

Gamma Distribution

Paul E. Johnson <pauljohn@ku.edu>

January 4, 2006

1 Mathematical Description

Gamma is a probability model for a continuous variable on $[0, \infty)$. It can look like a “ski slope” or it can look single-peaked.

It has 2 parameters, shape and scale. I hate to put in symbols for these because I always forget which is which. So let’s just use the words.

In some books, the scale is specified as $1/rate$. Many models set $scale=1$, so $rate=1$, and there is no difference.

If x_i is Gamma distributed, the probability density function is:

$$f(x_i) = \frac{1}{scale^{shape} \Gamma(shape)} x_i^{(shape-1)} e^{-\left(\frac{x_i}{scale}\right)}$$

If you read a different book, you might find it rearranged like so:

$$f(x_i) = \frac{1}{scale * \Gamma(shape)} \left[\frac{x_i}{scale} \right]^{(shape-1)} e^{-\left(\frac{x_i}{scale}\right)}$$

That’s really just the same distribution.

2 Mathematical detail

What is that function Γ ?

The function $\Gamma(s)$ is the Gamma function, formally defined as $\Gamma(s) = \int_0^\infty t^{s-1} e^{-t} dt$ if $s > 0$. The Gamma function arises in many statistical applications. The formula appears to be complicated, but just remember: its just the factorial function “extended” to take on values between the integers. Here’s why”

If you pick s as an integer, $\Gamma(s)$ is:

$$\Gamma(s) = (s - 1)! \text{ for } s = 1, 2, 3...$$

So, the value of $\Gamma(1) = 1$. And $\Gamma(2) = 1$, $\Gamma(3) = 2$, $\Gamma(4) = 6$, and so forth. And $\Gamma(20)$ is some impossibly huge number.

I never really worry very much about the value of Gamma, but in case you do, here’s a graph of it:

Figure 1: The Gamma Function

```
> xvals <- seq(0, 5, by = 0.05)
> gammaxvals <- gamma(xvals)
> plot(xvals, gammaxvals, type = "l", xlab = "x", ylab = expression(Gamma(x)),
+      main = "The Gamma Function")
```

