STAT 712 fa 2022 Lec 6 slides

Suite of ought-to-know probability distributions

Karl B. Gregory

University of South Carolina

These slides are an instructional aid; their sole purpose is to display, during the lecture, definitions, plots, results, etc. which take too much time to write by hand on the blackboard. They are not intended to explain or expound on any material.

Distributions related to Bernoulli trials

Let $X_i = \begin{cases} 1 & \text{success} \\ 0 & \text{failure} \end{cases}$ $\quad i = 1, 2, \dots$ indicate outcomes of indep. Bernoulli trials.

1. $Y = X_1 =$ success of single trial $\implies Y \sim \text{Bernoulli}(p)$,

$$p_Y(y) = p^y(1-p)^{1-y} \cdot \mathbf{1}(y \in \{0, 1\})$$

2. $Y = \sum_{i=1}^n X_i = \#$ successes in n trials $\implies Y \sim \text{Binomial}(n, p)$,

$$p_Y(y) = \binom{n}{y} p^y(1-p)^{n-y} \cdot \mathbf{1}(y \in \{0, 1, 2, \dots, n\})$$

3. $Y = \min\{i : X_i = 1\} = \#$ of trial of 1st success $\implies Y \sim \text{Geometric(p)}$,

$$p_Y(y) = (1-p)^{y-1}p \cdot \mathbf{1}(y \in \{1, 2, \dots\})$$

4. $Y = \min\{i : \sum_{j=1}^i X_j = r\} = \#$ of trial of rth succ. $\implies Y \sim \text{negBin}(r, p)$,

$$p_Y(y) = \binom{y-1}{r-1}(1-p)^{y-r}p^r \cdot \mathbf{1}(y \in \{r, r+1, r+2, \dots\})$$

Exercise: Derive these pmfs; then find mean, variance, and mgf of each.

Hypergeometric distribution

Draw $K \geq 0$ marbles from a bag of $N \geq 0$ marbles, of which $M \geq 0$ are red.

Then $X = \#$ red marbles drawn $\implies X \sim \text{Hypergeometric}(N, M, K)$, has pmf

$$p_X(x) = \frac{\binom{M}{x}\binom{N-M}{K-x}}{\binom{N}{K}} \cdot \mathbf{1}(x \in \{\max(K-(N-M), 0), \dots, \min(K, M)\}).$$

Exercise: For $X \sim \text{Hypergeometric}(N, M, K)$, find $\mathbb{E}X$ and $\text{Var}\, X$.

Discrete uniform distribution and empirical distribution

1. If X takes the values $1, \dots, K$ with prob. $1/K$ then $X \sim \text{discUnif}(K)$.

$$p_X(x) = \frac{1}{K} \cdot \mathbf{1}(x \in \{1, \dots, K\})$$

2. If X takes the values $x_1, \dots, x_n$ with prob. $1/n$ then $X \sim \text{empDist}(x_1, \dots, x_n)$.

$$p_X(x) = \frac{1}{n} \cdot \mathbf{1}(x \in \{x_1, \dots, x_n\})$$

Exercise: Give the mean and variance of $X \sim \text{empDist}(x_1, \dots, x_n)$.

Poisson distribution

For some $\lambda > 0$, let $Y_n \sim \text{Binomial}(n, \lambda/n)$ for $n = 1, 2, \ldots$ Then

$$\lim_{n\to\infty} M_{Y_n}(t) = M_Y(t) = e^{\lambda(e^t - 1)} \quad \text{for all } t \in \mathbb{R},$$

where $M_Y(t)$ is the mgf of $Y \sim \text{Poisson}(\lambda)$, which has pmf

$$p_Y(y) = \frac{e^{-\lambda}\lambda^y}{y!} \cdot \mathbf{1}(y \in \{0, 1, 2, \ldots\}).$$

So for large n, $Y_n \sim \text{Binomial}(n, \lambda/n)$ behaves like $Y \sim \text{Poisson}(\lambda)$.

Often posited for # of occurrences of an event per unit of time/space, where

- the events occur independently from one another, and
- are as likely to occur in any time/space interval as in any other.

Exercise: Find mean and variance and derive mgf.

