Applied Statistics

Probability Density Functions (PDFs)





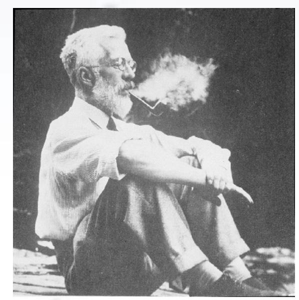
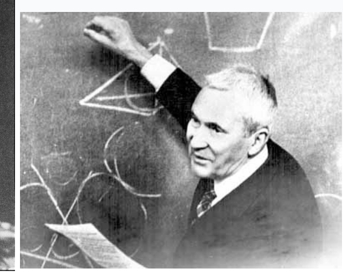

Troels C. Petersen & Mathias L. Heltberg (NBI)

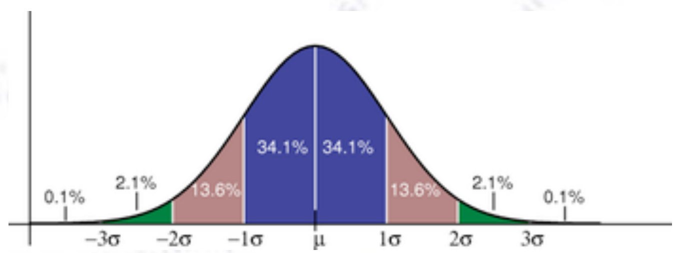

"Statistics is merely a quantisation of common sense"

Probability Density Functions

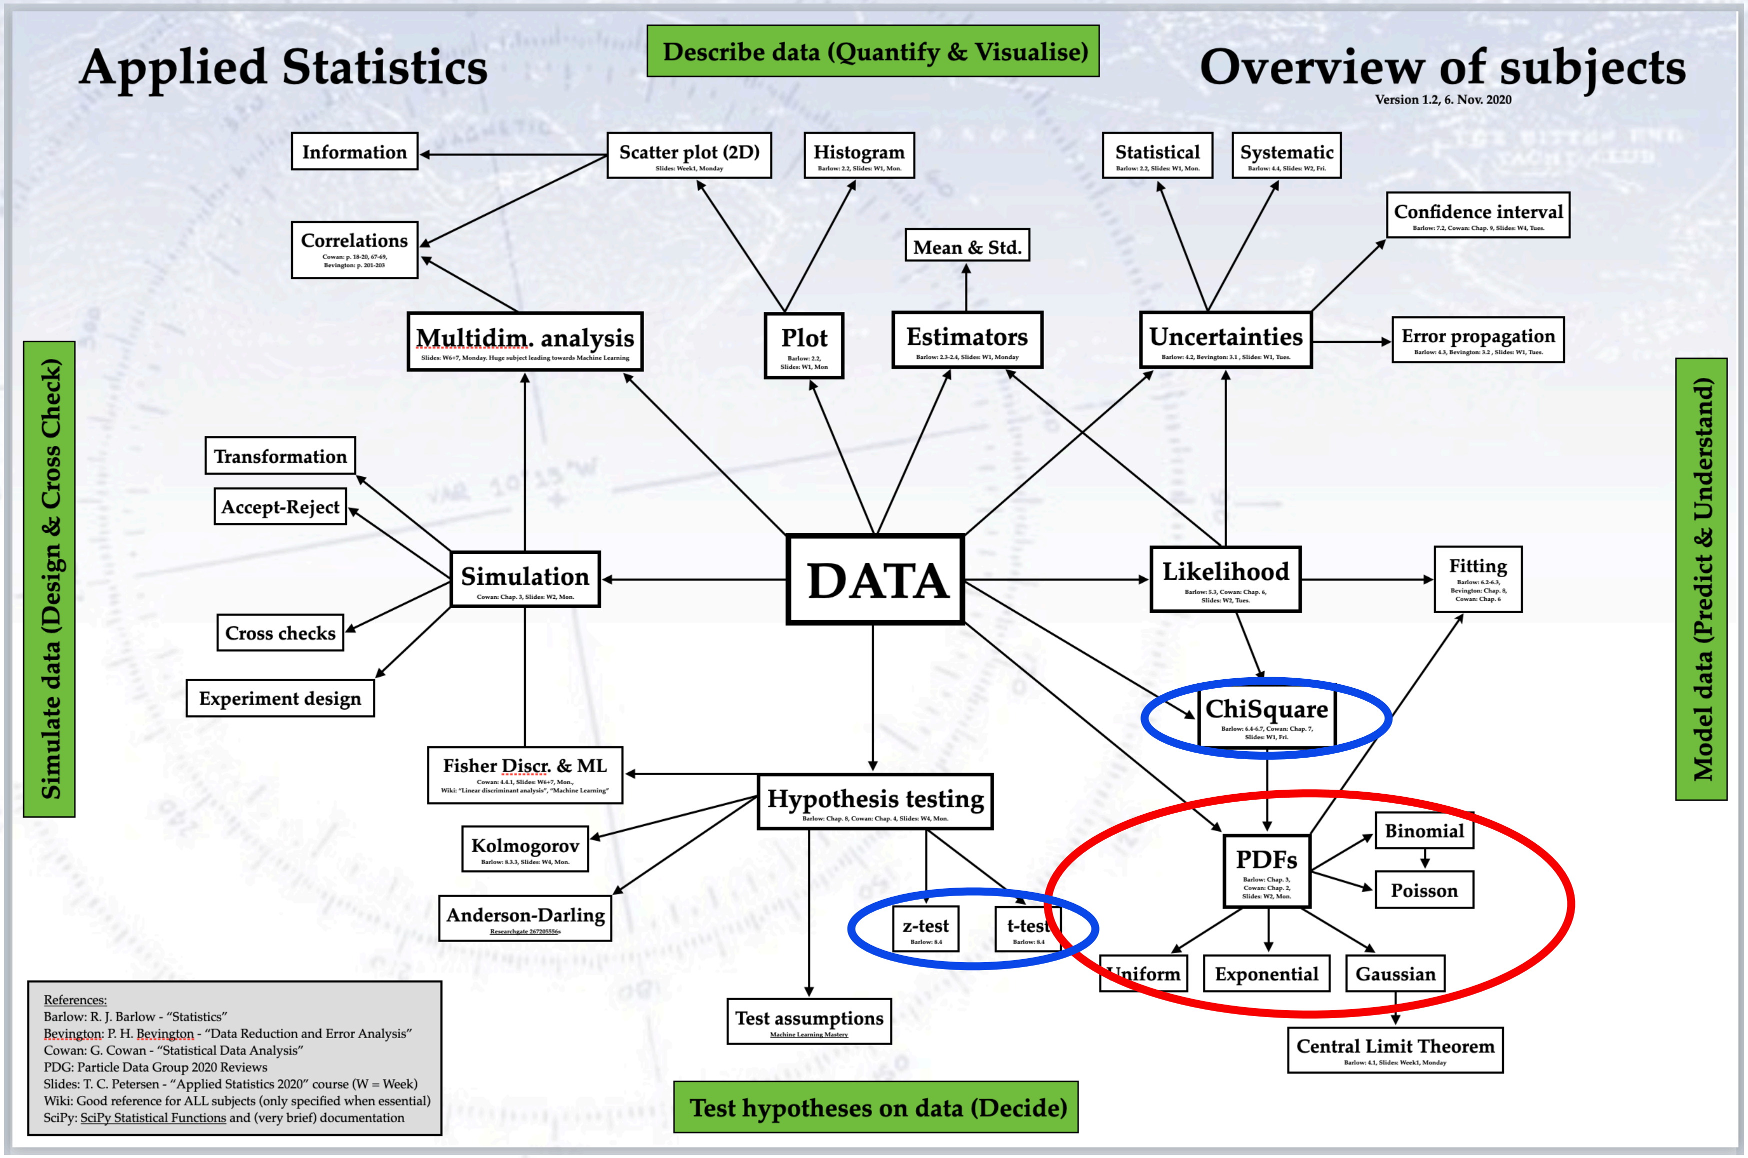

Probability Density Functions

A Probability Density Function (PDF) f(x) describes the probability of an outcome x:

probability to observe x in the interval [x, x+dx] = f(x) dx

PDFs are required to be normalised:

$$\int_S f(x) dx = 1$$

The expectation value (aka. mean) and the variance (i.e. standard deviation squared) are then defined as follows:

$$\mu = \int_{-\infty}^{\infty} x f(x) dx$$

$$\sigma^2 = \int_{-\infty}^{\infty} (x - \mu)^2 f(x) dx$$

Probability Density Functions

Just a word on the probability lingo:

In the literature it is often we use large letters for a random variable X. This means an *outcome* for an event! If I roll a die, we say that X takes on values in {1,2,3,4,5,6}.

Now the small letters x, is a normal variable that is a real number. So we could write: P(X<x), which translated means that we calculate the probability that in one event, we obtain a variable with the value smaller than the real value x.

Probability Density Functions

<u>Example:</u>

Consider a uniform distribution:

$$f(x) = \begin{cases} 1 & x \in [0,1] \\ 0 & else \end{cases}$$

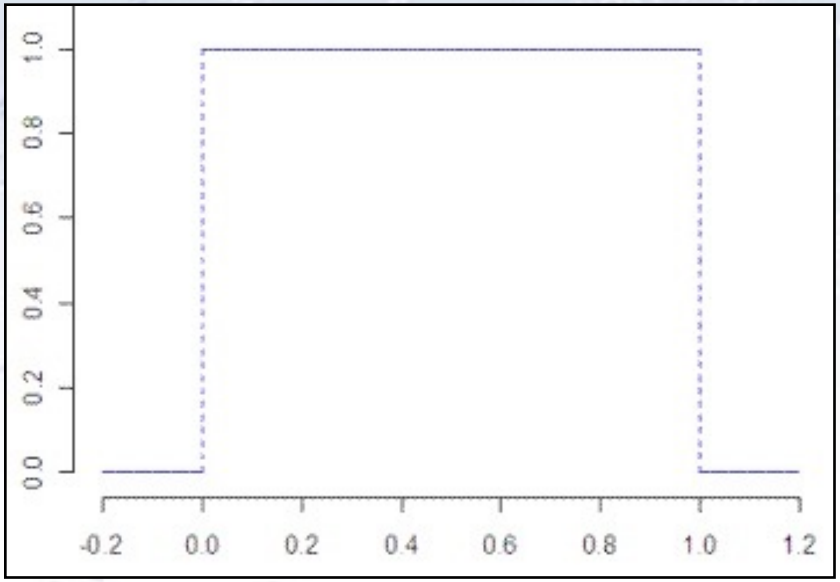

Calculating the mean and width:

$$\mu = \int_{-\infty}^{\infty} x f(x) dx = \int_0^1 x dx = [\frac{1}{2}x^2]_0^1 = \frac{1}{2}$$

$$\sigma^2 = \int_{-\infty}^{\infty} (x-\mu)^2 f(x) dx = \int_0^1 (x - \frac{1}{2})^2 dx =$$

$$[\frac{1}{3}x^3 - \frac{1}{2}x^2 + \frac{1}{4}x]_0^1 = \frac{1}{3} - \frac{1}{2} + \frac{1}{4} = \frac{1}{12}$$

Probability Density Functions

The number of PDFs is infinite, and nearly so is the list of known ones:

Discrete distributions [edit source | edit beta]

With finite support [edit source | edit beta]

