

# Inverse Gaussian Distribution

Paul E. Johnson <pauljohn@ku.edu>

June 10, 2013

The Inverse Gaussian distribution is an exponential distribution. It is one of the distributions implemented in R's Generalized Linear Model routines. To my surprise, there are whole books dedicated to this distribution (V. Seshadri, *The Inverse Gaussian Distribution: A Case Study in Exponential Families*, Oxford University Press, 1994; R. S. Chhikara and J. L. Folks, *The Inverse Gaussian Distribution: Theory, Methodology, and Applications*, New York: Dekker, 1989).

Articles on insurance problems and the stock market often claim that observations follow an Inverse Gaussian distribution. It has one mode in the interior of the range of possible values and it is skewed to the right, sometimes with an extremely long tail. The fact that extremely large outcomes can occur even when almost all outcomes are small is a distinguishing characteristic.

## 1 Mathematical Description

There are Inverse Gaussian distributions in several R packages. Run

```
help.search("inverse gaussian")
```

to see for yourself.

In VGAM, the documentation for `inv.gaussianff` matches the information in in the package `statmod`'s documentation on `dinvgauss`. So let's follow that approach. The distribution of  $x_i$  is described by two characteristics, a mean  $\mu > 0$  and precision  $\lambda > 0$ . The probability density function is

$$p(x; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} e^{-\frac{\lambda(x-\mu)^2}{2\mu^2 x}}, 0 < x < \infty$$

If you would like to take  $\lambda$  out of the square root, you can put this down as

$$p(x; \mu, \lambda) = \frac{\lambda^2}{\sqrt{2\pi x^3}} e^{-\frac{(\sqrt{\lambda}(x-\mu))^2}{2\mu^2 x}}, 0 < x < \infty$$

The expected value is  $\mu$  and the variance of this version of the inverse Gaussian distribution is

$$\text{Var}(x) = \frac{\mu^3}{\lambda}$$

The skewness and kurtosis are, respectively,

$$3\sqrt{\frac{\mu}{\lambda}} \quad \text{and} \quad \frac{15\mu}{\lambda}$$

In all honesty, I have no intuition whatsoever about what the appearance of this probability model might be. It does not have a kernel smaller than the density itself, so we can't just throw away part on the grounds that it is a normalizing constant or a factor of proportionality.

And to make matters worse, the variance depends on the mean.

Tidbit: If  $\mu = 1$  this is called the Wald distribution.

## 2 Illustrations

The probability density function of a Inverse Gaussian distribution with  $\mu = 1$  and  $\lambda = 2$  is shown in Figure 1. The R code which produces that figure is:

```
library(statmod)
mu <- 1
lambda <- 2
xrange <- seq(from=0.0, to=2*mu+5/lambda, by=0.02)
mainlabel <- expression(paste("IG(", mu, ", ", lambda, ")",
  sep=""))
xprob <- dinvgauss(xrange, mu = mu, lambda = lambda)
plot(xrange, xprob, type = "l", main = mainlabel, xlab
     = "possible values of x", ylab = "probability of x
  ")
```

How would one describe Figure 1?

