MATH 559: BAYESIAN THEORY AND METHODS

SLICE SAMPLING

Suppose that

$$\pi(x) = \frac{\exp\{-h(x)\}}{\displaystyle\int_{-\infty}^{\infty} \exp\{-h(t)\}\, dt} \propto \exp\{-h(x)\}$$

is a pdf to be sampled. The slice sampler for this problem writes

$$\pi_{X,U}(x,u) = \frac{\mathbb{1}\{u \leq \exp\{-h(x)\}\}}{\displaystyle\int_{-\infty}^{\infty} \exp\{-h(t)\}\, dt}$$

for auxiliary variable U defined on $(0,1)$. Note that the marginal for X is π.

$$\int_0^1 \pi_{X,U}(x,u)\, du = \frac{1}{\displaystyle\int_{-\infty}^{\infty} \exp\{-h(t)\}\, dt} \int_0^{\exp\{-h(x)\}} du = \frac{\exp\{-h(x)\}}{\displaystyle\int_{-\infty}^{\infty} \exp\{-h(t)\}\, dt} = \pi(x)$$

For this model we have the following full conditionals.

- $\pi_{U|X}(u|x) \equiv Uniform(0, \exp\{-h(x)\})$
- $\pi_{X|U}(x|u)$ is uniform on the segment of $\mathbb{R}$ defined by

$$h(x) \leq -\log u$$

For initialization of the corresponding Gibbs sampler, we need to have a configuration $(x^{(0)}, u^{(0)})$ for which $\pi_{X,U}(x^{(0)}, u^{(0)}) > 0$.

EXAMPLE: Sampling from the Normal distribution
Suppose we wish to produce a sample from the pdf