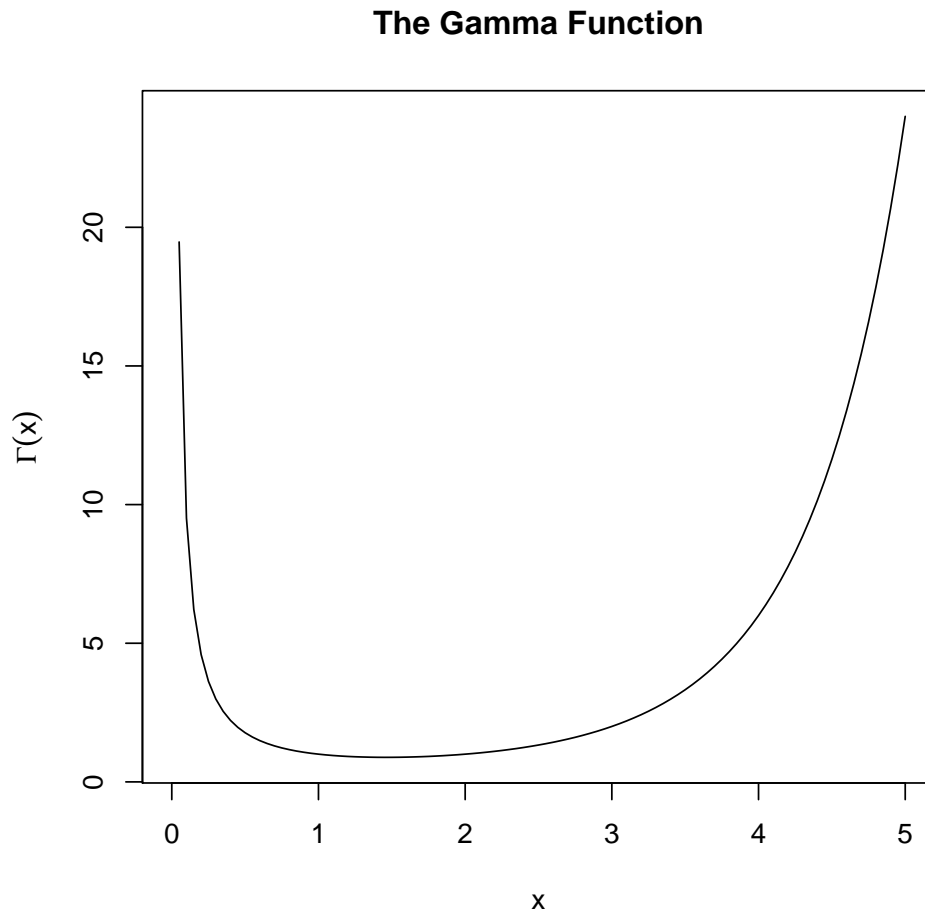
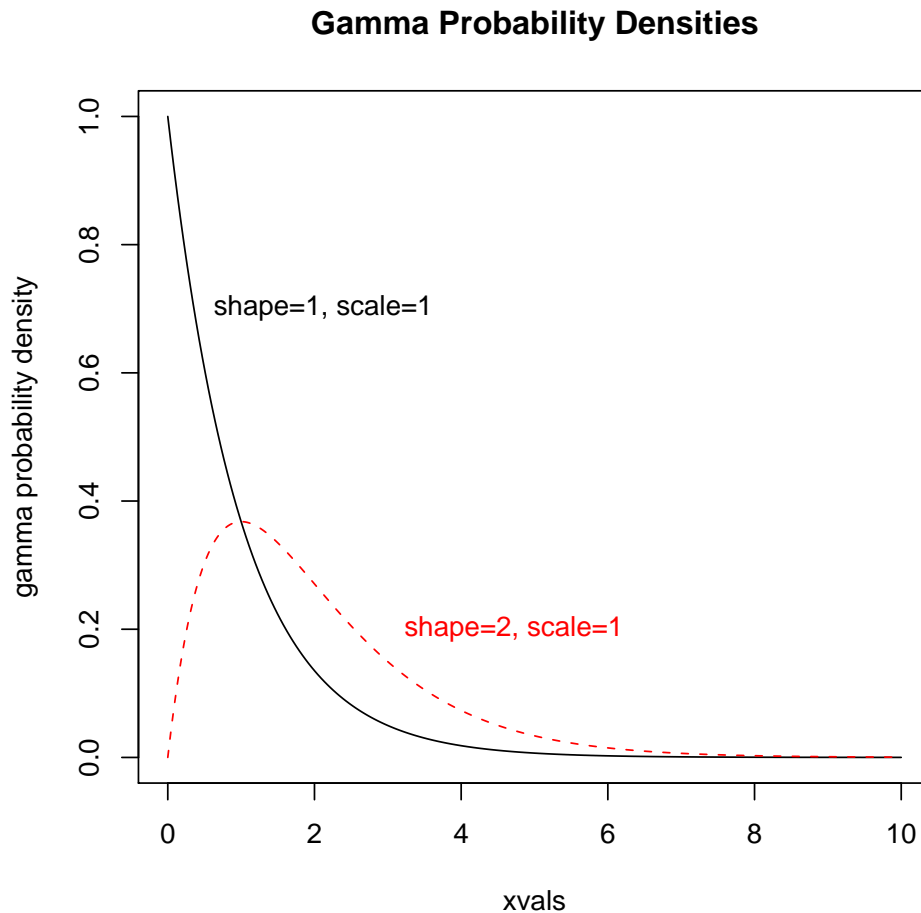