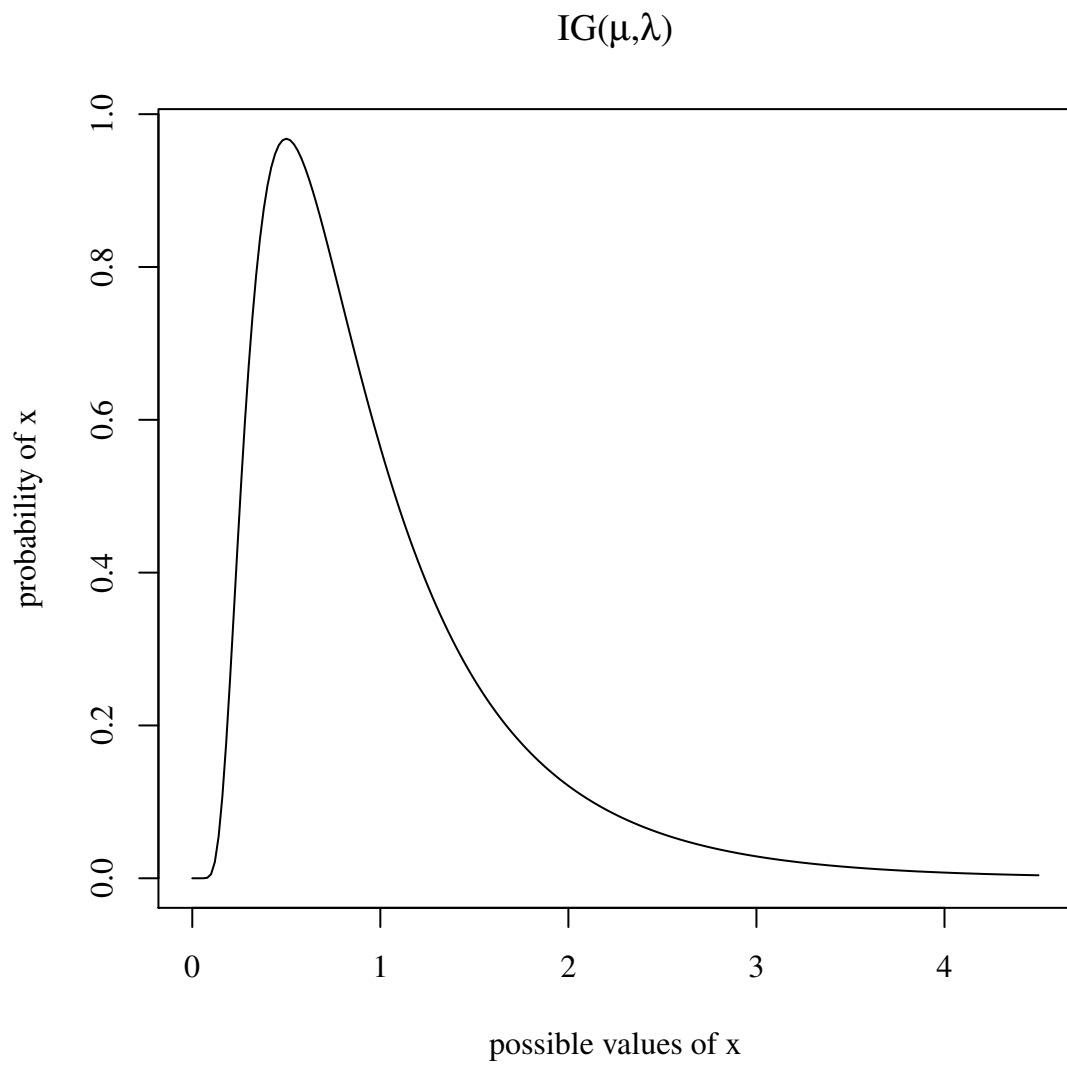
1. Single-peaked
2. Not symmetric— tail to the right

How does this distribution change in appearance if  $\mu$  and  $\lambda$  are changed?

Let's do some experimentation. The following R code creates an array of figures with 4 rows and 3 columns with various values of  $\mu$  and  $\lambda$  that is displayed in Figure 2.

```
par(mfrow=c(4,3))
for (i in 1:4) {
  for (j in 1:3) {
    mu <- 3*i
    lambda <- 20*j
```

Figure 1: Inverse Gaussian Distribution



```

xrange <- seq(from = 0.0 , to = 3*mu, by = 0.02)
mainlabel<- paste("IG(",mu," ",lambda,")", sep = " ")
xprob <- dinvgauss(xrange , mu = mu, lambda = lambda)
plot(xrange , xprob , type = "l" , main = mainlabel , xlab =
      "possible values of x" , ylab = "probability of x")
}
}

```

I had a very hard time believing that these calculations were correct. The graph shows almost no impact of changing parameters.

The mistake I had made was assuming that the Inverse Gamma's tail will be cut off in a way that makes a nice picture. It turns out that, when the  $\lambda$  is small, then there can be extremely huge values observed. I found that by plotting histograms of random samples from various parameter settings.

For example, if we draw 1000 observations from Inverse Gamma with  $\mu = 12$  and  $\lambda = 2$ , look what happens in Figure 3:

Can you describe this variety in a nutshell?

Sometimes the IG has a very long tail, stretching far to the right, and it makes the expected value a very poor description of the modal observation.

Probably the best illustration I have found for this model is presented in Figure 5.