$$\pi(x) \propto \exp\{-x^2/2\} \qquad x \in \mathbb{R}$$

which is a standard distribution that can be sampled directly. We write

$$\pi(x) \propto \exp\{-h(x)\}$$

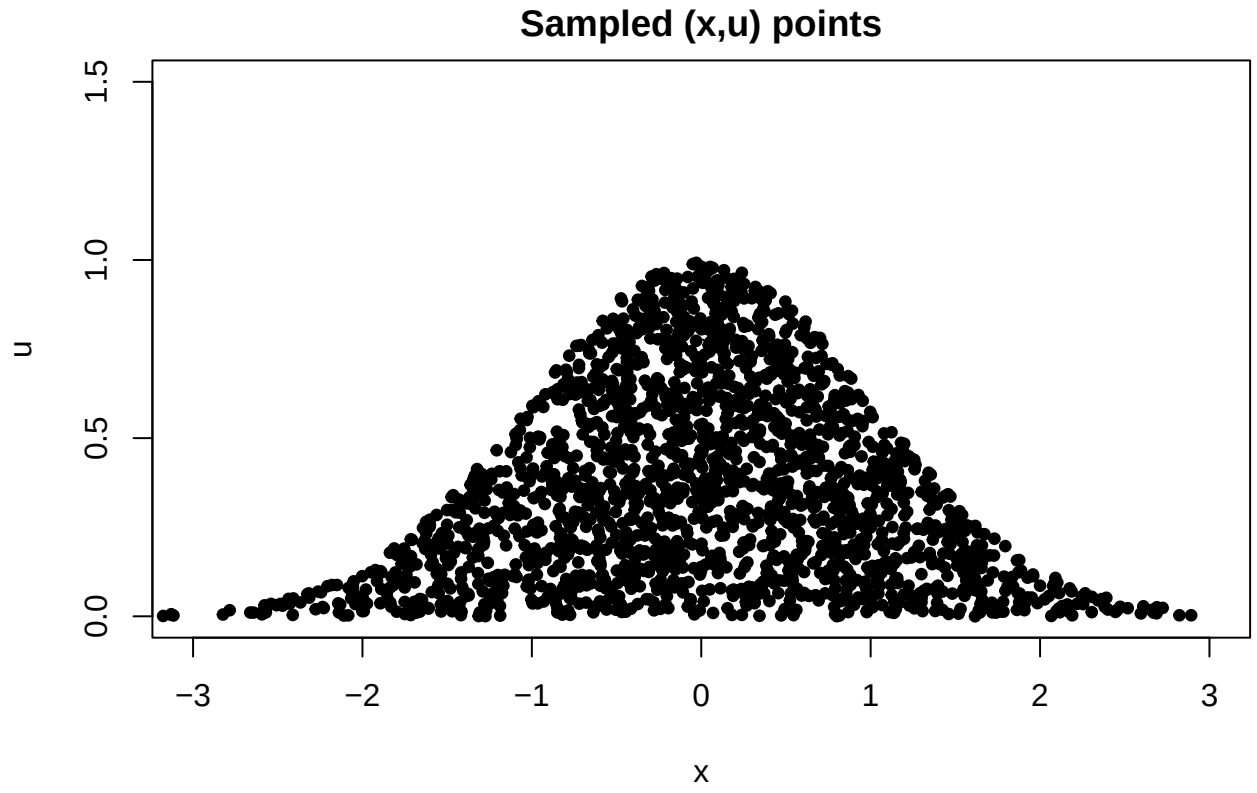
with $h(x) = x^2/2$. Then

- $\pi_{U|X}(u|x) \equiv Uniform(0, \exp\{-h(x)\})$
- $\pi_{X|U}(x|u)$ is uniform on the segment defined by

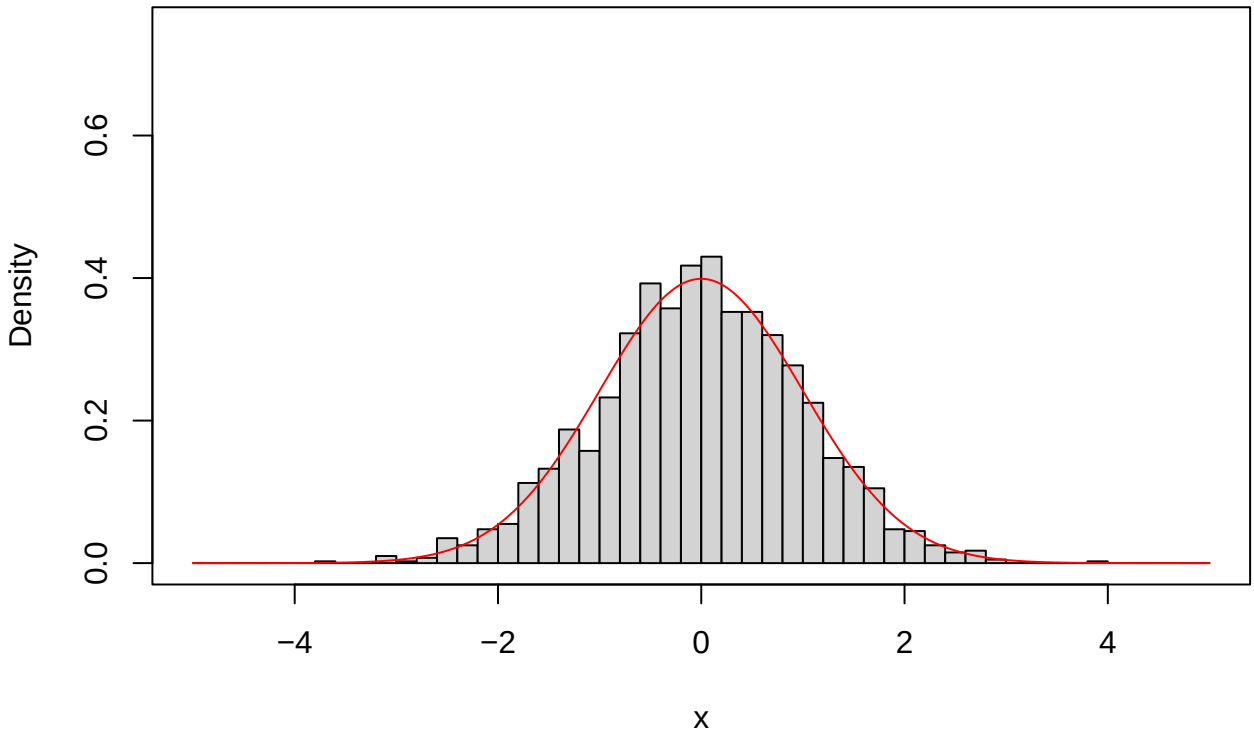
$$x^2 \leq -2\log u \qquad \Longleftrightarrow \qquad -\sqrt{-2\log u} \leq x \leq \sqrt{-2\log u}$$

```
x<-0
h<-function(x){return(x^2/2)}
u<-runif(1,0,exp(-h(x)))
par(mar=c(4,4,2,2))
plot(x,u,xlim=range(-3,3),ylim=range(0,1.5))
nits<-2000
xsamp<-rep(0,nits)
```

