# Working with the ought-to-know continuous rv distributions in R

*Karl Gregory*

*9/27/2017*

## Normal

The Normal($\mu, \sigma^2$) pdf is

$$f_X(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right], \quad -\infty < x < \infty.$$

In R, we can type `?dnorm` into the console to pull up the R documentation about the Normal distribution functions. We see that we can compute the height of the density curve $f_X(x; \mu, \sigma^2)$ as `dnorm(x = x, mean = mu, sd = sigma)`. Note that R takes the parameter $\sigma$, which is the standard deviation, instead of $\sigma^2$.

### Normal cdf

We can compute

$$F_X(x; \mu, \sigma^2) = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} \frac{1}{\sigma} \exp\left[-\frac{(t-\mu)^2}{2\sigma^2}\right] dt, \quad -\infty < x < \infty$$

in R as `pnorm(q = x , mean = mu, sd = sigma)`.

### Normal quantiles

We can compute

$$Q_X(u; \mu, \sigma^2) = \inf\{x : F_X(x; \mu, \sigma^2) \geq u\}, \quad 0 < u < 1$$

in R as `qnorm(p = u, mean = mu, sd = sigma)`.

### Plots of different Normal pdfs

Let's plot the pdfs of some different Normal distributions:

```
# create a sequence of x values at which to compute the Normal pdfs
x.seq <- seq(-10,10, length=500)

# create an empty plot to which we can add several lines
plot(NA,xlim=c(-6,10),ylim=c(0,1.2),ylab="Normal(mu,sigma^2) pdfs",xlab="x")

# plot Normal pdf against sequence of x values for different choices of mu and sigma
lines(dnorm(x.seq,mean = 0, sd = 1)~x.seq,col=1,lwd=2) # mu = 1, sigma = 1
lines(dnorm(x.seq,mean = 0, sd = 1/2)~x.seq,col=2,lwd=2) # mu = 1, sigma = 1/2
lines(dnorm(x.seq,mean = 0, sd = 2)~x.seq,col=3,lwd=2) # mu = 0, sigma = 2
lines(dnorm(x.seq,mean = 2, sd = 2/5)~x.seq,col=4,lwd=2) # mu = 2, sigma = 2/5
lines(dnorm(x.seq,mean = -5, sd = 1/3)~x.seq,col=5,lwd=2) # mu = -5, sigma = 1/3
lines(dnorm(x.seq,mean = 3, sd = 3)~x.seq,col=6,lwd=2) # mu = 3, sigma = 3

# add a legend to the plot
```

```
legend( x = 3,y=1.1, legend=c("mu = 0, sigma = 1",
                                 "mu = 0, sigma = 1/2",
                                 "mu = 0, sigma = 2",
                                 "mu = 2, sigma = 2/5",
                                 "mu = -5, sigma = 1/3",
                                 "mu = 3, sigma = 3"),col=c(1,2,3,4,5,6),lwd=2,bty="n")
```