In the Webpages for the old classic program Dataplot (<http://www.itl.nist.gov/div898/software/dataplo>) I found this interesting comment. That page uses the parameter gamma in place of lambda, otherwise the formula is the same. It says, "The inverse Gaussian distribution is symmetric and moderate tailed for small gamma. It is highly skewed and long tailed for large gamma. It approaches normality as gamma approaches zero." I can't find any evidence in favor of this characterization in Figure 6. The opposite is more likely true, making me suspect that in the Dataplot author's mind, the parameter gamma ( $\gamma$ ) might have at one time been  $1/\lambda$ . This leads me to caution studentst that if you want to be confident about one of these results, it is not sufficient to just take the word of a randomly chosen Web site.

To see the effect of tuning  $\lambda$  up and down, consider figure 6. This shows pretty clearly that if you think of  $\lambda$  as a precision parameter and you have high precision, then the observations are likely to be tightly clustered and symmetric about the mean. When you lose precision, as  $\lambda$  gets smaller, then the strong long tail to the right emerges.

### 3 Why do people want to use this distribution?

We want a distribution that can "reach up high" and admit some extreme values.

It is pretty easy to estimate  $\mu$  and  $\lambda$  by maximum likelihood.

An alternative distribution with this general shape is the three parameter Weibull distribution, which is more difficult to estimate (W.E. Bardsley, "Note on the Use of the Inverse Gaussian Distribution for Wind Energy Applications," Journal of Applied Meteorology, 19(9): 1126-1130).

