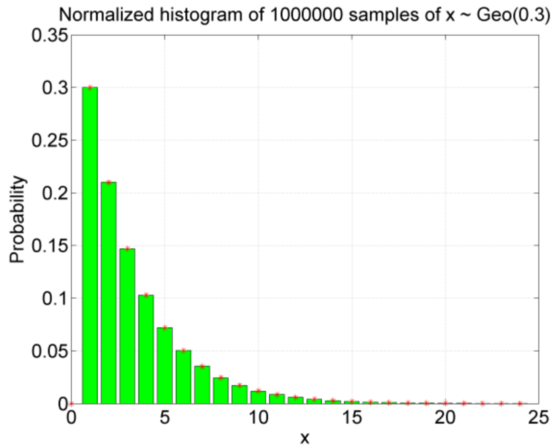

Binomial distribution

The random variable x is the number of successes out of n independent binary (success or failure) trials.
p is a fixed probability of success in each independent trial.

$$x \sim \mathrm{B}\left(n,p\right)$$
$$p(x,n,p) = \binom{n}{x} p^x \left(1-p\right)^{n-x}$$
$$M_x(t) = \left(1-p+pe^t\right)^n$$
$$\mu = np$$
$$\sigma^2 = np(1-p)$$

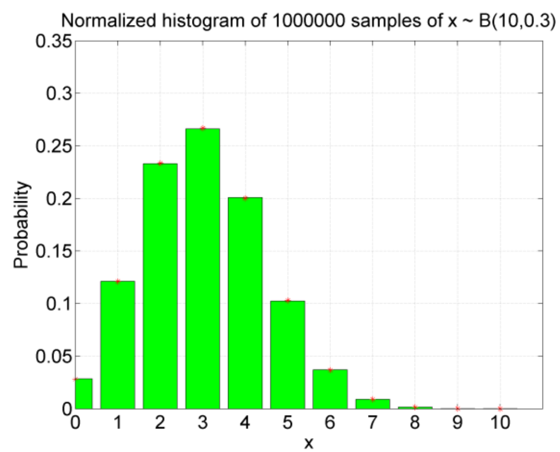

Poisson distribution

The random variable x is the number occurrences (e.g. goals, telephone calls ....) in a set interval of time, given a mean rate of occurrence λ.

$$x \sim \mathrm{Po}\left(\lambda\right)$$
$$p(x,\lambda) = \frac{\lambda^x e^{-\lambda}}{x!}$$
$$M_x(t) = e^{\lambda\left(e^t-1\right)}$$
$$\mu = \lambda$$
$$\sigma^2 = \lambda$$

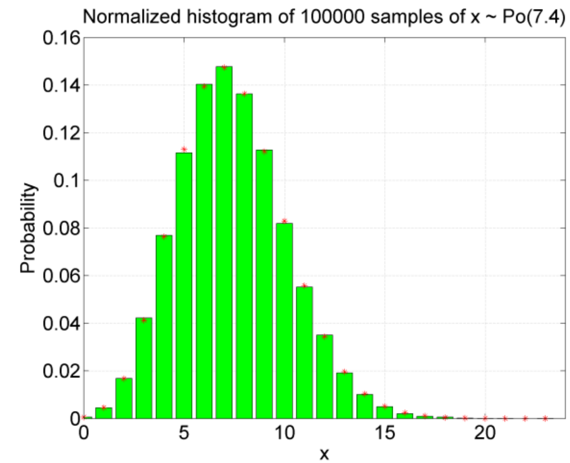

Random distribution

$$x_i \sim \mathrm{U}(0,1)$$
$$p_n = \frac{x_n}{\sum_{j=0}^{N-1} x_j}$$
$$\mu = \sum_{n=0}^{N-1} n \times p_n$$
$$\sigma^2 = \sum_{n=0}^{N-1} n^2 \times p_i - \mu^2$$