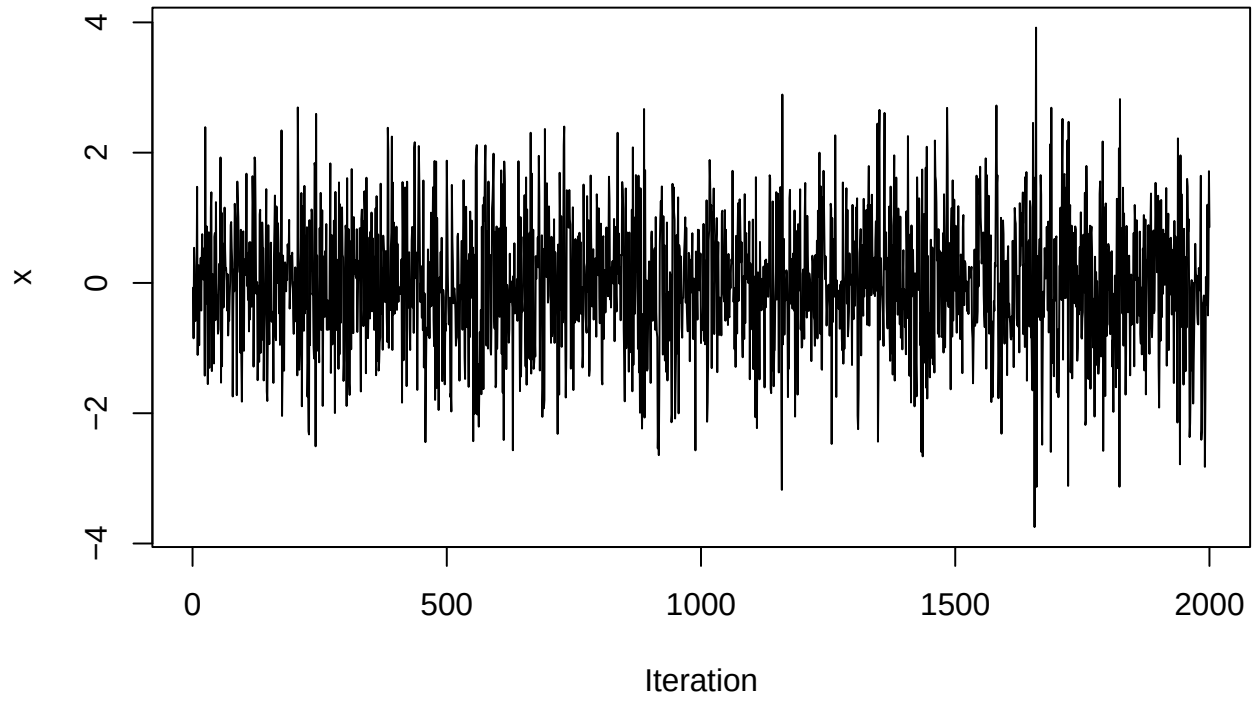
```
for(iter in 1:nits){
        x<-runif(1,-sqrt(-2*log(u)),sqrt(-2*log(u)))
        u<-runif(1,0,exp(-h(x)))
        points(x,u,pch=19,cex=0.75)
        xsamp[iter]<-x
}
title('Sampled (x,u) points')
```