Figure 2: Variety of Inverse Gaussian

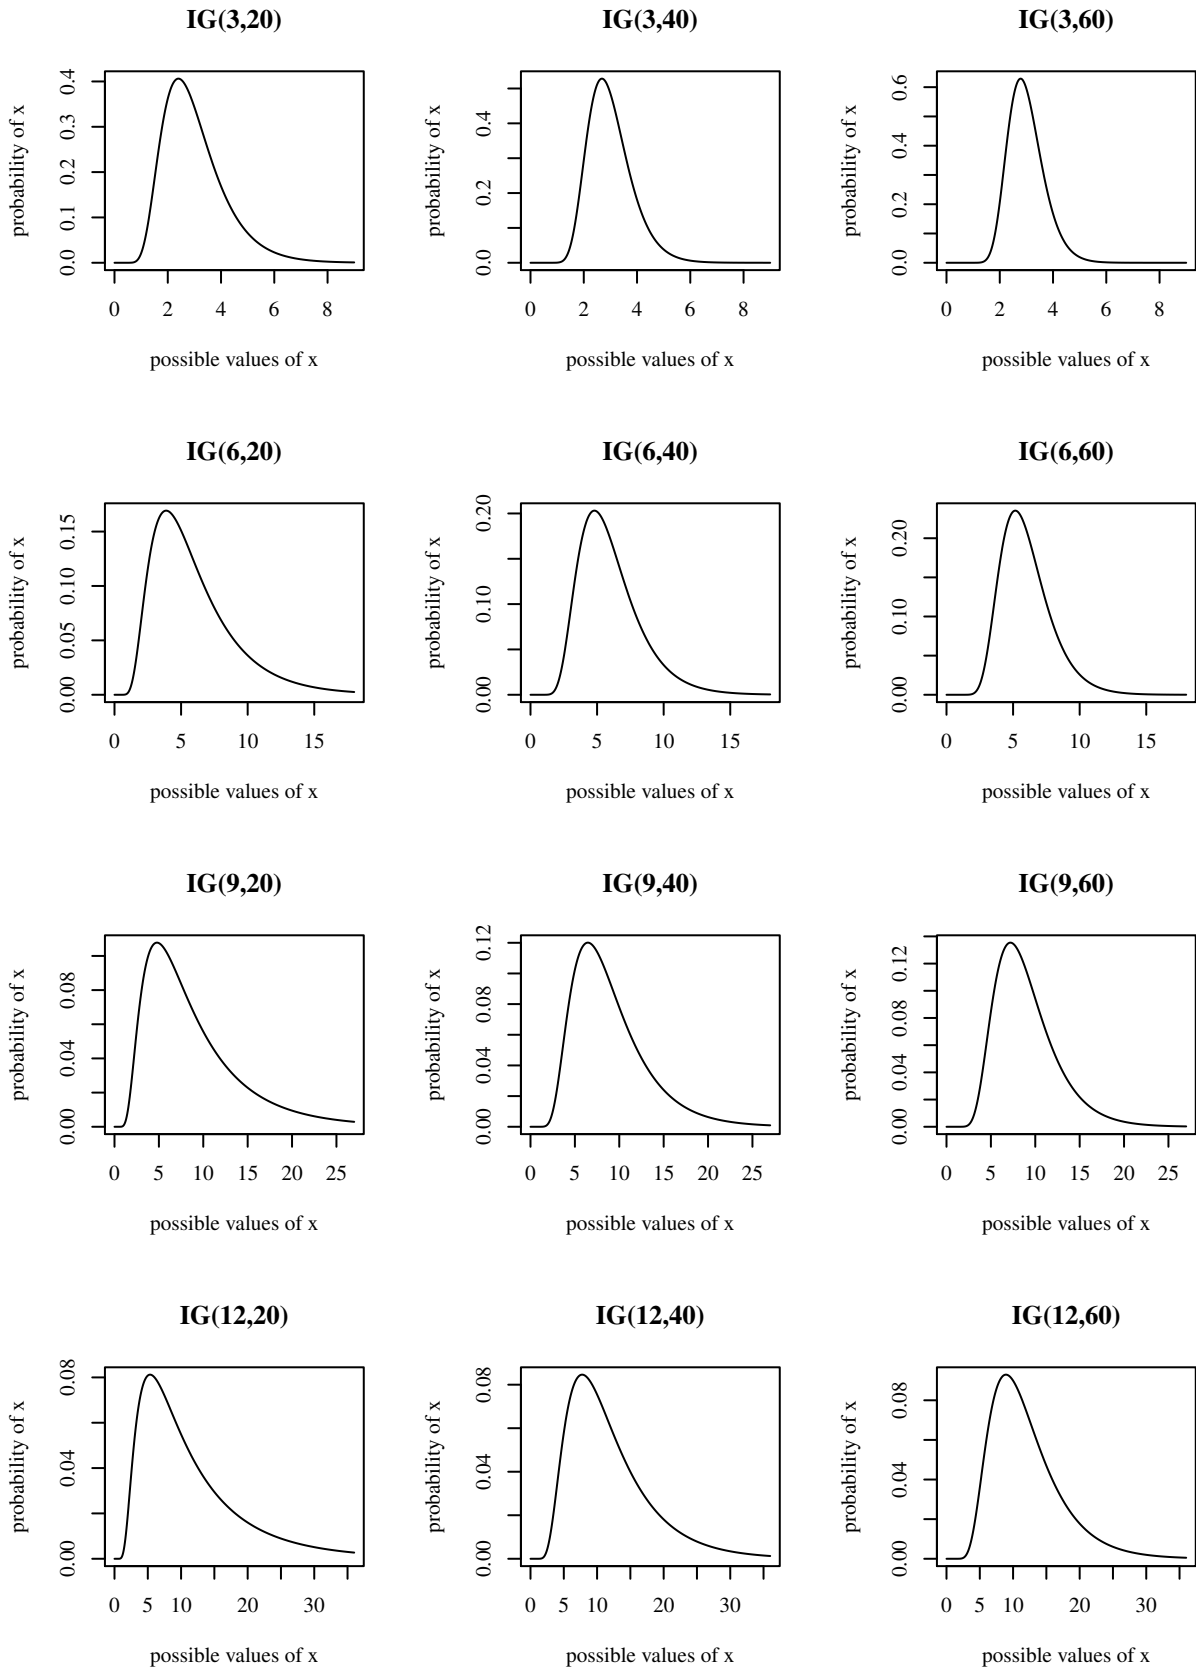


Figure 3: Sample from Inverse Gaussian with  $\mu = 12$  and  $\lambda = 2$

