Normal Random Variable

X is a normal (aka Gaussian) random variable with mean μ and variance σ^2 (written $X \sim \mathcal{N}(\mu, \sigma^2)$) if it has the density:

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

Let's get some intuition for that density...

Is $\mathbb{E}[X] = \mu$?

Yes! Plug in $\mu - k$ and $\mu + k$ and you'll get the same density for every k. The density is symmetric around $\boldsymbol{\mu}$. The expectation must be μ.

Standardize

To turn $\mathrm{X} \sim \mathcal{N}(\mu, \sigma^2)$ into $\mathrm{Y} \sim \mathcal{N}(0,1)$ you want to set $Y = \frac{X-\mu}{\sigma}$

Why standardize?

The density is a mess. The CDF does not have a pretty closed form.

But we're going to need the CDF a lot, so...