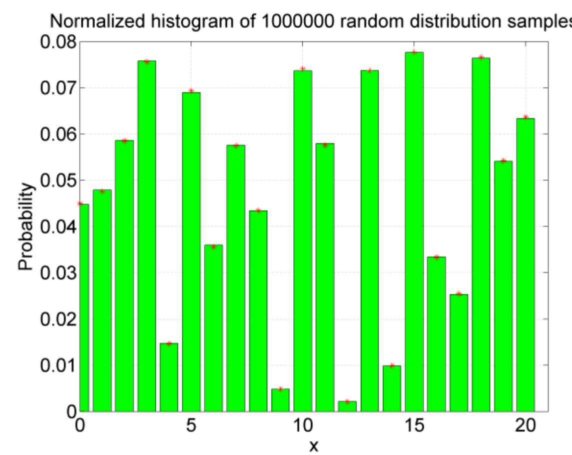

Generating random integers from discrete probability distributions

The fact that the sum of the probabilities in a discrete distribution must sum to unity can be used to generate random integers, assuming it is possible to generate a random number within the range [0,1].

Use the probabilities to form the edges of a series of 'boxes' which span the interval [0,1]. For every random fraction ~ U(0,1), determine the box number which encloses the fraction. This box number is the random variable.

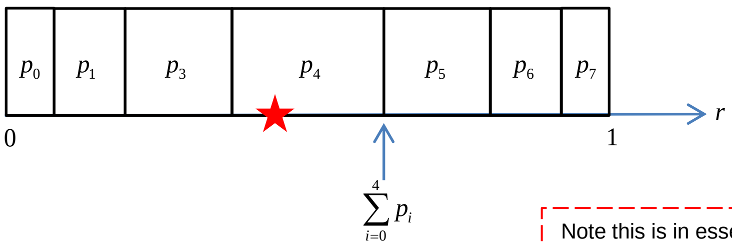

Example: consider a discrete distribution of eight possible probabilities, for the random integers 0...7 as show above. The widths of the boxes correspond to the probabilities.

A random number $r \sim \mathrm{U}(0,1)$ is chosen. ★
This happens to be in the range:

$$\sum_{i=0}^{3} p_i \le r < \sum_{i=0}^{4} p_i$$ Cumulative distribution function

so in this case the random number $x = 4$ is selected.

Note this is in essence the same process as random number generation from continuous distributions. i.e. the output of the **inverse cumulative distribution function** with input being a random number from U(0,1)

MATLAB code to generate the Binomial distribution example

```
%Number of binomially distributed numbers
N = 1e6;

%Success probability
p = 0.3;

%Number of binary trials
n = 10;

%Fontsize for graphs
fsize = 18;

%Distribution name
name = 'binomial';

%Graph title string
title_str = ['Normalized histogram of ',num2str(N),...
    ' samples of x ~ B(',num2str(n),',',num2str(p),')'];

%Determine binomial probabilities P
P = zeros(1,n+1);
for k=0:n
    P(k+1) = nchoosek(n,k)*(p^k)*(1-p)^(n-k);
end

%Determine array of cumulative probabilities
CP = cumsum(P);
if CP(end)~=1
    CP = [CP,1];
end

%Generate N uniformly distributed random numbers in range [0,1]
z = rand(1,N);

%Define an array of cumulative probability, one element shifted.
%This is used to efficiently determine the 'bin' corresponding to z
CP0 = [0,CP(1:end-1)];

%For each number, determine the 'bin' number of a division of the interval
%[0,1] by the cumulative probabilities. If the first bin is zero, then
%these numbers should be distributed according to the probability
%distribution used to generate the cumulative probabilities.
x = zeros(1,N);
for k=1:length(CP)
    test = ( z >= CP0(k) ) & ( z <= CP(k) );
    x( test ) = k - 1;
end

%Ignore any values from final bin, if CP is longer than P
if numel(CP)>numel(P)
    x( x==numel(P) ) = [];
end

%Determine a normalized histogram of x values
xx = 0 : length(P) - 1;
h = hist( x, xx )/N;

%Plot normalized histogram
b = bar(xx,h);
set(b,'facecolor','g');
xlabel('x','fontsize',fsize);
ylabel('Probability','fontsize',fsize);
title(title_str,'fontsize',fsize);
grid on;
box on;
set(gca,'fontsize',fsize);
xlim([0,max(xx)+1]);

%Overlay probability distribution
hold on
plot(xx,P,'r*');

%Print PNG of graph
print(gcf,[name,'.png'],'-dpng','-r300');
close(gcf);
```