- The pdf of the Uniform(a, b) distribution is given by

$$f_X(x; a, b) = \frac{1}{b-a}\mathbf{1}(a < x < b) \quad \text{for } x \in \mathbb{R},$$

  for $a < b$.
- Parameters:
  - a is the lower bound of the support
  - b is the upper bound of the support
- The Uniform$(0, 1)$ pdf is $f_X(x) = \mathbf{1}(0 < x < 1)$.

pdfs of several Normal distributions

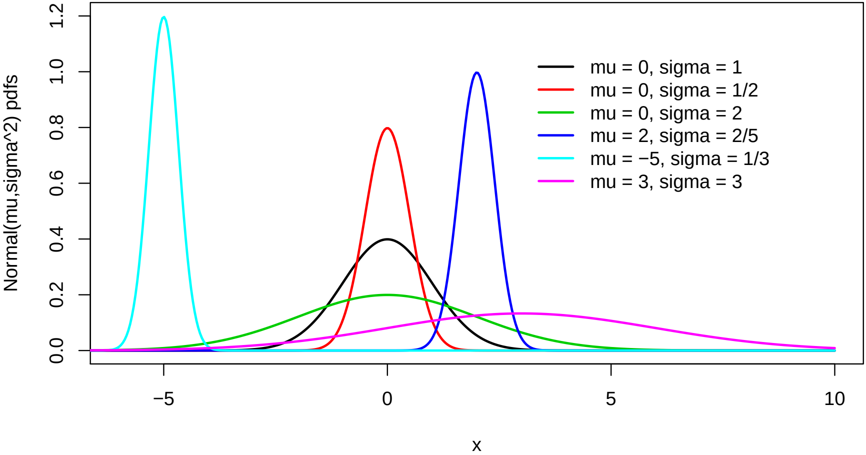

- The pdf of the Normal(μ, σ^2) distribution is given by

$$f_X(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] \quad \text{for } x \in \mathbb{R}.$$

- Parameters:
  - $\mu \in \mathbb{R}$ is a *location parameter*
  - $\sigma > 0$ is a *scale parameter*
- If $X \sim$ Normal(μ, σ^2), then
  - $\mathbb{E}X = \mu$
  - $\operatorname{Var} X = \sigma^2$
  - X has mgf $M_X(t) = e^{\mu t + \sigma^2 t^2/2}$ for all $t \in \mathbb{R}$
- The pdf and cdf of the Normal$(0, 1)$ distribution get special notation:

$$\phi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2}$$

$$\Phi(z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt \quad \text{for } z \in \mathbb{R}.$$

Exercises: Let $X \sim \text{Normal}(\mu, \sigma^2)$ and $Z \sim \text{Normal}(0, 1)$. Show

1. $M_X(t) = e^{\mu t + \sigma^2 t^2/2}$ for all $t \in \mathbb{R}$
2. $X \overset{d}{=} \sigma Z + \mu$
3. $\mathbb{E}Z = 0$
4. $\text{Var}\, Z = 1$
5. $\mathbb{E}X = \mu$
6. $\text{Var}\, X = \sigma^2$
7. $\int_{-\infty}^{\infty} \phi(z) dz = 1$

pdfs of several Gamma distributions

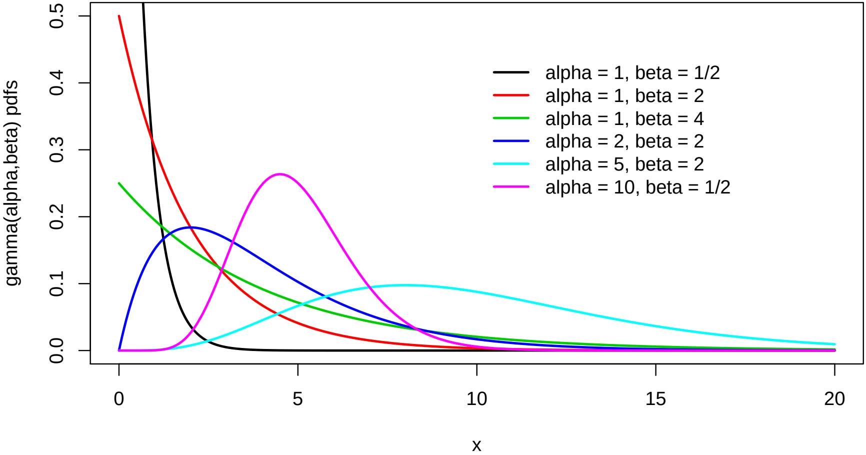

- The pdf of the Gamma(α, β) distribution is given by

$$f_X(x;\alpha,\beta) = \frac{1}{\Gamma(\alpha)\beta^\alpha} x^{\alpha-1} \exp\left[-\frac{x}{\beta}\right] \quad \text{for } x > 0.$$

- Parameters:
  - $\alpha > 0$ is a *shape parameter*
  - $\beta > 0$ is a *scale parameter*
- If $X \sim \text{Gamma}(\alpha, \beta)$, then
  - $\mathbb{E}X = \alpha\beta$
  - $\text{Var}\, X = \alpha\beta^2$
  - X has mgf $M_X(t) = (1 - \beta t)^{-\alpha}$ for $t < 1/\beta$.

The gamma distributions are brought to you by the gamma function.