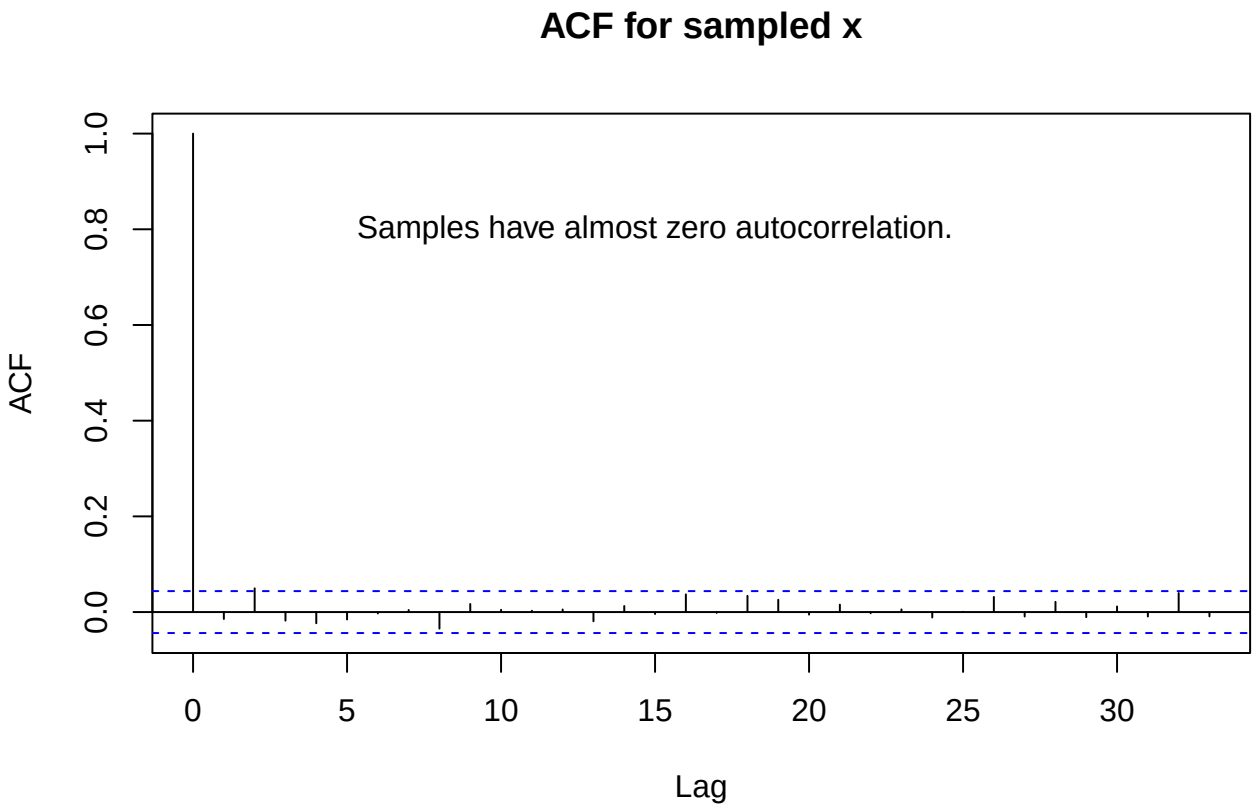

```
hist(xsamp,freq=FALSE,br=seq(-5,5,by=0.2),
        ylim=range(0,0.75),main='',xlab='x');box()
xv<-seq(-5,5,by=0.01);lines(xv,dnorm(xv),col='red')
```




```
par(mar=c(4,4,3,2))
plot(xsamp,type='l',xlab='Iteration',ylab='x')
```



```r
acf(xsamp,main='ACF for sampled x');text(15,0.8,'Samples have almost zero autocorrelation.')
```