- The Bernoulli distribution, which takes value 1 with
- The Rademacher distribution, which takes value 1
- The binomial distribution, which describes the numb
- The beta-binomial distribution, which describes the
- The degenerate distribution at x_0, where X is certa random variables in the same formalism.
- The discrete uniform distribution, where all element shuffled deck.
- The hypergeometric distribution, which describes t there is no replacement.
- The Poisson binomial distribution, which describes
- Fisher's noncentral hypergeometric distribution
- Wallenius' noncentral hypergeometric distribution
- Benford's law, which describes the frequency of th

With infinite support [edit source | edit beta]

- The beta negative binomial distribution
- The Boltzmann distribution, a discrete distribution i analogue. Special cases include:
  - The Gibbs distribution
  - The Maxwell–Boltzmann distribution
- The Borel distribution
- The extended negative binomial distribution
- The extended hypergeometric distribution
- The generalized log-series distribution
- The generalized normal distribution
- The geometric distribution, a discrete distribution w
- The hypergeometric distribution
- The logarithmic (series) distribution
- The negative binomial distribution or Pascal distribu
- The parabolic fractal distribution
- The Poisson distribution, which describes a very la Poisson, the hyper-Poisson, the general Poisson b
  - The Conway–Maxwell–Poisson distribution, a tw
- The Polya-Eggenberger distribution
- The Skellam distribution, the distribution of the diff
- The skew elliptical distribution
- The skew normal distribution
- The Yule–Simon distribution
- The zeta distribution has uses in applied statistics
- Zipf's law or the Zipf distribution. A discrete power-
- The Zipf–Mandelbrot law is a discrete power law dis

Continuous distributions [edit source | edit beta]

Supported on a bounded interval [edit source | edit

- The Arcsine distribution on [a,b], which is a speci
- The Beta distribution on [0,1], of which the unifor
- The Logitnormal distribution on (0,1).
- The Dirac delta function although not strictly a fu but the notation treats it as if it were a continuous
- The continuous uniform distribution on [a,b], wher
  - The rectangular distribution is a uniform distrib
- The Irwin-Hall distribution is the distribution of the
- The Kent distribution on the three-dimensional sph
- The Kumaraswamy distribution is as versatile as t
- The logarithmic distribution (continuous)
- The PERT distribution is a special case of the bet
- The raised cosine distribution on $[\mu - s, \mu + s]$
- The reciprocal distribution
- The triangular distribution on [a, b], a special case
- The truncated normal distribution on [a, b].
- The U-quadratic distribution on [a, b].
- The von Mises distribution on the circle.
- The von Mises-Fisher distribution on the N-dimens
- The Wigner semicircle distribution is important in t

Supported on semi-infinite intervals, usually [0,∞)

- The Beta prime distribution
- The Birnbaum–Saunders distribution, also known
- The chi distribution
  - The noncentral chi distribution
- The chi-squared distribution, which is the sum of t
  - The inverse-chi-squared distribution
  - The noncentral chi-squared distribution
  - The Scaled-inverse-chi-squared distribution
- The Dagum distribution
- The exponential distribution, which describes the
- The F-distribution, which is the distribution of the ratio of two chi-squared variates which are not no
  - The noncentral F-distribution
- Fisher's z-distribution
- The folded normal distribution
- The Fréchet distribution
- The Gamma distribution, which describes the time
  - The Erlang distribution, which is a special case
  - The inverse-gamma distribution
- The generalized Pareto distribution
- The Gamma/Gompertz distribution
- The Gompertz distribution
- The half-normal distribution

- Hotelling's T-squared distribution
- The inverse Gaussian distribution, also kn
- The Lévy distribution
- The log-Cauchy distribution
- The log-gamma distribution
- The log-Laplace distribution
- The log-logistic distribution
- The log-normal distribution, describing vari
- The Mittag–Leffler distribution
- The Nakagami distribution
- The Pareto distribution, or "power law" dist
- The Pearson Type III distribution
- The phased bi-exponential distribution is c
- The phased bi-Weibull distribution
- The Rayleigh distribution
- The Rayleigh mixture distribution
- The Rice distribution
- The shifted Gompertz distribution
- The type-2 Gumbel distribution
- The Weibull distribution or Rosin Rammler grinding, milling and crushing operations.

Supported on the whole real line [edit sour

- The Behrens–Fisher distribution, which ari
- The Cauchy distribution, an example of a resonance energy distribution, impact and
- Chernoff's distribution
- The Exponentially modified Gaussian distri
- The Fisher–Tippett, extreme value, or log-
  - The Gumbel distribution, a special case
- Fisher's z-distribution
- The generalized logistic distribution
- The generalized normal distribution
- The geometric stable distribution
- The Holtsmark distribution, an example of
- The hyperbolic distribution
- The hyperbolic secant distribution
- The Johnson SU distribution
- The Landau distribution
- The Laplace distribution
- The Lévy skew alpha-stable distribution or distribution, Lévy distribution and normal d
- The Linnik distribution
- The logistic distribution
- The map-Airy distribution
- The normal distribution, also called the Ga independent, identically distributed variabl
- The Normal-exponential-gamma distributior
- The Pearson Type IV distribution (see Pea
- The skew normal distribution

- Student's t-distribution, useful for estimating u
  - The noncentral t-distribution
- The type-1 Gumbel distribution
- The Voigt distribution, or Voigt profile, is the c
- The Gaussian minus exponential distribution is

With variable support [edit source | edit beta]

- The generalized extreme value distribution has parameter
- The generalized Pareto distribution has a sup
- The Tukey lambda distribution is either suppor
- The Wakeby distribution

Mixed discrete/continuous distributions [edi

- The rectified Gaussian distribution replaces n

Joint distributions [edit source | edit beta]

For any set of independent random variables the

Two or more random variables on the same sar

- The Dirichlet distribution, a generalization of t
- The Ewens's sampling formula is a probability
- The Balding–Nichols model
- The multinomial distribution, a generalization c
- The multivariate normal distribution, a generali
- The negative multinomial distribution, a genera
- The generalized multivariate log-gamma distrib

Matrix-valued distributions [edit source | edit b

- The Wishart distribution
- The inverse-Wishart distribution
- The matrix normal distribution
- The matrix t-distribution

Non-numeric distributions [edit source | edit

- The categorical distribution

newton distribution

Miscellaneous distributions [edit source | edit

- The Cantor distribution
- The generalized logistic distribution family
- The Pearson distribution family
- The phase-type distribution

And surely more!

Probability Density Functions

The number of PDFs is infinite, and nearly so is the list of known ones:

Discrete distributions [edit source | edit beta]

With finite support [edit source | edit beta]

- The Bernoulli distribution, which takes value 1 with
- The Rademacher distribution, which takes value 1
- The binomial distribution, which describes the numb
- The beta-binomial distribution, which describes the
- The degenerate distribution at x_0, where X is certa random variables in the same formalism.
- The discrete uniform distribution, where all element shuffled deck.
- The hypergeometric distribution, which describes t there is no replacement.
- The Poisson binomial distribution, which describes
- Fisher's noncentral hypergeometric distribution
- Wallenius' noncentral hypergeometric distribution

Continuous distributions [edit source | edit beta]

Supported on a bounded interval [edit source | edit

- The Arcsine distribution on [a,b], which is a speci
- The Beta distribution on [0,1], of which the uniforr
- The Logitnormal distribution on (0,1).
- The Dirac delta function although not strictly a fur but the notation treats it as if it were a continuous
- The continuous uniform distribution on $[a,b]$, wher
  - The rectangular distribution is a uniform distrib
- The Irwin-Hall distribution is the distribution of the
- The Kent distribution on the three-dimensional sph
- The Kumaraswamy distribution is as versatile as t
- The logarithmic distribution (continuous)
- The PERT distribution is a special case of the bet
- The raised cosine distribution on $[\mu - s, \mu + s]$

- Hotelling's T-squared distribution
- The inverse Gaussian distribution, also kn
- The Lévy distribution
- The log-Cauchy distribution
- The log-gamma distribution
- The log-Laplace distribution
- The log-logistic distribution
- The log-normal distribution, describing vari
- The Mittag–Leffler distribution
- The Nakagami distribution
- The Pareto distribution, or "power law" dist
- The Pearson Type III distribution
- The phased bi-exponential distribution is c
- The phased bi-Weibull distribution
- The Rayleigh distribution
- The Rayleigh mixture distribution
- The Rice distribution

- Student's t-distribution, useful for estimating u
  - The noncentral t-distribution
- The type-1 Gumbel distribution
- The Voigt distribution, or Voigt profile, is the c
- The Gaussian minus exponential distribution is

With variable support [edit source | edit beta]

- The generalized extreme value distribution has parameter
- The generalized Pareto distribution has a sup
- The Tukey lambda distribution is either suppor
- The Wakeby distribution

Mixed discrete/continuous distributions [edi

- The rectified Gaussian distribution replaces ne

Joint distributions [edit source | edit beta]

For any set of independent random variables the

https://docs.scipy.org/doc/scipy/reference/stats.html

- The Gibbs distribution
- The Maxwell–Boltzmann distribution
- The Borel distribution
- The extended negative binomial distribution
- The extended hypergeometric distribution
- The generalized log-series distribution
- The generalized normal distribution
- The geometric distribution, a discrete distribution w
- The hypergeometric distribution
- The logarithmic (series) distribution
- The negative binomial distribution or Pascal distribu
- The parabolic fractal distribution
- The Poisson distribution, which describes a very la Poisson, the hyper-Poisson, the general Poisson b
  - The Conway–Maxwell–Poisson distribution, a tw
- The Polya-Eggenberger distribution
- The Skellam distribution, the distribution of the diff
- The skew elliptical distribution
- The skew normal distribution
- The Yule–Simon distribution
- The zeta distribution has uses in applied statistics
- Zipf's law or the Zipf distribution. A discrete power-
- The Zipf–Mandelbrot law is a discrete power law dis

Supported on semi-infinite intervals, usually [0,∞)

- The Beta prime distribution
- The Birnbaum–Saunders distribution, also known a
- The chi distribution
  - The noncentral chi distribution
- The chi-squared distribution, which is the sum of t
  - The inverse-chi-squared distribution
  - The noncentral chi-squared distribution
  - The Scaled-inverse-chi-squared distribution
- The Dagum distribution
- The exponential distribution, which describes the t
- The F-distribution, which is the distribution of the ratio of two chi-squared variates which are not no
  - The noncentral F-distribution
- Fisher's z-distribution
- The folded normal distribution
- The Fréchet distribution
- The Gamma distribution, which describes the time
  - The Erlang distribution, which is a special case
  - The inverse-gamma distribution
- The generalized Pareto distribution
- The Gamma/Gompertz distribution
- The Gompertz distribution
- The half-normal distribution

- resonance energy distribution, impact and
- Chernoff's distribution
- The Exponentially modified Gaussian distri
- The Fisher–Tippett, extreme value, or log-
  - The Gumbel distribution, a special case
- Fisher's z-distribution
- The generalized logistic distribution
- The generalized normal distribution
- The geometric stable distribution
- The Holtsmark distribution, an example of
- The hyperbolic distribution
- The hyperbolic secant distribution
- The Johnson SU distribution
- The Landau distribution
- The Laplace distribution
- The Lévy skew alpha-stable distribution or distribution, Lévy distribution and normal d
- The Linnik distribution
- The logistic distribution
- The map-Airy distribution
- The normal distribution, also called the Ga independent, identically distributed variabl
- The Normal-exponential-gamma distributior
- The Pearson Type IV distribution (see Pea
- The skew normal distribution

- The generalized multivariate log-gamma distrib

Matrix-valued distributions [edit source | edit

- The Wishart distribution
- The inverse-Wishart distribution
- The matrix normal distribution
- The matrix t-distribution

Non-numeric distributions [edit source | edit

- The categorical distribution

newton distribution

Miscellaneous distributions [edit source | edit

- The Cantor distribution
- The generalized logistic distribution family
- The Pearson distribution family
- The phase-type distribution

And surely more!