**Density plot, N=1000, mu= 12 lambda=2**






Figure 4: Random Sample From Various Inverse Gaussians

**mu= 12 lambda=2**



**mu= 12 lambda=5**



**mu= 12 lambda=20**





Figure 5: More Inverse Gaussian Distributions

**IG(2, various lambda)**



**IG(10, various lambda)**

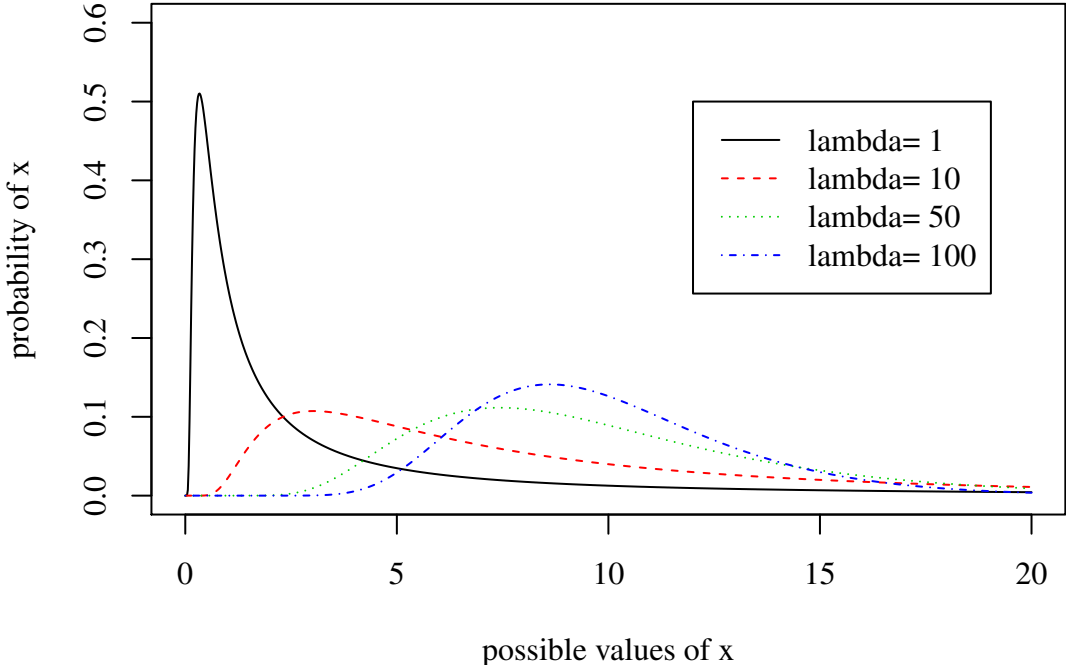
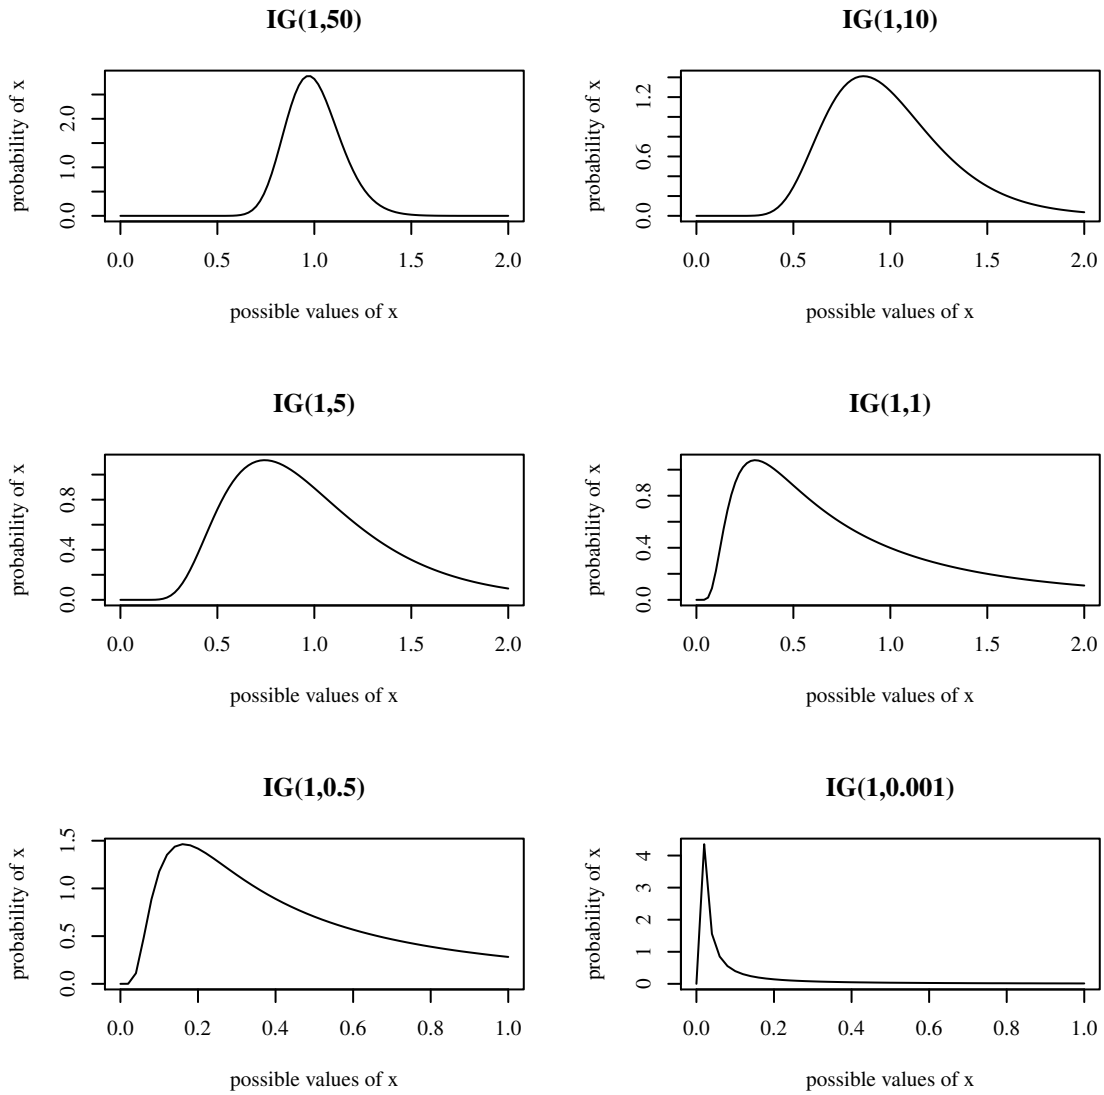