EXAMPLE: Sampling an awkward distribution
Suppose we wish to produce a sample from the pdf

$$\pi(x) \propto x^{\alpha-1} \exp\left\{-\left[\beta x + \frac{\gamma}{x}\right]\right\} \qquad x > 0$$

for parameters $\alpha, \beta, \gamma > 0$. The normalizing constant is not trivial to compute in general. Here

$$h(x) = -(\alpha - 1)\log x + \beta x + \frac{\gamma}{x}$$

say. Note that the minimizing value of $h(x)$ solves

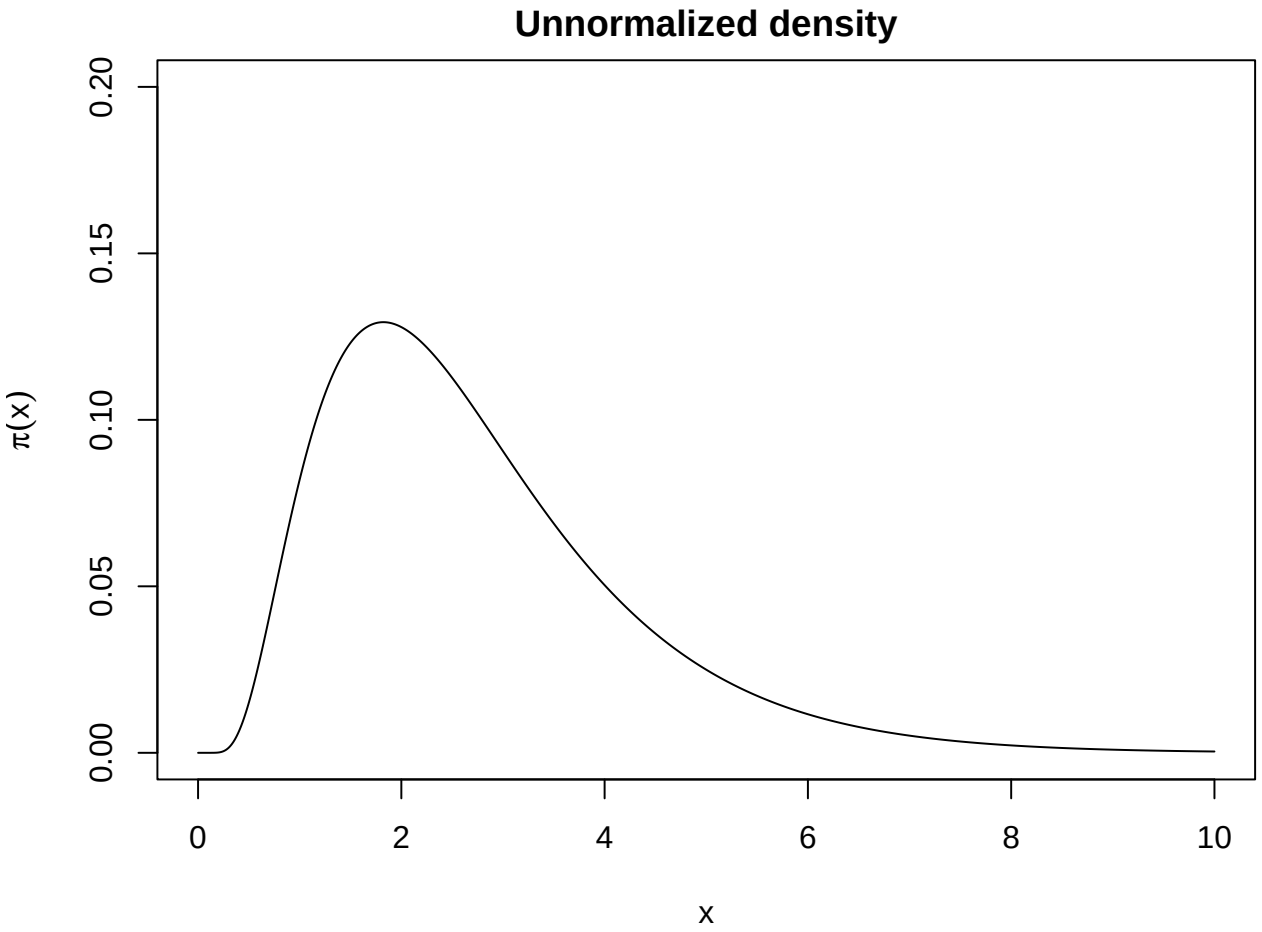
$$-\frac{(\alpha-1)}{x} + \beta - \frac{\gamma}{x^2} = 0$$