Probability Density Functions

The number of PDFs is infinite, and nearly so is the list of known ones:

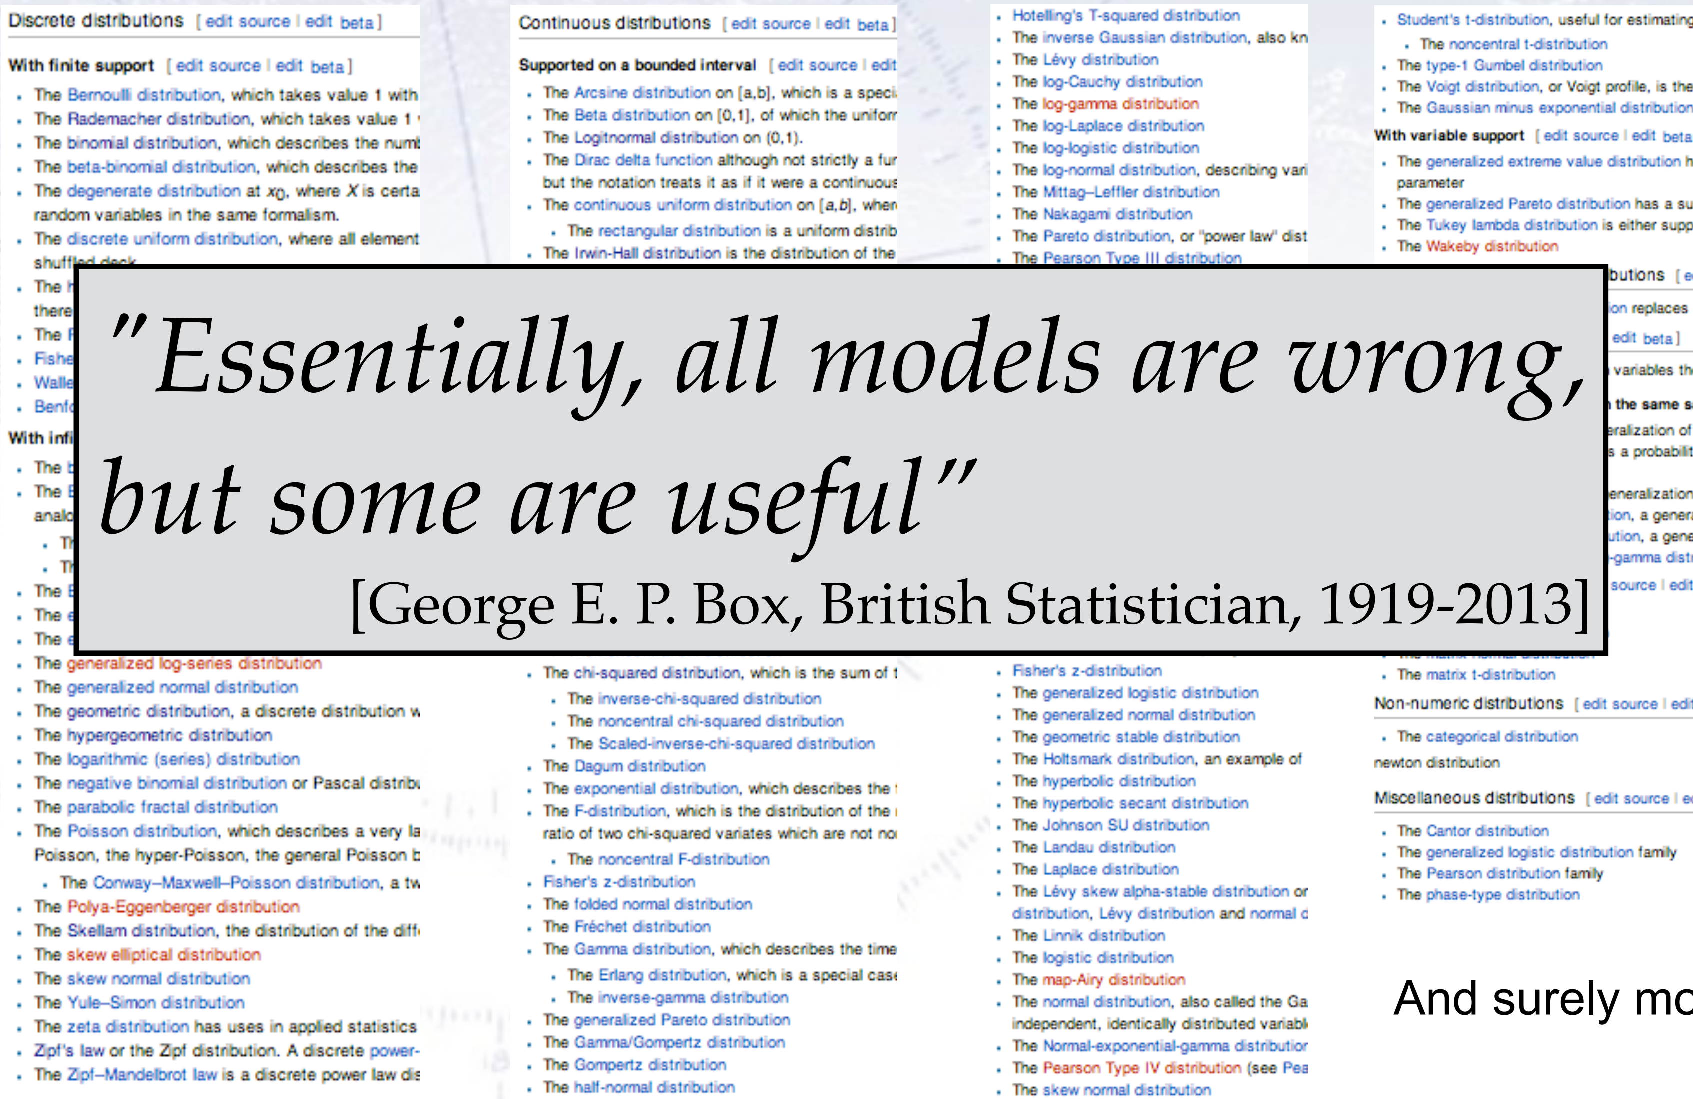

"Essentially, all models are wrong, but some are useful"

[George E. P. Box, British Statistician, 1919-2013]

And surely more!

Why PDF's

We basically want to extract information about the system we are studying based on the data. Why not just estimate a mean and a variance and call it a day?

PDFs makes us able to ask what the *probability* of a certain event given the underlying model - and this allows for new discoveries!

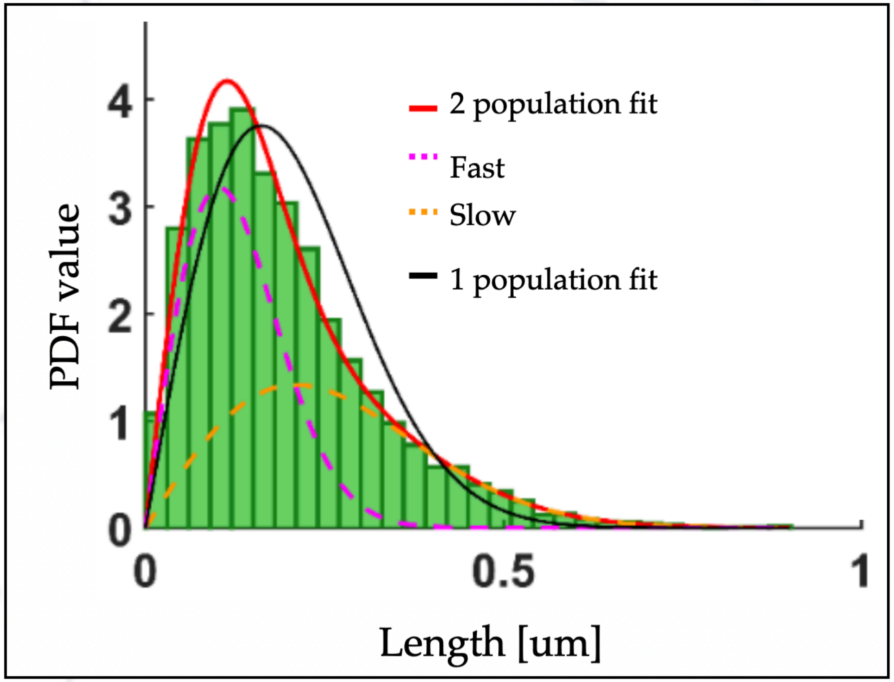

Probability Density Functions

An almost complete list of those we will deal with in this course is:

- **Gaussian** (aka. Normal)
- **Poisson**
- **Binomial** (and also Multinomial)
- Students t-distribution
- Uniform
- ChiSquare
- Exponential
- Error function (integral of Gaussian)

See Barlow chap.3 and Cowan chap.2

You should already know most of these, and the rest will be explained.

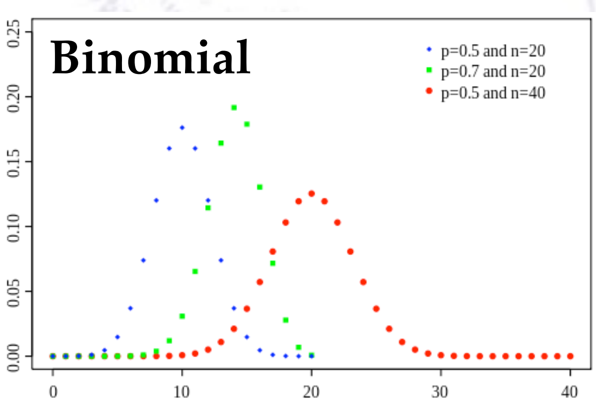

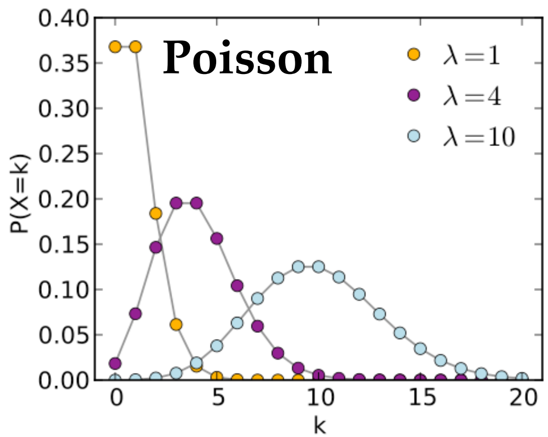

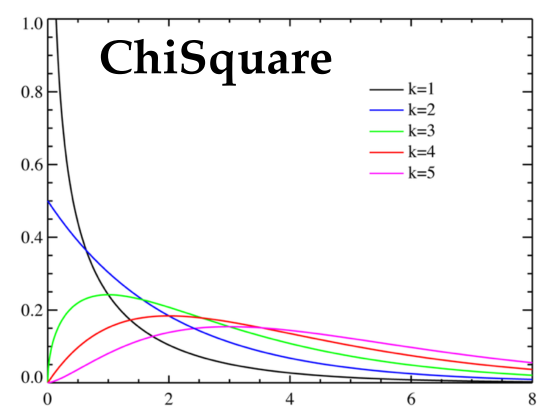

Binomial, Poisson, Gaussian

$$f(n;N,p) = \frac{N!}{n!(N-n)!}p^n(1-p)^{N-n}$$

Given **N trials** each with **p chance of success**, how many **successes n** should you expect in total?