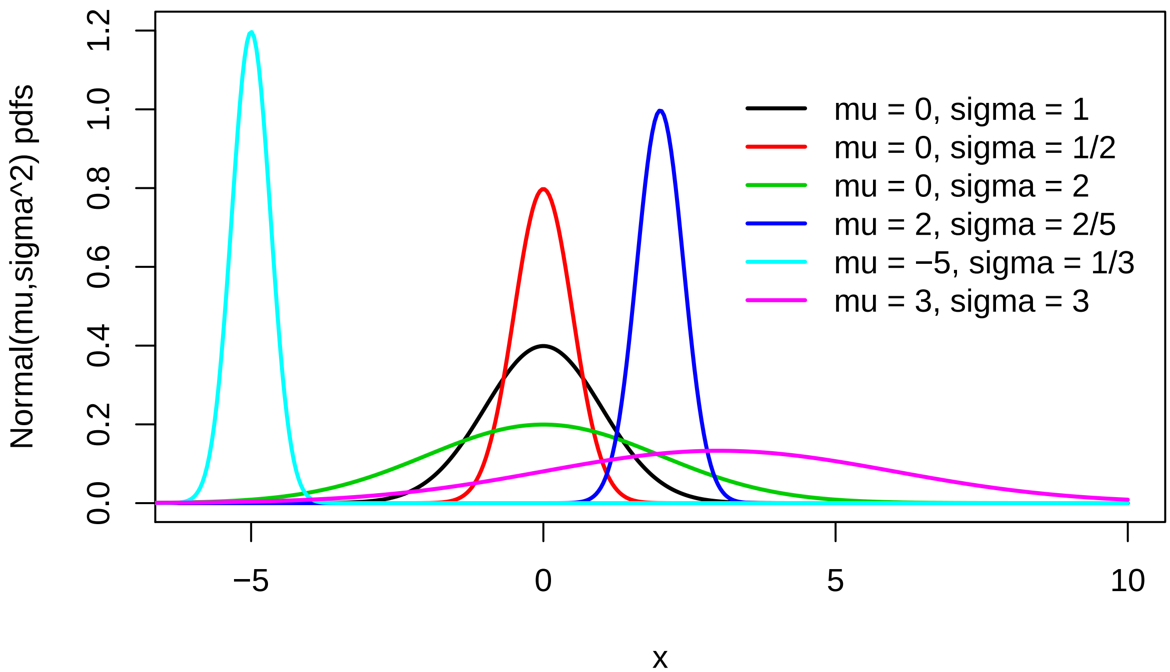


### Gamma

The gamma$(\alpha, \beta)$ pdf is

$$f_X(x; \alpha, \beta) = \begin{cases} \frac{1}{\Gamma(\alpha)\beta^\alpha} x^{\alpha-1} \exp\left[-\frac{x}{\beta}\right] & x > 0 \\ 0 & x \leq 0. \end{cases}$$

In R, we can type `?dgamma` into the console to pull up the R documentation about the gamma distribution functions. We see that we can compute the height of the density curve $f_X(x; \alpha, \beta)$ as `dgamma(x = x, shape = alpha, scale = beta)`.

#### Gamma cdf

We can compute

$$F_X(x; \alpha, \beta) = \begin{cases} \int_0^x \frac{1}{\Gamma(\alpha)\beta^\alpha} t^{\alpha-1} \exp\left[-\frac{t}{\beta}\right] dt & 0 < x < \infty \\ 0 & -\infty < x \leq 0 \end{cases}$$

in R as `pgamma(q = x , shape = alpha, scale = beta)`.

#### Gamma quantiles

We can compute

$$Q_X(u; \alpha, \beta) = \inf\{x : F_X(x; \alpha, \beta) \geq u\}, \quad 0 < u < 1$$

in R as `qgamma(p = u, shape = alpha, scale = beta)`.

**Plots of different Gamma pdfs**

Let's plot the pdfs of some different gamma distributions:

```
# create a sequence of x values at which to compute the gamma pdfs
x.seq <- seq(0,20, length=500)

# create an empty plot to which we can add several lines
plot(NA,xlim=c(0,20),ylim=c(0,.5),ylab="gamma(alpha,beta) pdfs",xlab="x")

# plot gamma pdf against sequence of x values for different choices of alpha and beta
lines(dgamma(x.seq,shape = 1, scale = 1/2)~x.seq,col=1,lwd=2) # alpha = 1, beta = 1/2
lines(dgamma(x.seq,shape = 1, scale = 2)~x.seq,col=2,lwd=2) # alpha = 1, beta = 2
lines(dgamma(x.seq,shape = 1, scale = 4)~x.seq,col=3,lwd=2) # alpha = 1, beta = 4
lines(dgamma(x.seq,shape = 2, scale = 2)~x.seq,col=4,lwd=2) # alpha = 2, beta = 2
lines(dgamma(x.seq,shape = 5, scale = 2)~x.seq,col=5,lwd=2) # alpha = 5, beta = 2
lines(dgamma(x.seq,shape = 10, scale = 1/2)~x.seq,col=6,lwd=2) # alpha = 10, beta = 1/2

# add a legend to the plot
legend( x = 10,y=.45, legend=c("alpha = 1, beta = 1/2",
                               "alpha = 1, beta = 2",
                               "alpha = 1, beta = 4",
                               "alpha = 2, beta = 2",
                               "alpha = 5, beta = 2",
                               "alpha = 10, beta = 1/2"),col=c(1,2,3,4,5,6),lwd=2,bty="n")
```