that is, the quadratic

$$\beta x^2 - (\alpha - 1)x - \gamma = 0$$

and we need the positive root, $x_{\min} = ((\alpha-1) + \sqrt{(\alpha-1)^2 + 4\beta\gamma})/(2\beta)$

```r
al<-2; be<-1; gam<-1.5
#Positive root
hminx<-((al-1)+sqrt((al-1)^2+4*be*gam))/(2*be)
xv<-seq(0,10,by=0.001)
hxv<--(al-1)*log(xv)+be*xv+gam/xv
yv<-exp(-hxv)
par(mar=c(4,4,2,2))
plot(xv,yv,type='l',xlab='x',ylab=expression(pi(x)),ylim=range(0,0.2))
title('Unnormalized density')
```

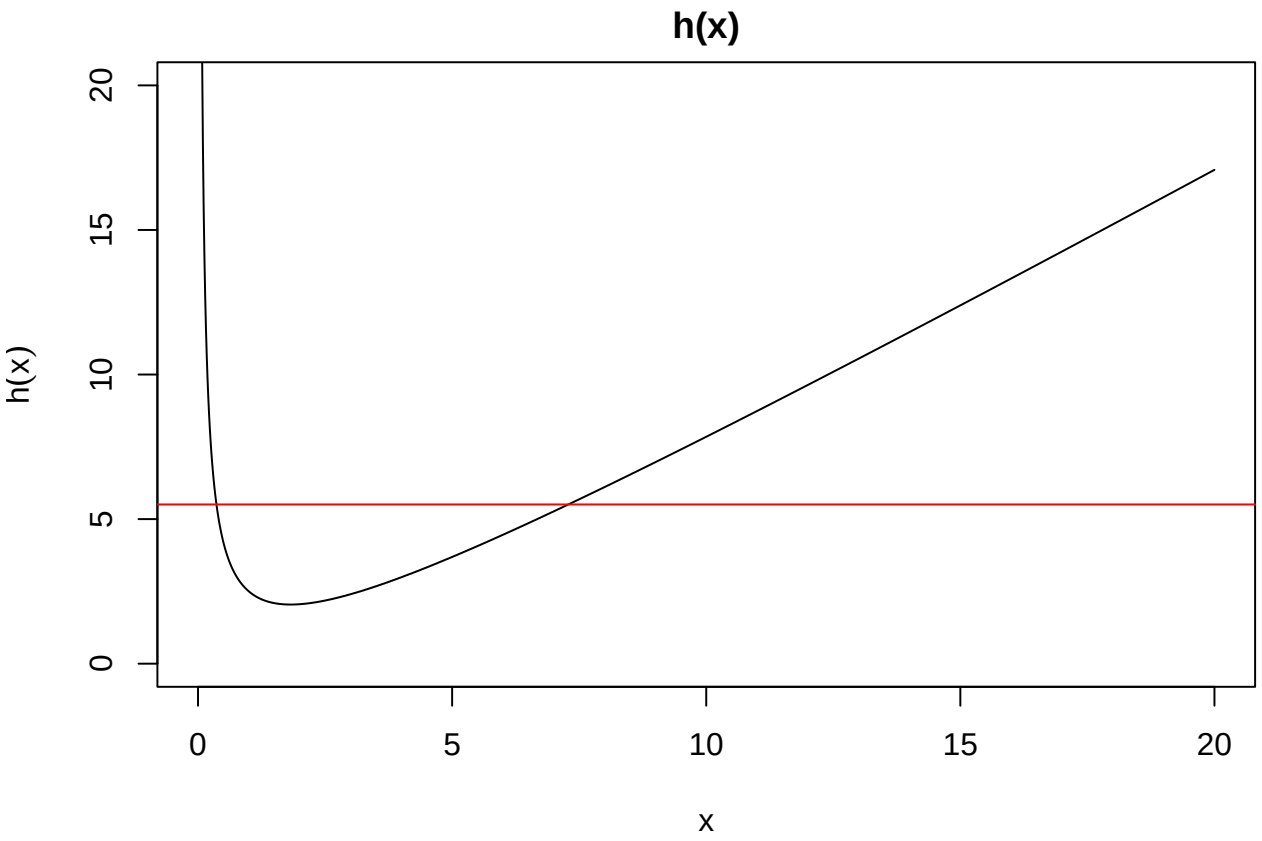


For this model we have the following full conditionals.