This distribution is… **Binomial:**

$$f(n;N,p) = \frac{N!}{n!(N-n)!}p^n(1-p)^{N-n}$$

Mean = Np
Variance = Np(1-p)

This means, that the error on a fraction f = n/N is:

$$\sigma(f) = \sqrt{\frac{f(1-f)}{N}}$$

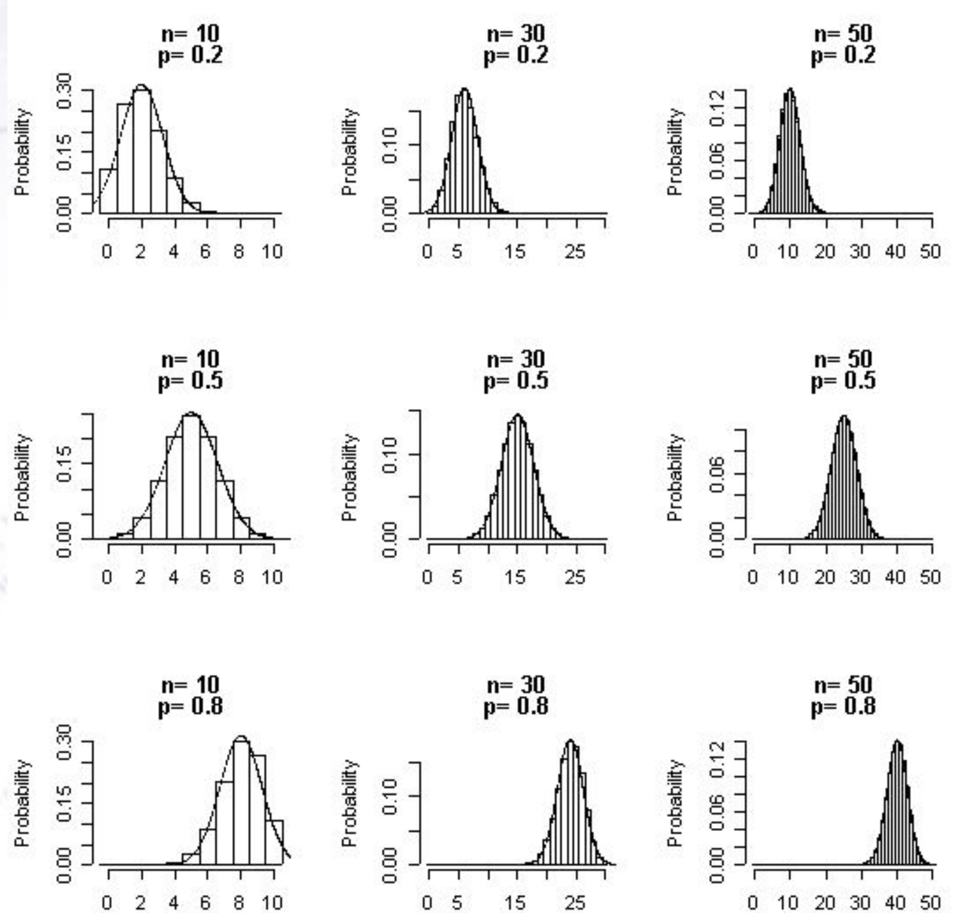

Binomial, Poisson, Gaussian

You count 100 cars, and see 15 red cars. What is your estimate of the fraction (i.e. probability) of red cars and its uncertainty?

a) 0.150 ± 0.050
b) 0.150 ± 0.026
c) 0.150 ± 0.036
d) 0.125 ± 0.030
e) 0.150 ± 0.081

From previous page: $$\sigma(f) = \sqrt{\frac{f(1-f)}{N}}$$

A friend tells you, that 8% of the cars on Blegdamsvej are red. What is the chance of that? Could he be right?

Binomial, Poisson, Gaussian

You count 100 cars, and see 15 red cars. What is your estimate of the fraction (i.e. probability) of red cars and its uncertainty?

a) 0.150 ± 0.050
b) 0.150 ± 0.026
c) 0.150 ± 0.036 ← (0.150 - 0.080) / 0.036 = 1.9 σ
d) 0.125 ± 0.030
e) 0.150 ± 0.081

From previous page: $$\sigma(f) = \sqrt{\frac{f(1-f)}{N}}$$

A friend tells you, that 8% of the cars on Blegdamsvej are red. What is the chance of that? Could he be right?

Binomial, Poisson, Gaussian

$$f(n;N,p) = \frac{N!}{n!(N-n)!}p^n(1-p)^{N-n}$$

$$(x+y)^4 = x^4 + 4x^3y + 6x^2y^2 + 4xy^3 + y^4.$$

0:	1								
1:	1	1							
2:	1	2	1						
3:	1	3	3	1					
4:	1	4	6	4	1				
5:	1	5	10	10	5	1			
6:	1	6	15	20	15	6	1		
7:	1	7	21	35	35	21	7	1	
8:	1	8	28	56	70	56	28	8	1

The binomial distribution was first introduced by the famous mathematician/physicist Jacob Bernoulli in 1713 (posthumously).

The binomial distribution basically consists of two elements: The binomial coefficient and the probabilities of n such events.

Even though a system has many outcomes, it is typically possible to refer to either "success" of "failure".

Assume the probability to have COVID19 is 1%. In a sample of 50 people the chance to have 1 or more infected is: 1-p(0) = 1-0.99^50 = 0.60

Binomial, Poisson, Gaussian

Requirements to be Binomial:

- Fixed number of trials, N
- Independent trials.
- Only two outcomes (success/failure).
- Constant probability of success/failure.

If number of possible outcomes is more than two ⇒ **Multinomial distribution**.

Examples of Binomial experiments:

- Tossing a coin 20 times counting number of tails.
- Asking 200 people if they watch sports on TV.
- Rolling a die to see if a 6 appear (Multinomial for all outcomes!).
- Asking 100 die-hards from Enhedslisten, if they would vote for Konservative at the next election!

Examples which aren't Binomial experiments:

- Rolling a die until a 6 appears (not fixed number of trials).
- Drawing 5 cards for a poker hand (no replacement ⇒ not independent)

Binomial, **Poisson**, Gaussian

If N → ∞ and p → 0, but Np → λ then a Binomial approaches a Poisson: (see Barlow 3.3.1)

$$f(n,\lambda) = \frac{\lambda^n}{n!}e^{-\lambda}$$

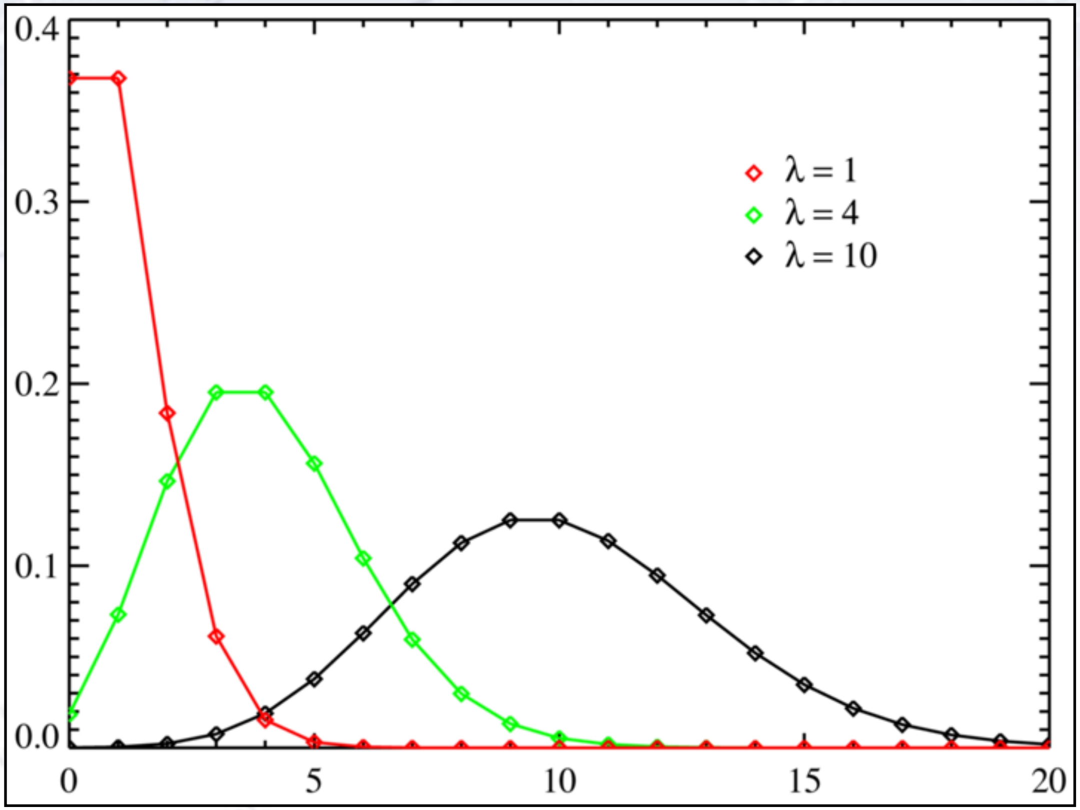

In reality, the approximation is already quite good at e.g. N=50 and p=0.1.

The Poisson distribution only has one parameter, namely λ.
Mean = λ
Variance = λ

So the error on a number is...

...the square root of that number!

Binomial, **Poisson**, Gaussian

The error on a (Poisson) number... is **the square root** of that number!!!

Binomial, **Poisson**, Gaussian

The error on a (Poisson) number... is **the square root** of that number!!!

A very useful case of this is the error to assign a bin in a histogram, if there is reasonable statistics (N_i > 5-20) in each bin.

The error on a (Poisson) number... is **the square root** of that number!!!

Note: The sum of two Poissons with λ_a and λ_b is a new Poisson with $\lambda = \lambda_a + \lambda_b$.
(See Barlow pages 33-34 for proof)

Binomial, **Poisson**, Gaussian

Examples of the Poisson distribution:

There are 122 deaths by horse kicks over 10 different regiments, over 20 years. What is the predicted number of deaths in a specific year?

Here we use exactly the Poisson distribution. First we estimate the mean value:

$$\mu = \frac{122}{20 * 10} = 0.61$$

This means that the probability that 0 will die is given by:

$$P(0) = e^{0}\frac{0.61^{0}}{0!} = 0.54$$

Binomial, Poisson, **Gaussian**

If $\lambda \rightarrow \infty$, the Poisson becomes a Gaussian… ...and $\lambda > 20$ is enough!