Figure 2: Gamma Density Functions



3 Illustrations

The probability density function of a the Gamma distribution changes quite a bit when one puts in different values of the parameters. If somebody knows some “interesting” parameter settings, then a clear, beautiful illustration of the Gamma can be produced. Consider the following code, which can be used to create the illustration of 2 possible Gamma density functions in Figure 2.

```
> xvals <- seq(0, 10, length.out = 1000)
> gam1 <- dgamma(xvals, shape = 1, scale = 1)
> gam2 <- dgamma(xvals, shape = 2, scale = 1)
> matplot(xvals, cbind(gam1, gam2), type = "l", ylab = "gamma probability density",
+         ylim = c(0, 1), main = "Gamma Probability Densities")
> text(0.4, 0.7, "shape=1, scale=1", pos = 4, col = 1)
> text(3, 0.2, "shape=2, scale=1", pos = 4, col = 2)
```

In many applications, such as the Negative Binomial regression model for overdispersed count data, Gamma density is used to insert observation-level heterogeneity. The scale parameter is set equal to 1. The following code will produce example density functions for a variety of shapes with the scale (and hence the rate) set at 1.

```
> par(mfcol = c(4, 2))
> for (i in 1:8) {
+   shape <- c(0.5, 1, 2, 3, 5, 10, 15, 30)
+   scale <- 1
+   xrange <- seq(from = 0, to = 50, by = 0.1)
+   mainlabel <- paste("Gamma(shape=", shape[i], ",scale=", scale,
+     ",)", sep = "")
+   xprob <- dgamma(xrange, shape = shape[i], scale = scale,
+     log = F)
+   plot(xrange, xprob, type = "l", main = mainlabel, xlab = "possible values of x",
+     ylab = "probability of x")
+ }
```

If you “assume away” the scale parameter by setting it equal to 1, then the probability density formula for the Gamma(scale=1, shape) simplifies to:

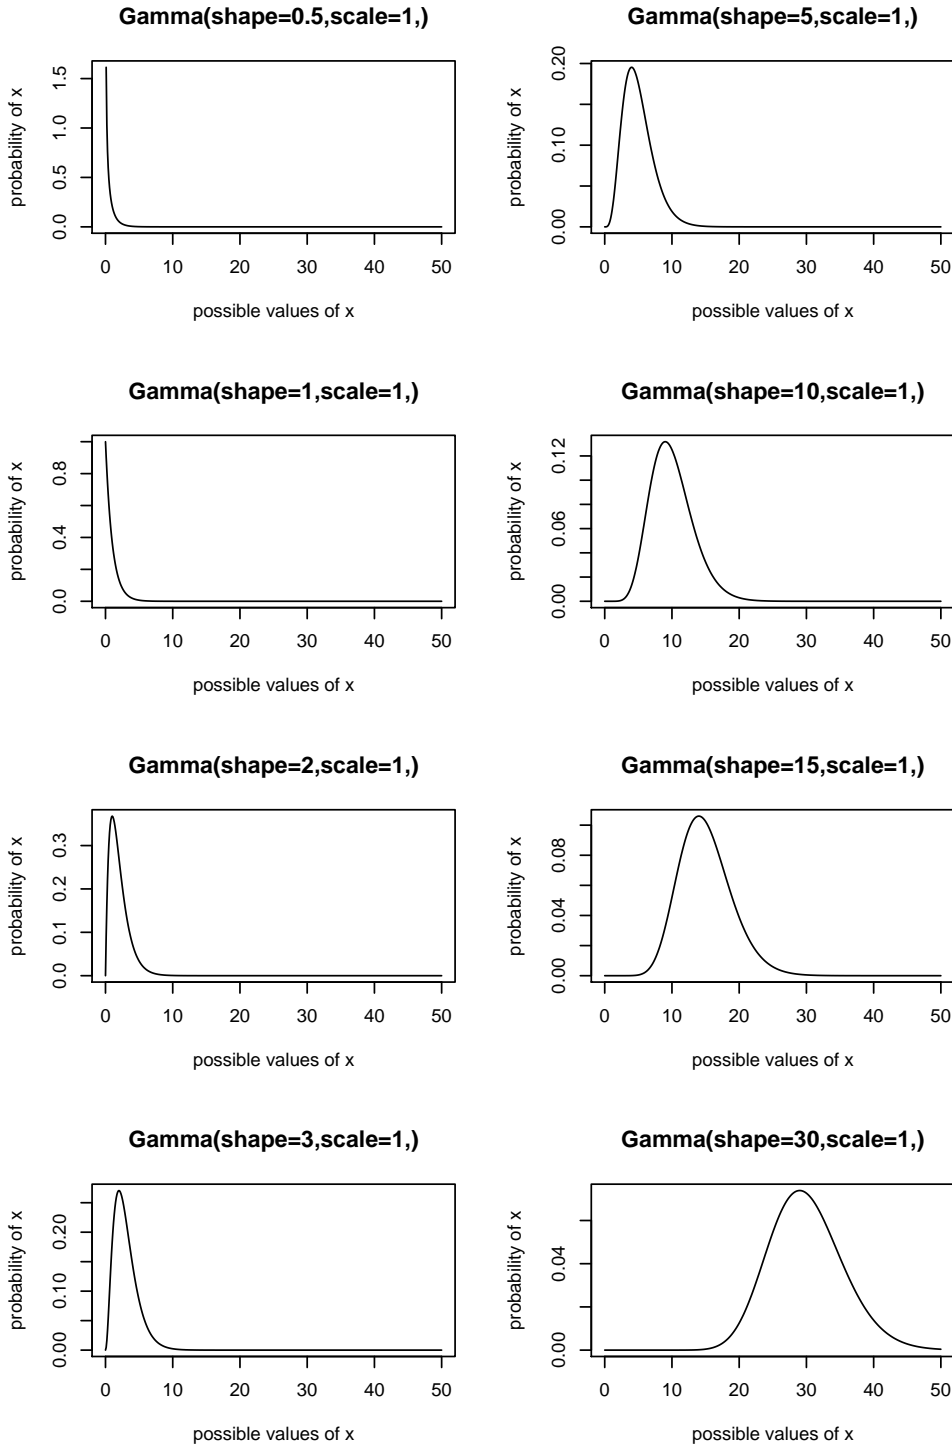
$$f(x_i) = \frac{1}{\Gamma(shape)} x_i^{(shape-1)} e^{-x_i} \quad x_i > 0$$