- $\pi_{U|X}(u|x) \equiv Uniform(0, \exp\{-h(x)\})$
- $\pi_{X|U}(x|u)$ is uniform on the segment defined by $h(x) \leq -\log u$. The segment is an interval in this case.

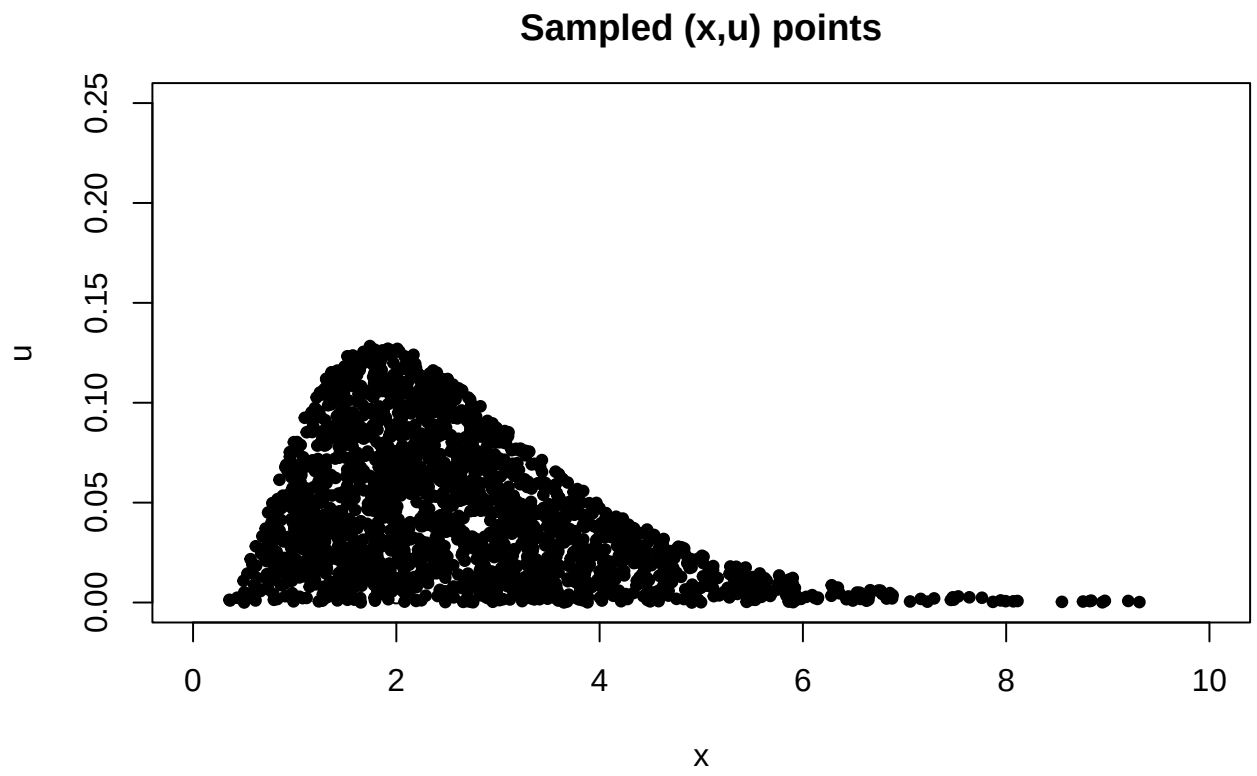
For example, suppose that $h(x) = 4$: we then sample $U \sim Uniform(0, \exp(-4))$

```
set.seed(9)
xv<-seq(0.001,20,by=0.001)
hxv<--(al-1)*log(xv)+be*xv+gam/xv
u<-runif(1,0,exp(-4))
par(mar=c(4,4,2,2))
plot(xv,hxv,type='l',xlab='x',ylab=expression(h(x)),ylim=range(0,20))
title('h(x)')
abline(h=-log(u),col='red')
```