For proof, see Barlow p.40

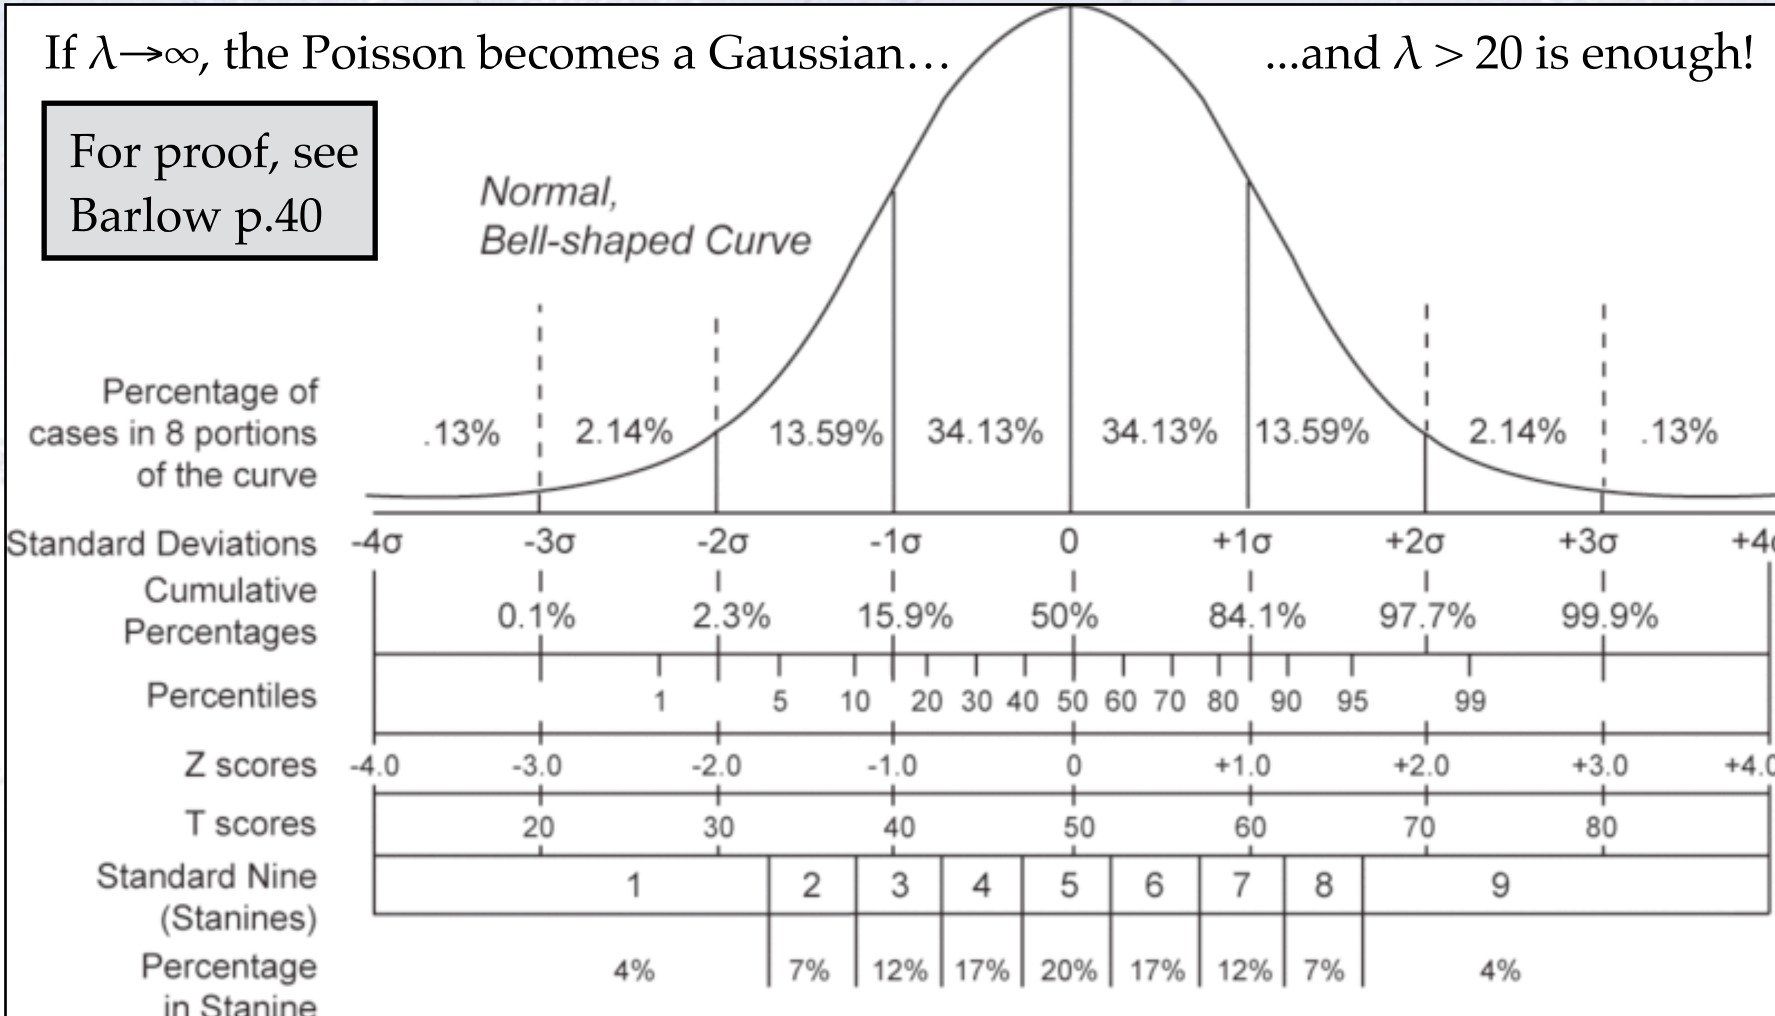

All fields encounter the Gaussian, and for this reason, its scale has many names!

Binomial, Poisson, **Gaussian**

If $\lambda \to \infty$, the Poisson becomes a Gaussian… ...and $\lambda > 20$ is enough!

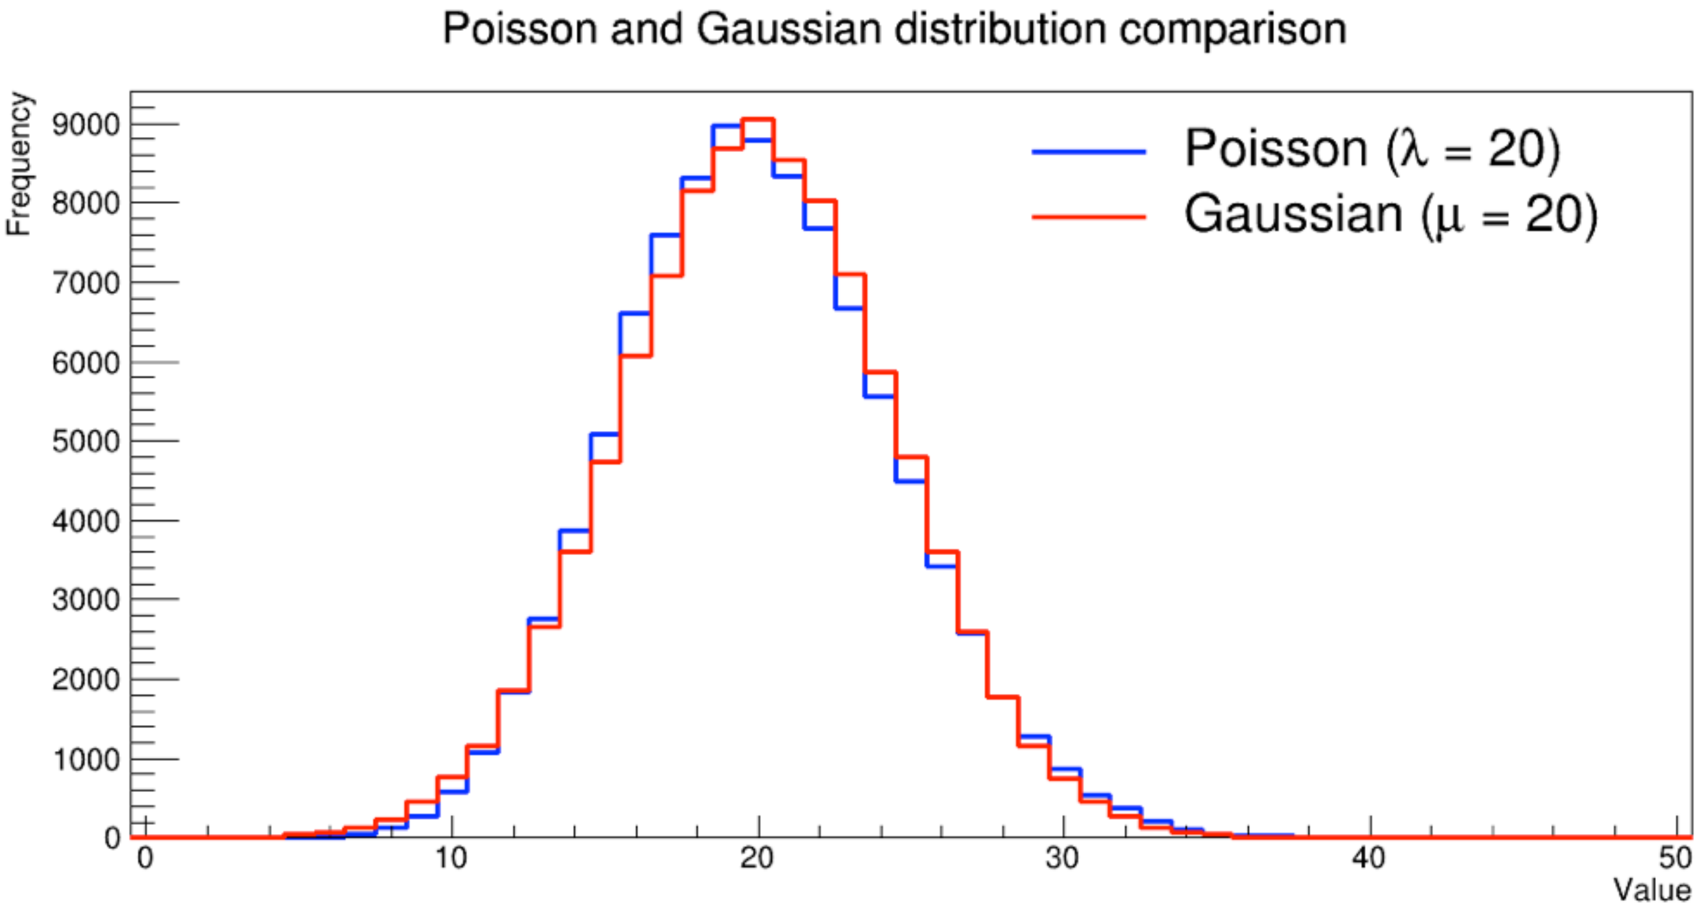

Binomial, Poisson, **Gaussian**

If $\lambda \to \infty$, the Poisson becomes a Gaussian… ...and $\lambda > 20$ is enough!