Gamma function

For any $\alpha \in \mathbb{C}$ with $\text{Re}(\alpha) > 0$, the *gamma function* is given by

$$\Gamma(\alpha) = \int_0^\infty u^{\alpha-1} e^{-u} du.$$

These are some of its properties:

1. $\Gamma(\alpha + 1) = \alpha\Gamma(\alpha)$ for all $\text{Re}(\alpha) > 0$.
2. $\Gamma(n) = (n-1)!$ for any integer $n > 0$.
3. $\Gamma(1/2) = \sqrt{\pi}$.

Exercise:

1. Prove the above properties.
2. Show $\mathbb{E}X = \alpha\beta$ if $X \sim \text{Gamma}(\alpha, \beta)$.

pdfs of several Exponential distributions

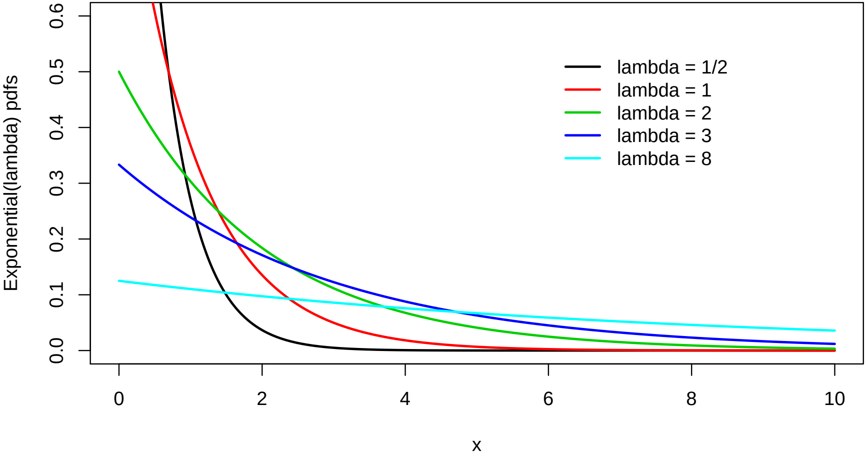

- The pdf of the Exponential(λ) distribution is given by

$$f_X(x;\lambda) = \frac{1}{\lambda}\exp\left[-\frac{x}{\lambda}\right] \quad \text{for } x > 0.$$

- Parameter:
  - $\lambda > 0$ is a *scale parameter*
- If $X \sim$ Exponential(λ), then
  - $\mathbb{E}X = \lambda$
  - $\text{Var}\, X = \lambda^2$
  - X has mgf $M_X(t) = (1 - \lambda t)^{-1}$ for $t < 1/\lambda$.
  - $X \sim \text{Gamma}(1, \lambda)$

Exercise: Find the cdf of the Exponential(λ) distribution.

pdfs of several Chi-squared distributions

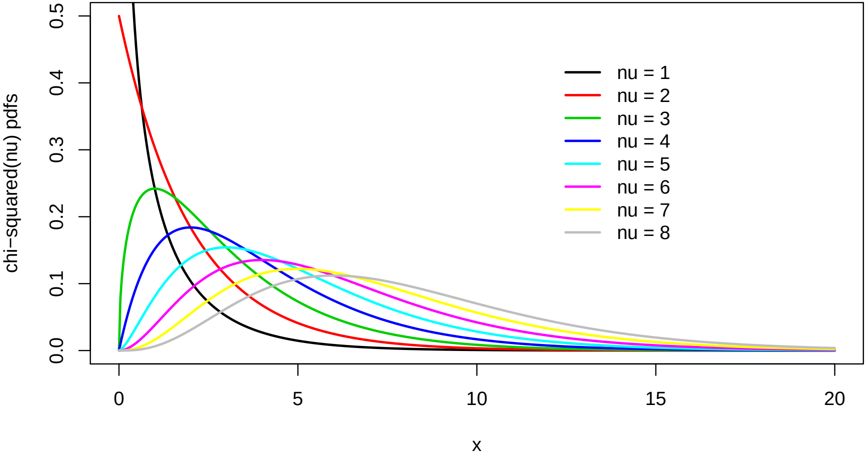

- The pdf of the Chi-squared(ν) distribution is given by

$$f_X(x;\nu) = \frac{1}{\Gamma(\nu/2)2^{\nu/2}} x^{\nu/2-1} \exp\left[-\frac{x}{2}\right] \quad \text{for } x > 0.$$