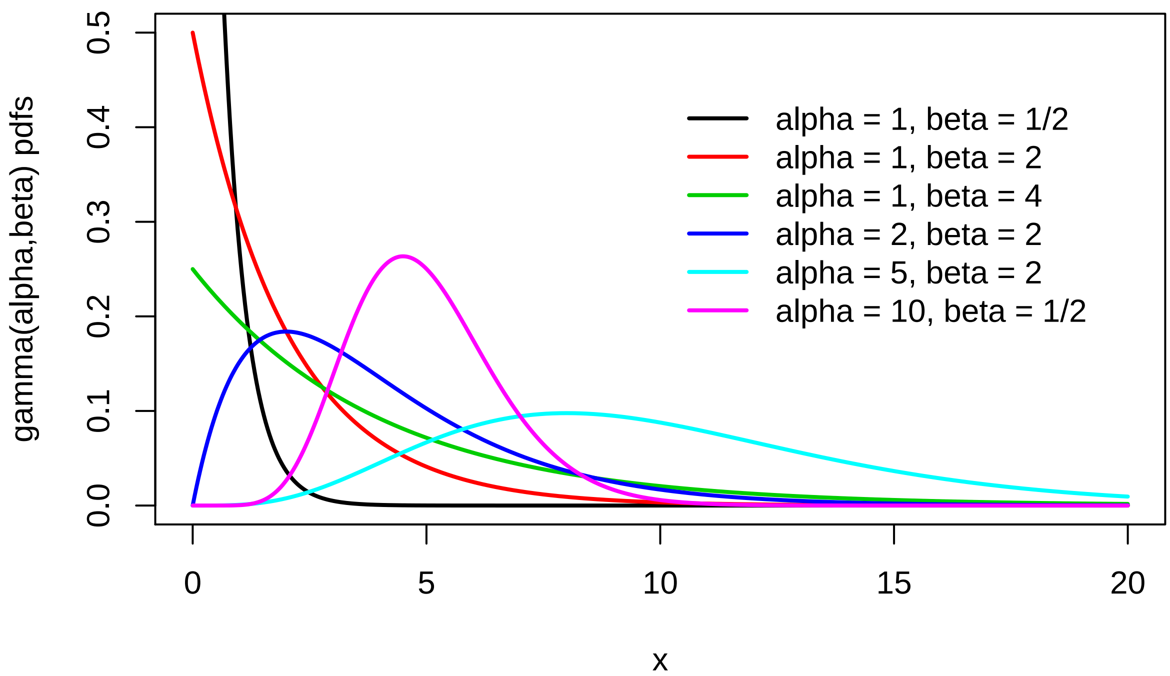


## Exponential

The Exponential($\lambda$) pdf is

$$f_X(x;\lambda) = \begin{cases} \frac{1}{\lambda}\exp\left[-\frac{x}{\lambda}\right] & x > 0 \\ 0 & x \leq 0. \end{cases}$$

In R, we can type `?dexp` into the console to pull up the R documentation about the Normal distribution functions. We see that we can compute the height of the density curve $f_X(x; \lambda)$ as `dexp(x = x, rate = 1/lambda)`. Note that R parameterizes the exponential distribution differently, such that we must put in $1/\lambda$ instead of $\lambda$.

### Exponential cdf

We can compute

$$F_X(x; \lambda) = \begin{cases} 1 - \exp[-\frac{x}{\lambda}] & 0 < x < \infty \\ 0 & -\infty < x \leq 0 \end{cases}$$

in R as `pexp(q = x , rate = 1/lambda)`.

### Exponential quantiles

We can compute

$$Q_X(u; \lambda) = \inf\{x : F_X(x; \lambda) \geq u\}, \quad 0 < u < 1$$

in R as `qexp(p = u, rate = 1/lambda)`.

### Plots of different Exponential pdfs

Let's plot the pdfs of some different exponential distributions:

```r
# create a sequence of x values at which to compute the exponential pdfs
x.seq <- seq(0,10, length=500)

# create an empty plot to which we can add several lines
plot(NA,xlim=c(0,10),ylim=c(0,.6),ylab="Exponential(lambda) pdfs",xlab="x")

# plot exponential pdf against sequence of x values for different choices of lambda
lines(dexp(x.seq,rate=1/(.5))~x.seq,col=1,lwd=2) # lambda = 1/2
lines(dexp(x.seq,rate=1/1)~x.seq,col=2,lwd=2) # lambda = 1
lines(dexp(x.seq,rate=1/2)~x.seq,col=3,lwd=2) # lambda = 2
lines(dexp(x.seq,rate=1/3)~x.seq,col=4,lwd=2) # lambda = 3
lines(dexp(x.seq,rate=1/8)~x.seq,col=5,lwd=2) # lambda = 4

# add a legend to the plot
legend( x = 6,y=.55, legend=c("lambda = 1/2",
                              "lambda = 1",
                              "lambda = 2",
                              "lambda = 3",
                              "lambda = 8"),col=c(1,2,3,4,5),lwd=2,bty="n")
```