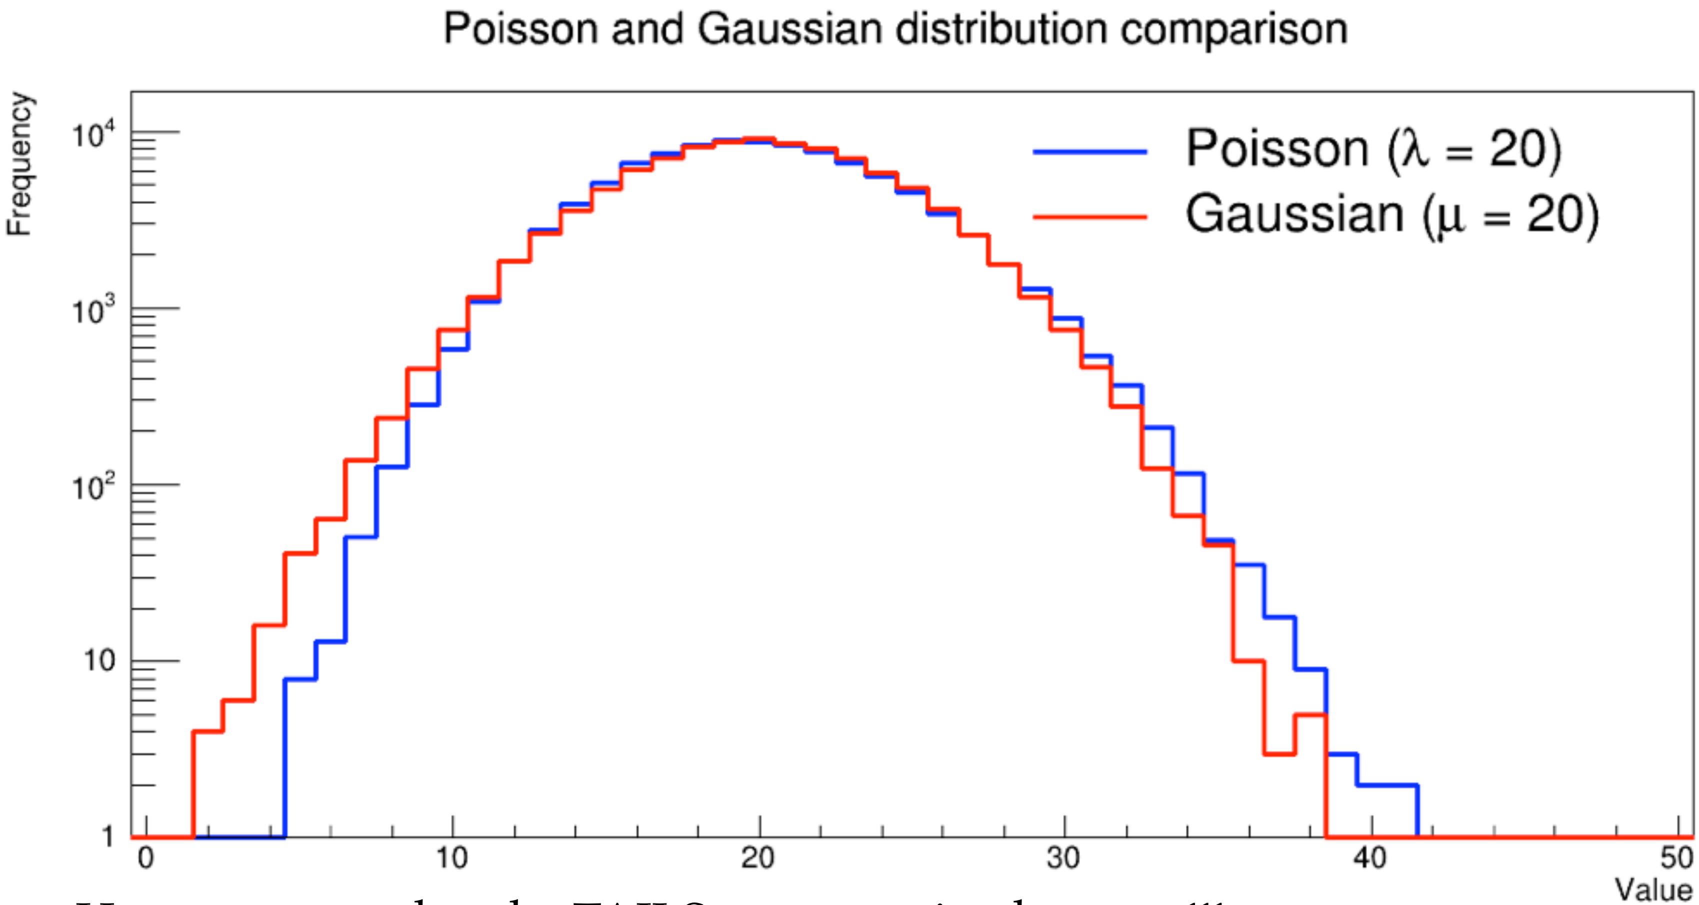

However, note that the TAILS are not quite the same!!!
This is the very reason for the difference between Chi2 and (binned) likelihood!

Binomial, Poisson, **Gaussian**

"If the Greeks had known it, they would have deified it."

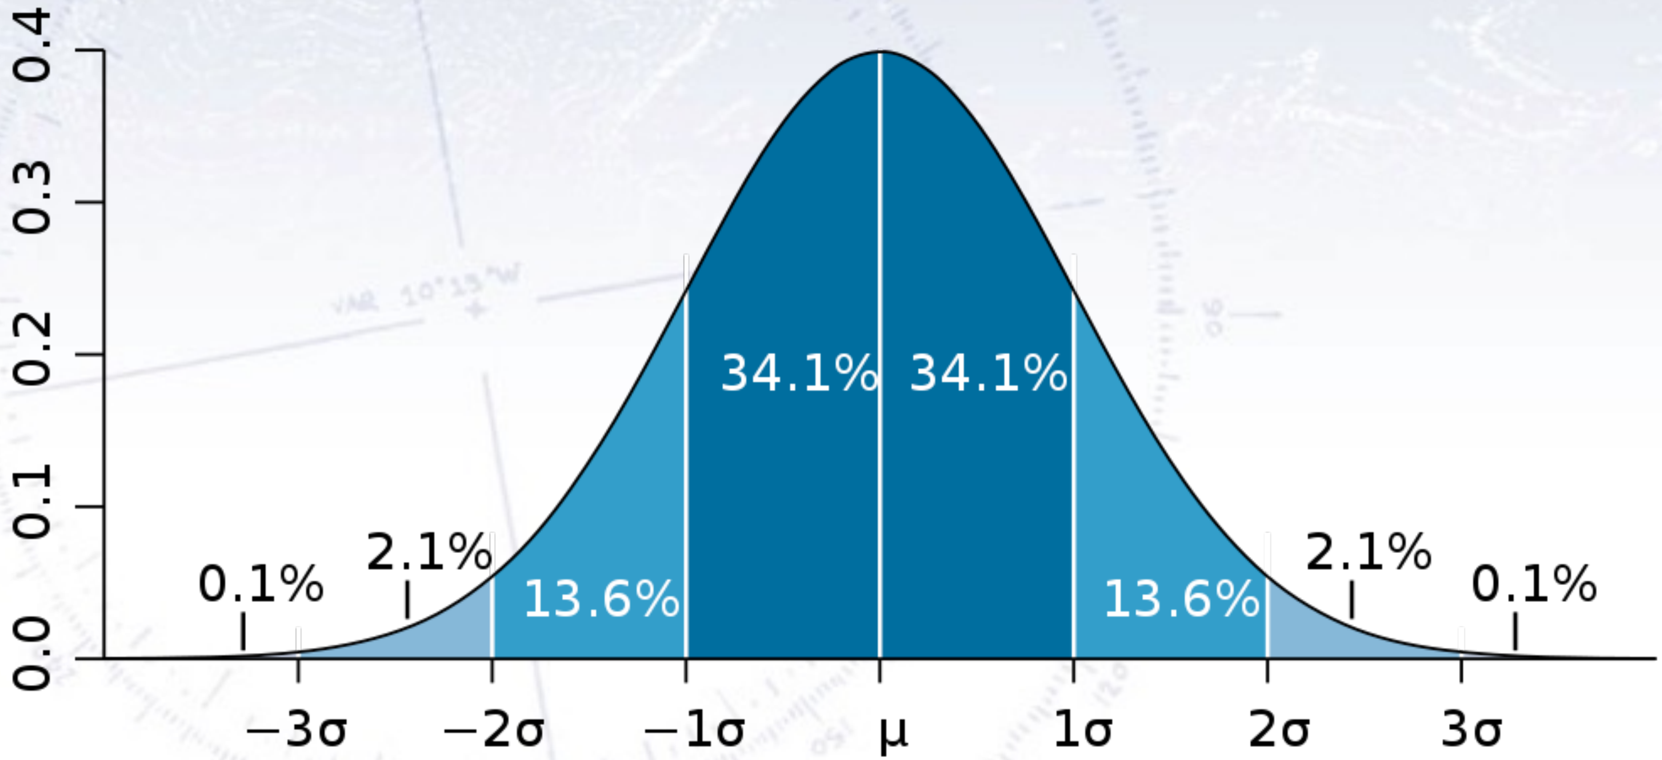

"If the Greeks had known it, they would have deified it. It reigns with serenity and in complete self-effacement amids the wildest confusion. ***The more huge the mob and the greater the apparent anarchy, the more perfect is its sway****. It is the supreme Law of Unreason. Whenever a large sample of chaotic elements are taken in hand and marshalled in the order of their magnitude, an unsuspected and most beautiful form of regularity proves to be latent all along." [Karl Pearson]*

Binomial, Poisson, **Gaussian**

The Gaussian **defines** the way we consider uncertainties.

Range	Inside	Outside
$\pm 1\sigma$	**68** %	32 %
$\pm 2\sigma$	**95** %	5 %
$\pm 3\sigma$	**99.7** %	0.3 %
$\pm 5\sigma$	99.99995 %	0.00005 %

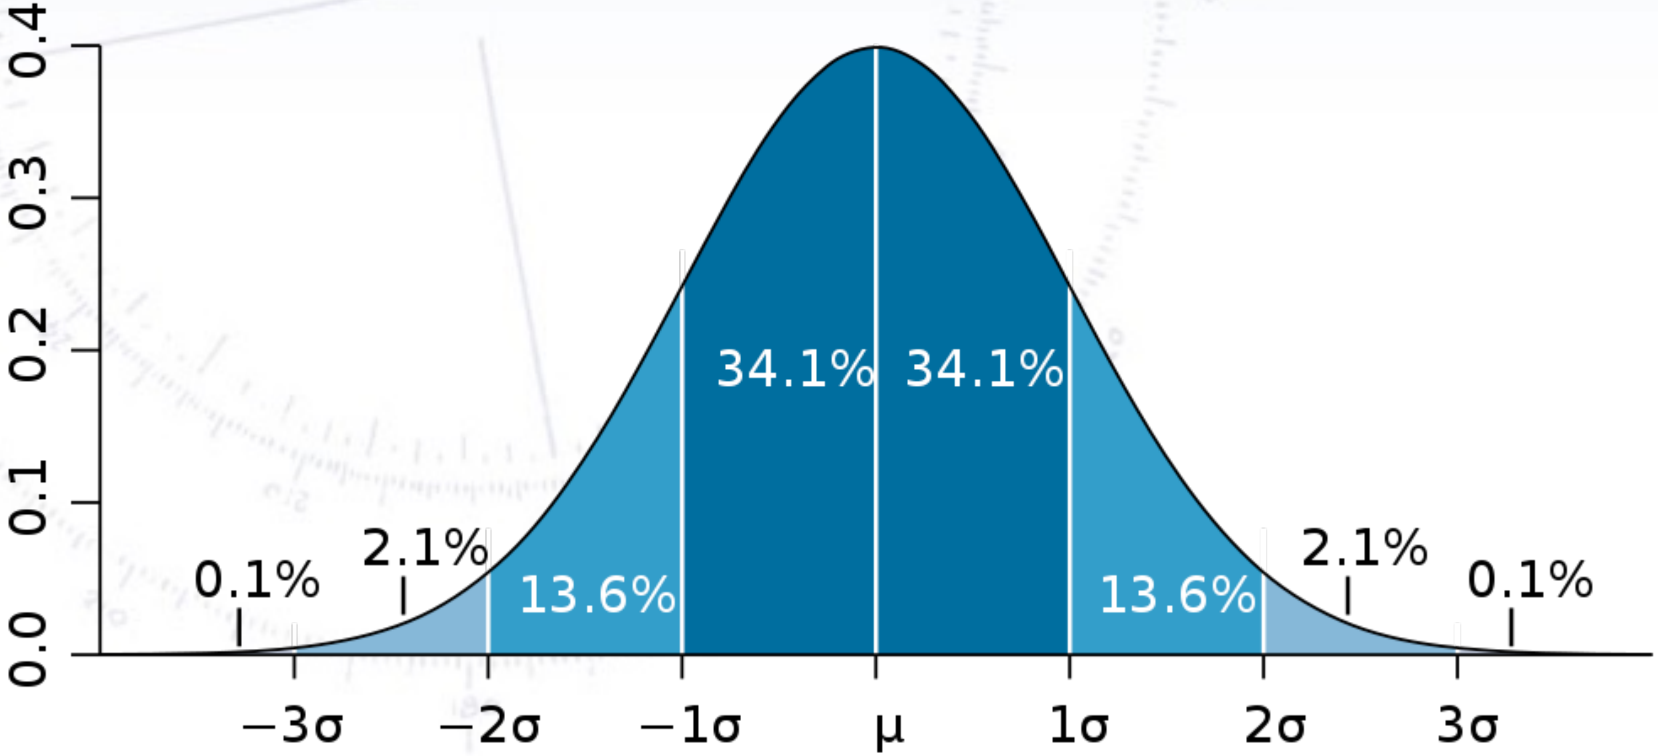

Student's t-distribution

Assume we are measuring a small sample, i.e. 5 measurements.