- Parameter:
  - $\nu > 0$ is called the *degrees of freedom*
- If $X \sim$ Chi-squared(ν), then
  - $\mathbb{E}X = \nu$
  - $\text{Var}\, X = 2\nu$
  - X has mgf $M_X(t) = (1-2t)^{-\nu/2}$ for $t < 1/2$.
  - $X \sim \text{Gamma}(\nu/2, 2)$

pdfs of several Beta distributions

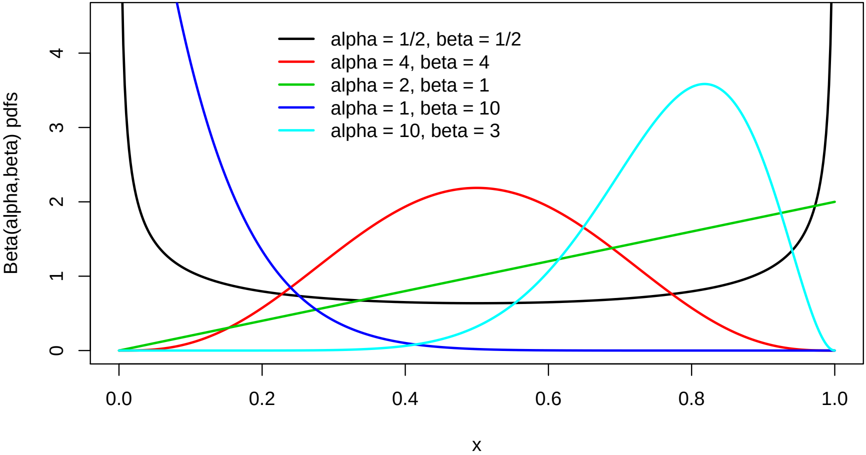

- The pdf of the Beta(α, β) distribution is given by

$$f_X(x; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1} \quad \text{for } x \in (0,1).$$

- Parameter:
  - $\alpha > 0$ is a *shape parameter*
  - $\beta > 0$ is a *shape parameter*
- If $X \sim \text{Beta}(\alpha, \beta)$, then
  - $\mathbb{E}X = \dfrac{\alpha}{\alpha+\beta}$
  - $\text{Var}\, X = \dfrac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$
  - X has mgf $M_X(t) = 1 + \sum_{k=1}^{\infty} \frac{t^k}{k!} \left(\prod_{r=0}^{k-1} \frac{\alpha+r}{\alpha+\beta+r} \right)$ for all $t \in \mathbb{R}$.
- The Beta$(1,1)$ distribution is the Uniform$(0,1)$ distribution.

The beta distributions are brought to you by the beta function.

Beta function

For any $\alpha, \beta \in \mathbb{C}$ with $\text{Re}(\alpha) > 0$ and $\text{Re}(\beta) > 0$, the *beta function* is given by

$$B(\alpha, \beta) = \int_0^1 u^{\alpha-1}(1-u)^{\beta-1} du = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

Exercise: Show $\mathbb{E}X = \alpha/(\alpha+\beta)$ if $X \sim \text{Beta}(\alpha, \beta)$.

We will now start talking about *families of distributions*.

Parametric family

For a pdf/pmf $f(\cdot;\theta)$ depending on $\theta = (\theta_1, \dots, \theta_d) \in \Theta$, the collection

$$\{f(\cdot;\theta) : \theta \in \Theta\}$$

of pdfs/pmfs is called a *parametric family* of pdfs/pmfs .

Example: The Beta parametric family is the set of pdfs

$$\left\{ f(x;\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}\mathbf{1}(0 < x < 1) : \alpha > 0, \beta > 0 \right\}.$$

Exponential family

A parametric family $\{f(\cdot;\theta) : \theta \in \Theta\}$ is called an *exponential family* if each member can be written as

$$f(x;\theta) = h(x)c(\theta)\exp\left(\sum_{i=1}^{k} w_i(\theta)t_i(x)\right), \quad x \in \mathbb{R}$$

for some real-valued functions

- $h(\cdot) \geq 0$ and $t_1(\cdot),\dots,t_k(\cdot)$ not depending on θ and
- $c(\cdot) \geq 0$ and $w_1(\cdot),\dots,w_k(\cdot)$ not depending on x.

Why though??

- Many common families of distributions are exponential families.
- Generalized linear models are based on exponential family distributions.

Exercise: Show that the Normal(μ, σ^2) pdfs for $\mu \in \mathbb{R}$ and $\sigma > 0$ are an exponential family.

Exercise: Show that the Poisson(λ) pmfs for $\lambda > 0$ are an exponential family.

Curved versus full exponential families

An exponential family $\{f(\cdot;\theta) : \theta \in \Theta\}$ is a *curved exponential family* if $d = \dim(\Theta) < k$, with k from the representation

$$f(x;\theta) = h(x)c(\theta)\exp\left(\sum_{i=1}^{k} w_i(\theta)t_i(x)\right), \quad x \in \mathbb{R}.$$

If $d = k$, then $\{f(\cdot;\theta) : \theta \in \Theta\}$ is a *full exponential family*.

Examples:

1. The Beta$(\alpha, 2\alpha)$ pdfs for all $\alpha > 0$.
2. The Normal(μ, μ^2) pdfs for all $\mu \in \mathbb{R}$.