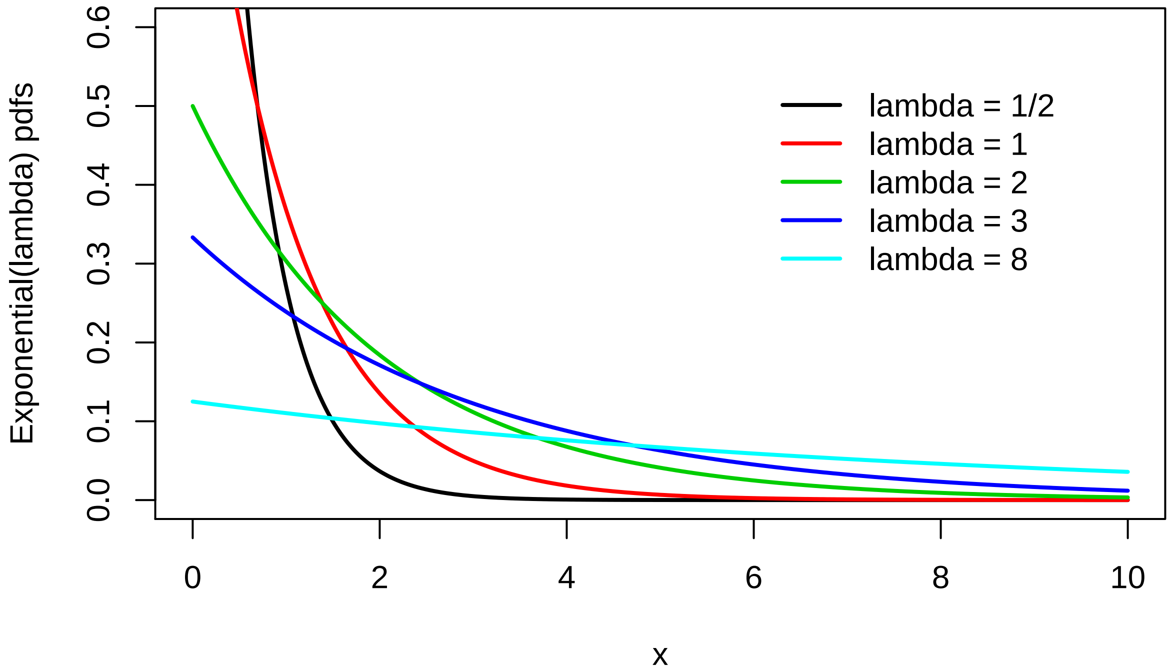


## Chi-squared

The Chi-squared($\nu$) pdf is

$$f_X(x;\nu) = \begin{cases} \frac{1}{\Gamma(\nu/2)2^{\nu/2}} x^{\nu/2-1} \exp\left[-\frac{x}{2}\right] & 0 < x < \infty \\ 0 & -\infty < x \leq 0. \end{cases}$$

In R, we can type `?dchisq` into the console to pull up the R documentation about the chi-squared distribution functions. We see that we can compute the height of the density curve $f_X(x;\nu)$ as `dchisq(x = x, df = nu)`. The `df` stands for "degrees of freedom", which is the name of the parameter we have denoted by $\nu$.

### Chi-squared cdf

We can compute

$$F_X(x;\nu) = \begin{cases} \int_0^x \frac{1}{\Gamma(\nu/2)2^{\nu/2}} t^{\nu/2-1} \exp\left[-\frac{t}{2}\right] dt & 0 < x < \infty \\ 0 & -\infty < x \leq 0 \end{cases}$$

in R as `pchisq(q = x , df = nu)`.

### Chi-squared quantiles

We can compute

$$Q_X(u;\nu) = \inf\{x : F_X(x;\nu) \geq u\}, \quad 0 < u < 1$$

in R as `qchisq(p = u, df = nu)`.