Now we don't even know the uncertainty on the measurements, but we assume that they are normally distributed. We can rewrite our variable of interest:

$$t = \frac{x - \mu}{\hat{\sigma}^2} = \frac{y}{\sqrt{\chi^2/n}}$$

This distribution we have seen before:

$$f(x) = \begin{cases} \frac{x^{k/2-1} e^{-x/2}}{2^{k/2}\Gamma(k/2)} & \text{for } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

When variance is unknown, estimating it from sample gives additional error:

Gaussian: $$z = \frac{x - \mu}{\sigma}$$

Student's: $$t = \frac{x - \mu}{\hat{\sigma}}$$

Student's t-distribution

Discovered by William Gosset (who signed "student"), student's t-distribution takes into account **lacking knowledge of the variance**.

$$\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\,\Gamma\left(\frac{\nu}{2}\right)}\left(1+\frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$$

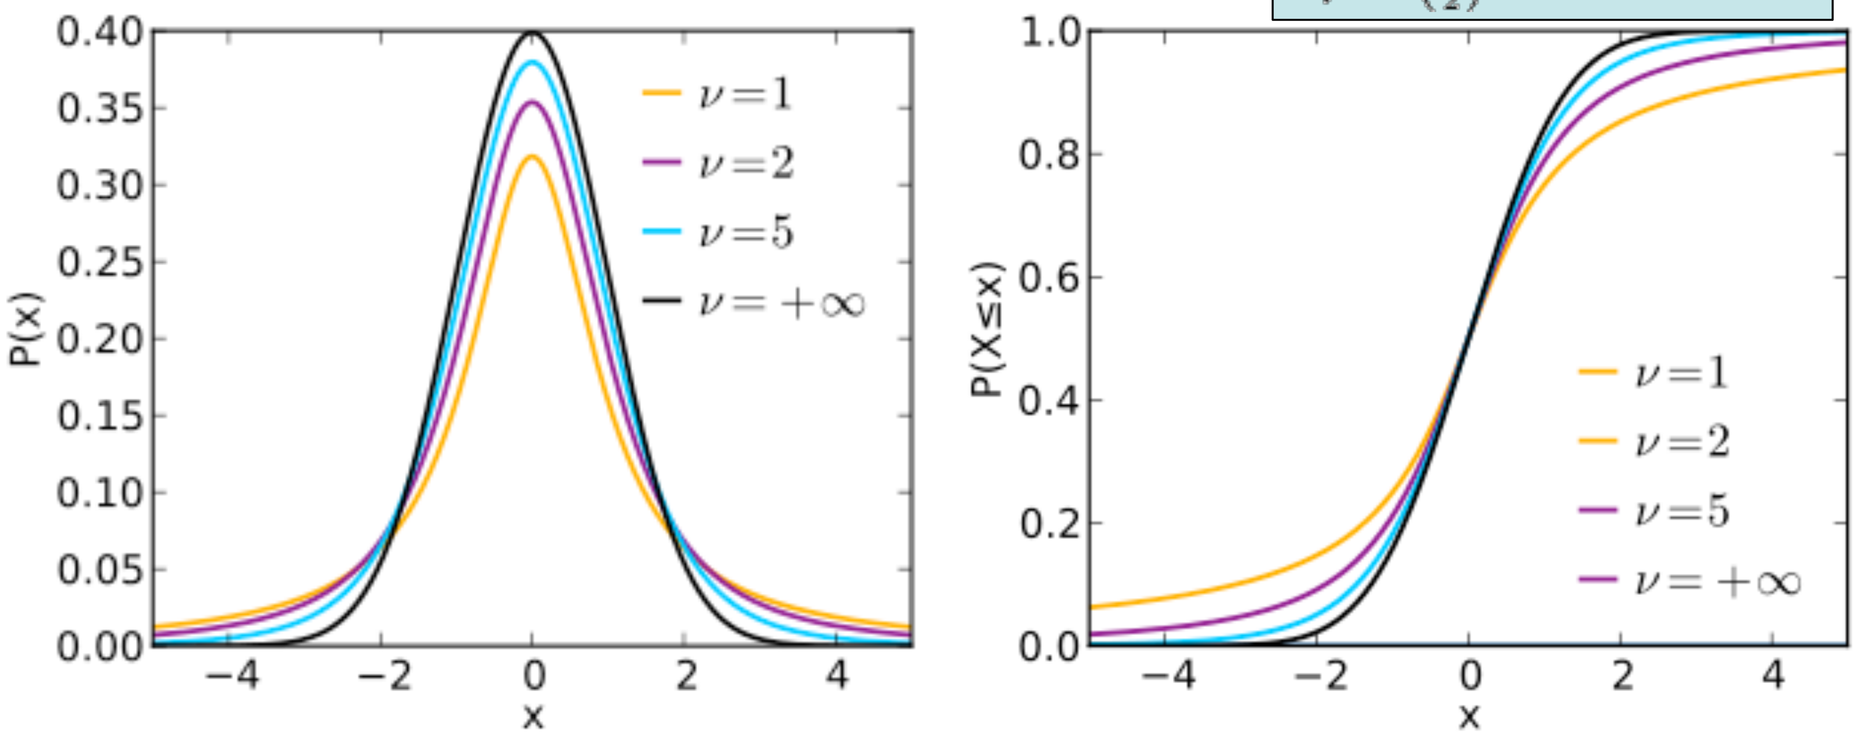

When variance is unknown, estimating it from sample gives additional error:

Gaussian: $$z = \frac{x-\mu}{\sigma}$$

Student's: $$t = \frac{x-\mu}{\hat{\sigma}}$$

Exponential distribution

One particularly important PDF is the exponential distribution.
This takes the shape:

$$f(x) = \lambda e^{-\lambda x}$$

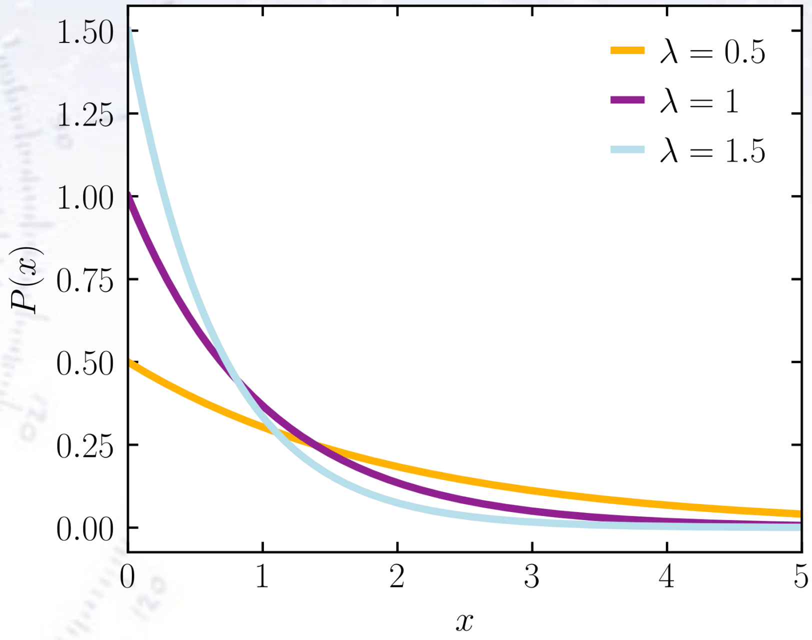

With this we can also calculate the mean value:

$$\begin{aligned}\mu &= \int_0^\infty x\lambda e^{-\lambda x} dx \\ &= \lambda\Big([-\frac{1}{\lambda} x e^{-\lambda x}]_0^\infty + \int_0^\infty \frac{1}{\lambda} e^{-\lambda x} dx\Big) \\ &= \frac{\lambda}{\lambda^2} = \frac{1}{\lambda}\end{aligned}$$

And also the variance by performing similar integration:

$$\begin{aligned}\sigma^2 &= \int_0^\infty (x-\mu)^2 \lambda e^{-\lambda x} dx \\ &= \frac{1}{\lambda^2}\end{aligned}$$

Exponential distribution

Why is the exponential distribution so important?

There are several good examples, but the one we will show here is the following:

If a process occurs with a constant rate, the distribution of time before the event happens will be exponentially distributed

This is really the case for many systems. Of course the most prominent example is the decay of atoms.

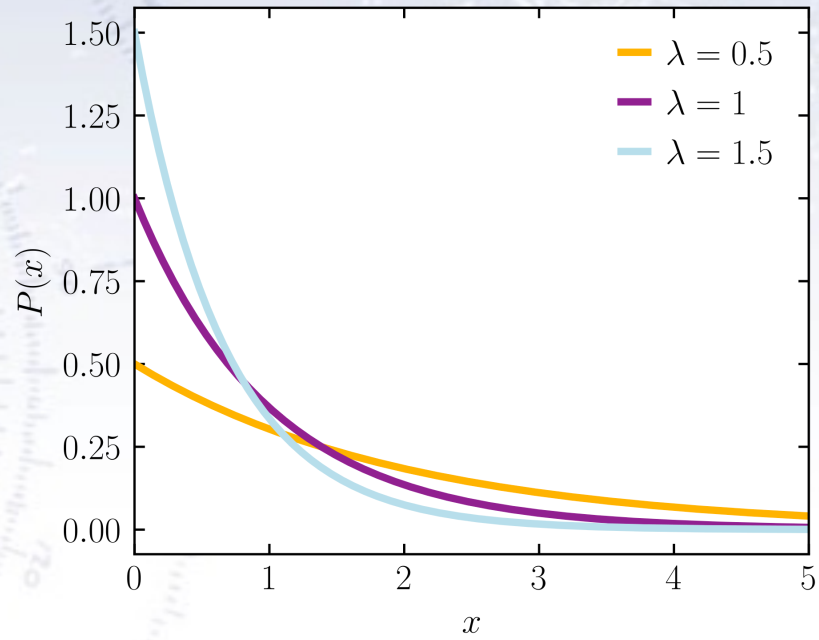

Exponential distribution

Assume that an atom decays with rate λ.