Figure 6: Shrinking lambda



## 4 Would you rather have Gamma or Inverse-Gaussian?

The Gamma and the Inverse-Gaussian share the property that they are possibly skewed to the right. If you choose the correct parameter values, you can make them practically indistinguishable.

However, there is a VERY WEIRD scaling property here. In order for the Inverse Gaussian to produce some extreme large values, it must have higher probability for large values. How can you make sense out of this strange result in Figure

Figure 7: Gamma or Inverse Gaussian?

```
xvals <- seq(0,20, length.out=200)
gam <- dgamma(xvals, shape=2, scale=1)
igaus <- dinvgauss(xvals, mu=2, lambda=5)
plot(xvals, gam, type="l", lty=1, main=" ")
lines(xvals, igaus, lty=2)
legend(6, .2, c("gamma(sh=2,sc=1)", "inv gauss(mu=2,lambda=5)"), lty
=1:2,)
```

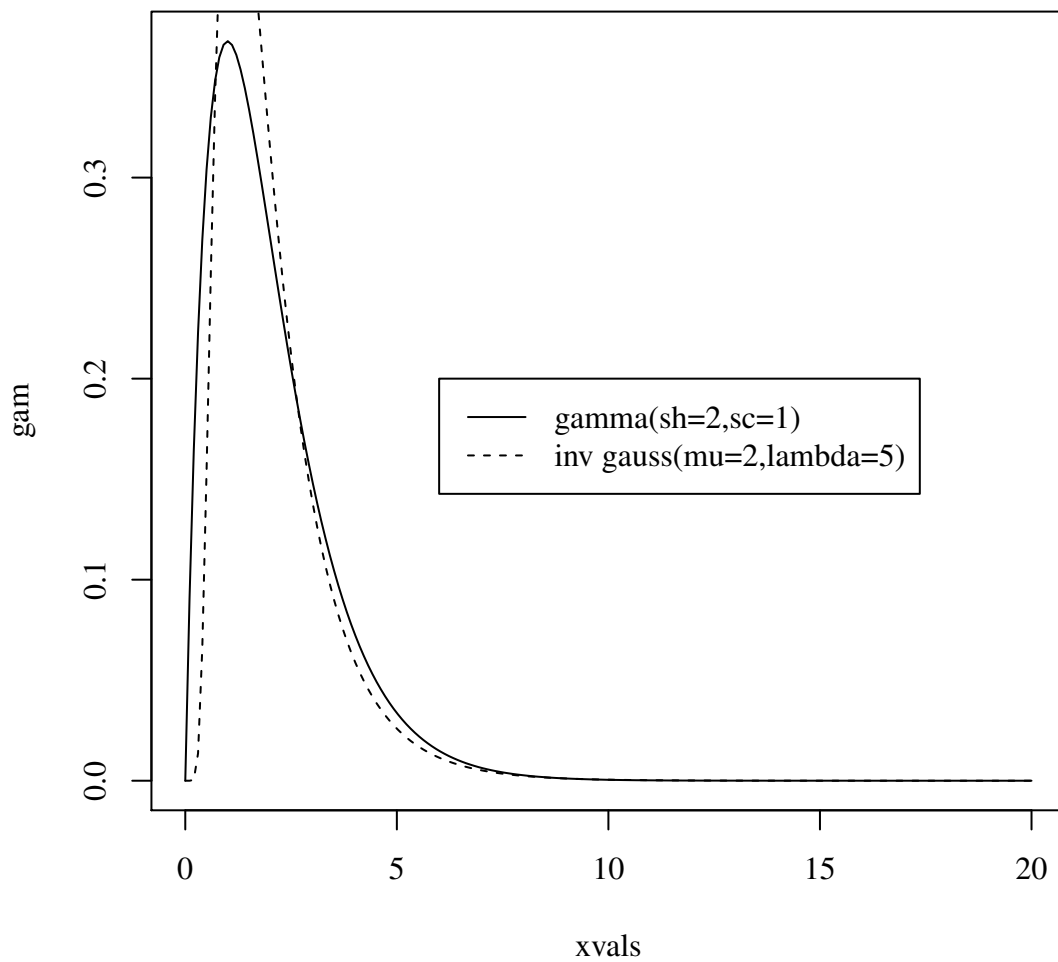


Figure 8: Compare Gamma and Inverse Gaussian

```

par(mfcol=c(3,1))
for (i in 1:3){
  minx <- 20 + 50 * (i-1)
  xvals <- seq(minx,300, length.out=1000)
  gam <- dgamma(xvals, shape = 2, scale = 1)
  igauss <- dinvgauss(xvals, mu = 2, lambda = 5)
  plot(xvals,gam, type="l", lty = 1, main="")
  lines(xvals, igauss, lty = 2)
  legend(150, 0.7*max(gam), c("gamma(sh=2,sc=1)", "inv gauss(
    mu=2,lambda=5)"), lty = 1:2,)
}

```






Figure 9: Compare Samples: Gamma and Inverse Gaussian


Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.0502	0.9770	1.7240	2.0160	2.7520	11.5400

[1] 2.067561

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.2959	1.1160	1.7240	2.0410	2.5270	10.5000

[1] 1.831236

**Gamma,sh=2,sc=1**



**Inv Gaus, mu=2,lambda=4**




Figure 10: Compare Samples: Gamma and Inverse Gaussian


Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.1192	1.9150	3.3560	4.0570	5.3900	22.7900

[1] 8.55414

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.648	2.062	3.125	3.825	4.802	18.150

[1] 6.715464

**Gamma, sh=2, sc=2**



**Inv Gauss, mu=4, lambda=8**

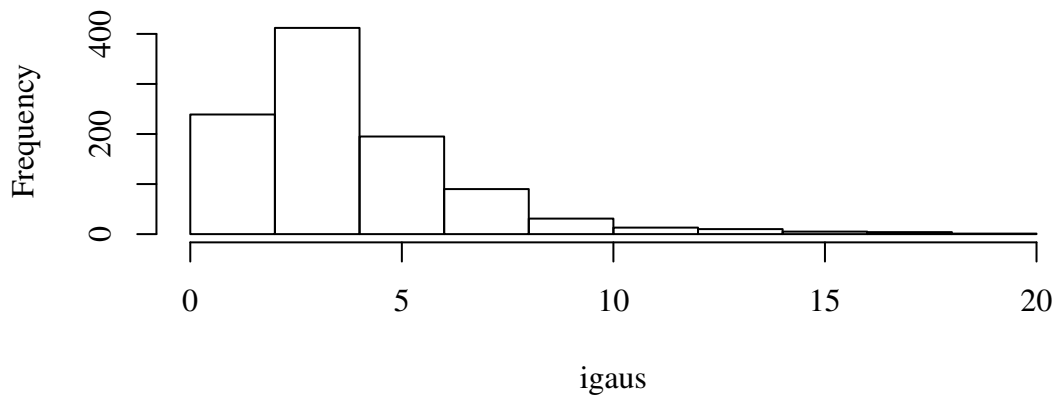


Figure 11: Compare Samples: Gamma and Inverse Gaussian


Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.709	6.868	10.600	11.860	15.730	42.260

[1] 43.43887

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
2.472	7.040	10.130	12.220	15.140	62.900

[1] 60.23548

**Gamma, sh=3,sc=4**



**Inv Gauss, mu=12,lambda=34**