### Plots of different Chi-squared pdfs

Let's plot the pdfs of some different gamma distributions:

```r
# create a sequence of x values at which to compute the chi-squared pdfs
x.seq <- seq(0,20, length=500)

# create an empty plot to which we can add several lines
plot(NA,xlim=c(0,20),ylim=c(0,.5),ylab="chi-squared(nu) pdfs",xlab="x")

# plot chi-squared pdf against sequence of x values for different choices of nu.
lines(dchisq(x.seq, df = 1)~x.seq,col=1,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 2)~x.seq,col=2,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 3)~x.seq,col=3,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 4)~x.seq,col=4,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 5)~x.seq,col=5,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 6)~x.seq,col=6,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 7)~x.seq,col=7,lwd=2) # nu = 1
lines(dchisq(x.seq, df = 8)~x.seq,col=8,lwd=2) # nu = 1

# add a legend to the plot
legend( x = 12,y=.45, legend=c("nu = 1",
                               "nu = 2",
                               "nu = 3",
                               "nu = 4",
                               "nu = 5",
                               "nu = 6",
                               "nu = 7",
                               "nu = 8"),col=c(1,2,3,4,5,6,7,8),lwd=2,bty="n")
```

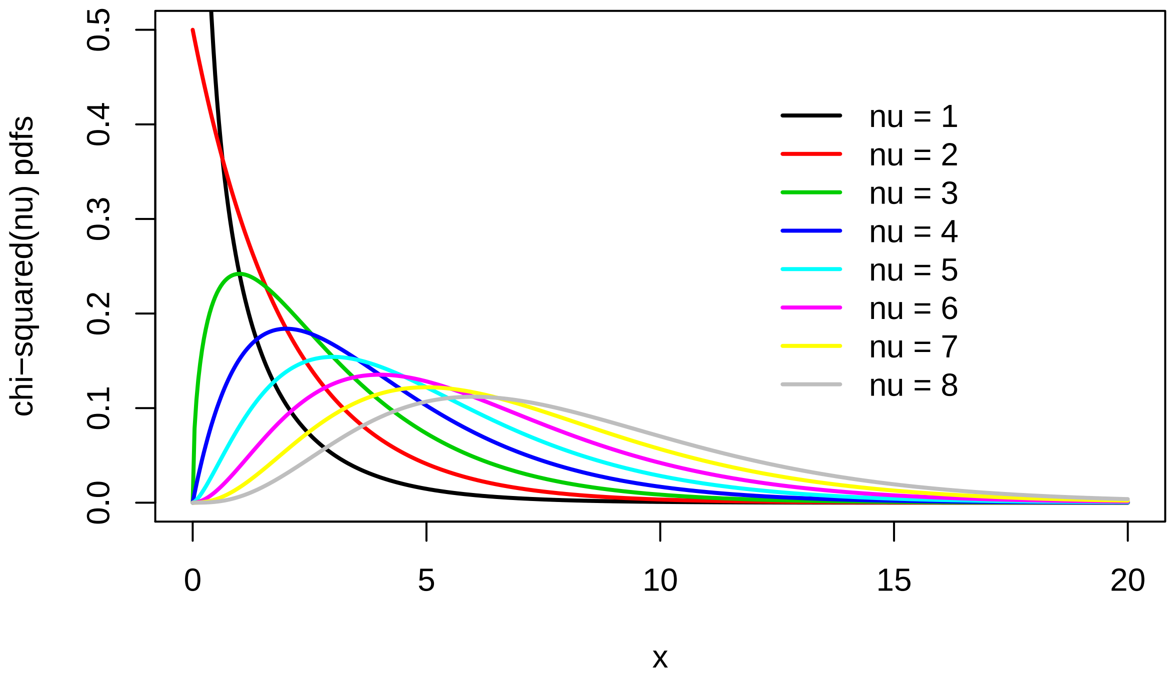


## Beta

The Beta$(\alpha, \beta)$ pdf is

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

In R, we can type `?dbeta` into the console to pull up the R documentation about the chi-squared distribution functions. We see that we can compute the height of the density curve $f_X(x;\alpha,\beta)$ as `dbeta(x = x, shape1 = alpha, shape2 = beta)`.

### Beta cdf

We can compute

$$F_X(x;\alpha,\beta) = \begin{cases} 1 & 1 \leq x < \infty \\ \int_0^x \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} t^{\alpha-1}(1-t)^{\beta-1}dt & 0 < x < 1 \\ 0 & -\infty < x \leq 0 \end{cases}$$

in R as `pbeta(q = x, shape1 = alpha, shape2 = beta)`.

### Beta quantiles

We can compute

$$Q_X(u;\alpha,\beta) = \inf\{x : F_X(x;\alpha,\beta) \geq u\}, \quad 0 < u < 1$$

in R as `qbeta(p = u, shape1 = alpha, shape2 = beta)`.