If shape=1, then this is an exponential distribution (because $\Gamma(1) = 1$). But, of course, the shape parameter can vary. If the scale can vary, what is the impact on the probability density? We can explore by making a grid of graphs with 3 different values for the scale parameter. Consider Figure 4, which results from this code:

```
> par(mfcol = c(4, 3))
> for (j in 1:3) {
+   for (i in 1:4) {
+     shape <- c(0.5, 1, 2, 10)
+     scale <- c(0.5, 1, 3)
+     xrange <- seq(from = 0, to = scale[j] * 20, by = 0.1)
+     mainlabel <- paste("shape=", shape[i], ",scale=", scale[j],
+       sep = "")
+     xprob <- dgamma(xrange, shape = shape[i], scale = scale[j],
+       log = F)
+     plot(xrange, xprob, type = "l", main = mainlabel, xlab = "possible values of
+       ylab = "probability of x")
+   }
+ }
```

The horizontal range of all of the graphs in each column is kept fixed, but the ranges used in the columns does differ. Hopefully, this gives you an idea why the scale parameter

Figure 3: Gammas Densities with scale=1



is often set at some arbitrary value, such as 1. On the surface, at least, the scale appears to be a “nuisance”, in the sense that its role can easily be understood and replicated later on. If one sets $scale=1$, and creates some variable x , then it appears one re-scale it “after the fact.”

4 Expected Value, Variance, and the role of the parameters

The Gamma probability distribution has these interesting properties:

$$E(x_i) = shape * scale$$

$$Var(x_i) = shape * scale^2$$

$$Std.Dev(x) = \sqrt{shape} \cdot scale$$

Figure 4: Gammas with Various Parameters