The probability that the atom will decay in a short time interval dt, must be:

$$p_{decay} = \lambda dt \quad p_{not} = (1-\lambda)dt$$

Now we can define a time that is the sum of a lot of small time intervals: $\tau = n \cdot dt$

$$p_{not}(\tau) = (1 - \frac{\tau}{n}\lambda))^n$$

And finally, the probability that an event occurs after exactly some time tau must be:

$$p(t_{first} = \tau) = (1 - \frac{\tau}{n}\lambda))^n \lambda dt = \lambda e^{-\lambda\tau} dt$$

Where we have used the definition of the exponential distribution

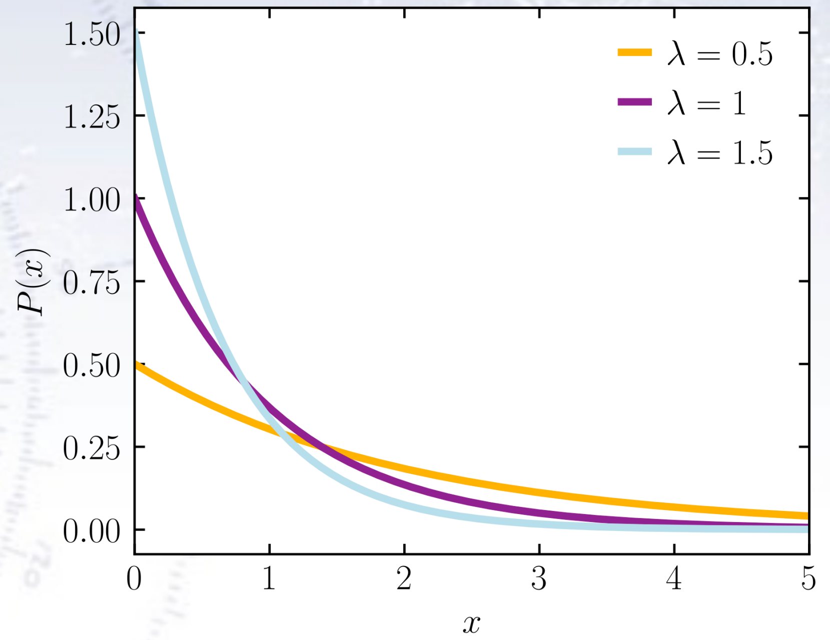

Rayleigh distribution

Let's assume that a particle moves in two dimensions. It diffuses so in both dimensions, it will take the step with distribution (for simplicity assuming std=1):

$$p(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2}$$

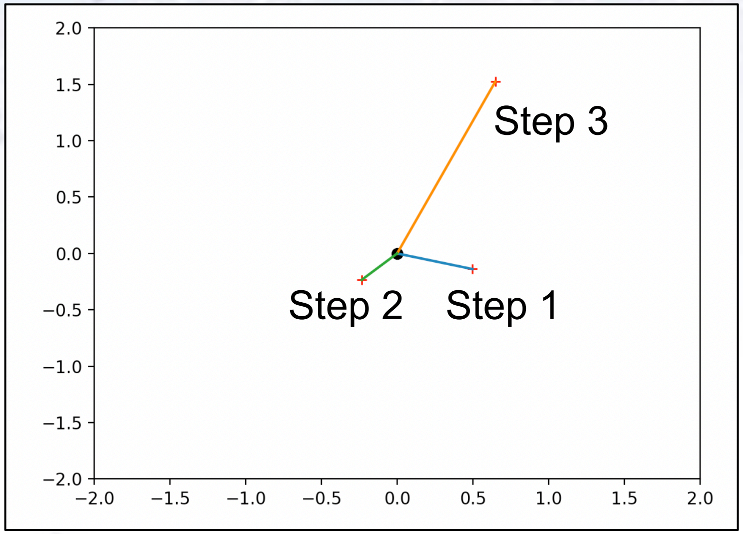

This means that the probability that the particle will be at exactly position x* *and* y* will be:

$$p(x, y) = \frac{1}{2\pi} e^{-\frac{1}{2}(x^2+y^2)}$$

This looks like the radius, but in order to compute to total probability of a radius, we must factor in the possibility of obtaining a particular radius.

Rayleigh distribution

To count the number of ways to obtain the radius r, is exactly the perimeter of a circle with radius r, i.e. $2\pi r$.

Multiplying with this we get:

$$p(r) = re^{-\frac{1}{2}r^2}$$

This is exactly the Rayleigh distribution for std=1.

If we simulate the 1000 of the steps, we see that we such a distribution in the r-coordinate.

Note that the must probable point is x=y=0, but that the probability of r=0 is zero.

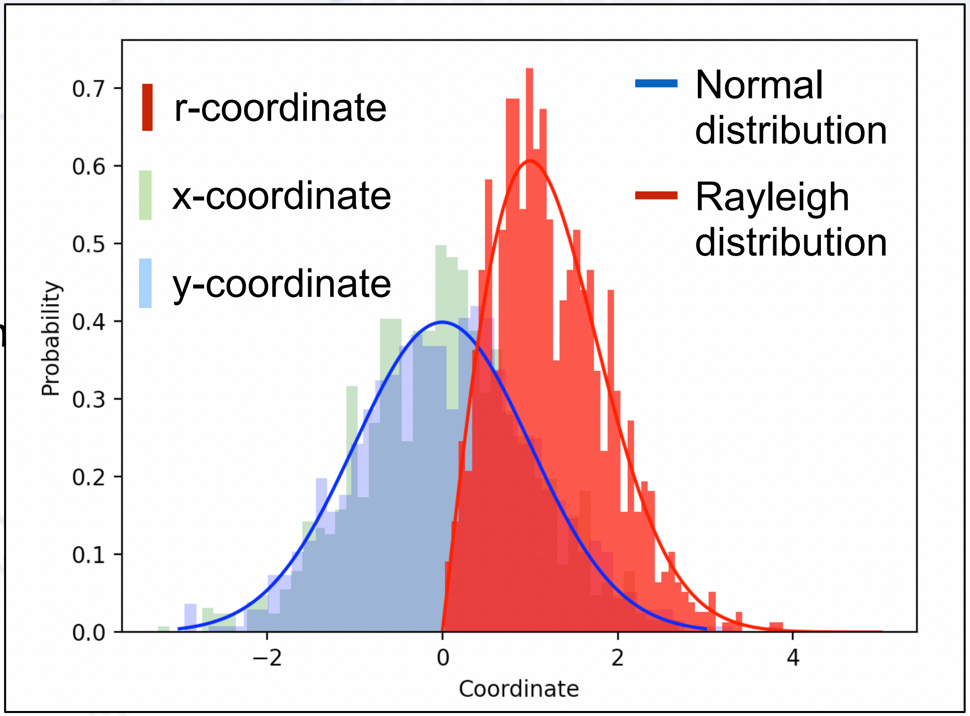

Some particular elements of PDF

Cumulative distributions functions

Completely basic to every PDF is the cumulative distribution function, the CDF.

The CDF is defined by:

$$F_X(x) = \int_{-\infty}^{x} f_X(t)\, dt.$$

In words, this means that it is the probability of getting x, or something below that value.

The CDF is used in many ways, and we will meet it again soon, when we discuss hypothesis testing and statistical tests.

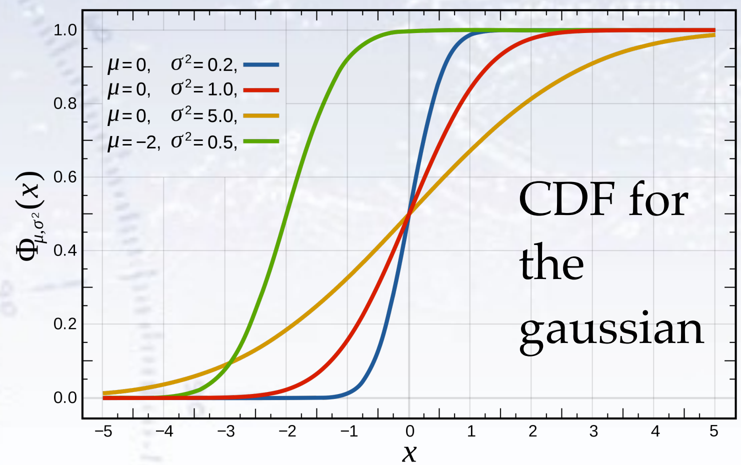

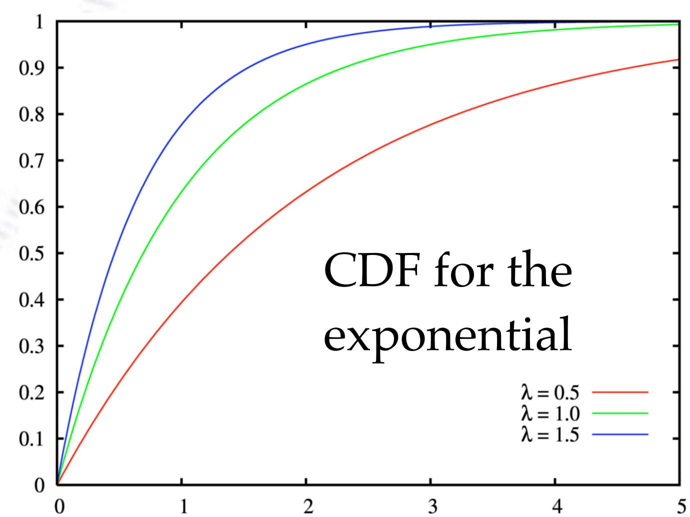

Convolution of probabilities

Suppose we draw two random numbers from a probability distribution f(x). How can we calculate the probability of their sum?

$$
\begin{aligned}
f(z) &= \int\int f(x_1)\cdot f(x_2)\delta(z-ax_1-ax_2)dx_1dx_2 \quad \text{we substitute } s = ax_1 \text{ so } \frac{ds}{dx_1} = a \\
&= \int\int f(\frac{s}{a})\cdot f(x_2)\delta(z-s-ax_2)\frac{1}{a}dsdx_2 \quad \text{we integrate using the } \delta \text{ function} \\
&= \frac{1}{a}\int f(\frac{z}{a}-x_2)\cdot f(x_2)dx_2
\end{aligned}
$$

For the gaussian for instance, we will get, that this is a new gaussian distribution.

Characteristic functions

Typically when studying statistics, one stumble upon the term characteristic function.

Every PDF has a characteristic function that is simply the Fourier transform of the PDF itself:

$$G_x(t) = \int_{-\infty}^{\infty} f(x)e^{itx}dx = \mathbb{E}(e^{ixt})$$

The characteristic function has some properties that makes it neat to use, for instance that it can generate all moments for the PDF itself (note that the first moment is the mean, the second is the variance etc.).

The gaussian has a characteristic function given by:

$$\begin{aligned}
G_z(t) &= \int e^{itz}\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}dz \text{ noting that } (z-it)^2 = z^2 - t^2 - itz \\
&= \int \frac{1}{\sqrt{2\pi}}e^{-\frac{(z-it)^2}{2}}e^{-\frac{t^2}{2}}dz \text{ substituting } u = z - it \\
&= \frac{1}{\sqrt{2\pi}}e^{-\frac{t^2}{2}}\int e^{-\frac{u^2}{2}}dz \text{ noting the integral is } \sqrt{2\pi}\text{(for the gaussian to be normalized)} \\
&= e^{-\frac{t^2}{2}}
\end{aligned}$$

Sum up of the lecture

All PDFs are normalized functions, that describe the probability of getting a value from a function.

The most fundamental are the binomial, the Poisson and the Gaussian.

Remember that the error on a (Poisson) number is **the square root** of that number!

Remember that the Gaussian distribution defines the uncertainties, that we report on experiments.