### Plots of different Beta pdfs

Let's plot the pdfs of some different beta distributions:

```r
# create a sequence of x values at which to compute the beta pdfs
x.seq <- seq(0,1, length=500)

# create an empty plot to which we can add several lines
plot(NA,xlim=c(0,1),ylim=c(0,4.5),ylab="Beta(alpha,beta) pdfs",xlab="x")

# plot beta pdf against sequence of x values for different choices of alpha and beta.
lines(dbeta(x.seq, shape1 = 1/2, shape2 = 1/2)~x.seq,col=1,lwd=2) # alpha = 1/2, beta = 1/2
lines(dbeta(x.seq, shape1 = 4, shape2 = 4)~x.seq,col=2,lwd=2) # alpha = 5, beta = 5
lines(dbeta(x.seq, shape1 = 2, shape2 = 1)~x.seq,col=3,lwd=2) # alpha = 2, beta = 1
lines(dbeta(x.seq, shape1 = 1, shape2 = 10)~x.seq,col=4,lwd=2) # alpha = 1, beta = 10
lines(dbeta(x.seq, shape1 = 10, shape2 = 3)~x.seq,col=5,lwd=2) # alpha = 10, beta = 3

# add a legend to the plot
legend( x = .2,y=4.5, legend=c("alpha = 1/2, beta = 1/2",
                               "alpha = 4, beta = 4",
                               "alpha = 2, beta = 1",
                               "alpha = 1, beta = 10",
                               "alpha = 10, beta = 3"),col=c(1,2,3,4,5),lwd=2,bty="n")
```

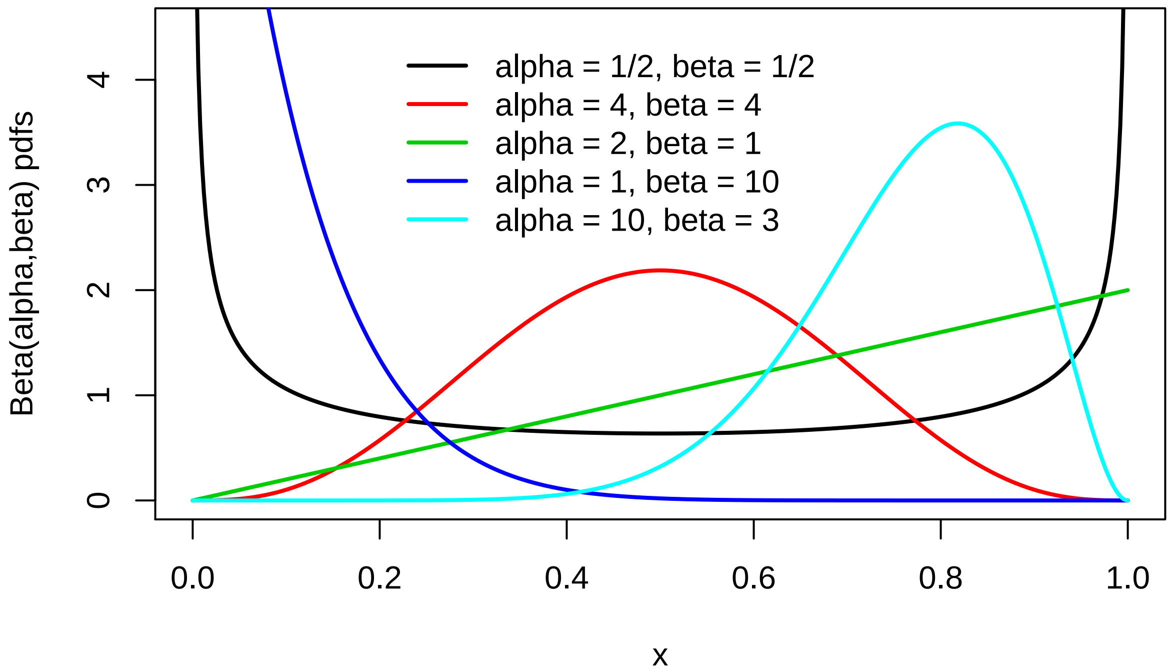
alpha = 1/2, beta = 1/2
alpha = 4, beta = 4
alpha = 2, beta = 1
alpha = 1, beta = 10
alpha = 10, beta = 3
Beta(alpha,beta) pdfs
0
1
2
3
4
0.0
0.2
0.4
0.6
0.8
1.0
x