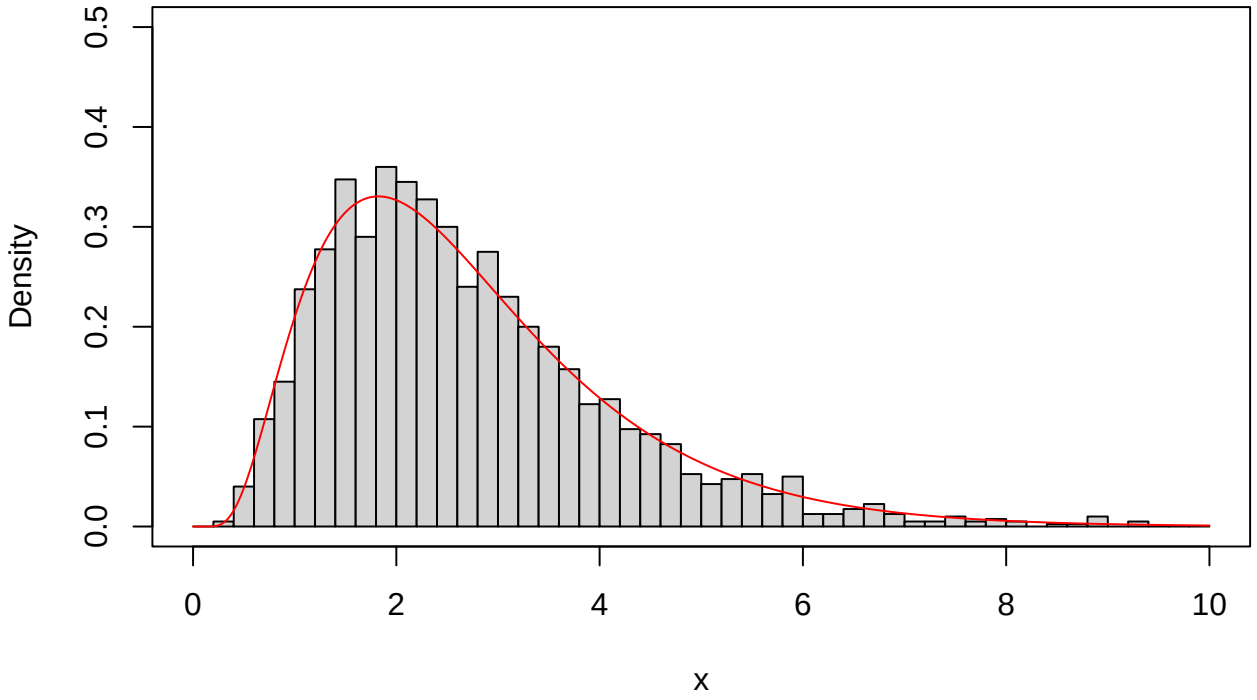


The value of U generated is $u = 0.004060$, and $-\log u = 5.506880$. We then generate X conditional on $U = u$ uniformly on the line segment connecting the two points where the horizontal line at $-\log u$ intersects $h(x)$. We may use the root-finding function `uniroot` and search either side of the minimizing x value $x_{\min}$.

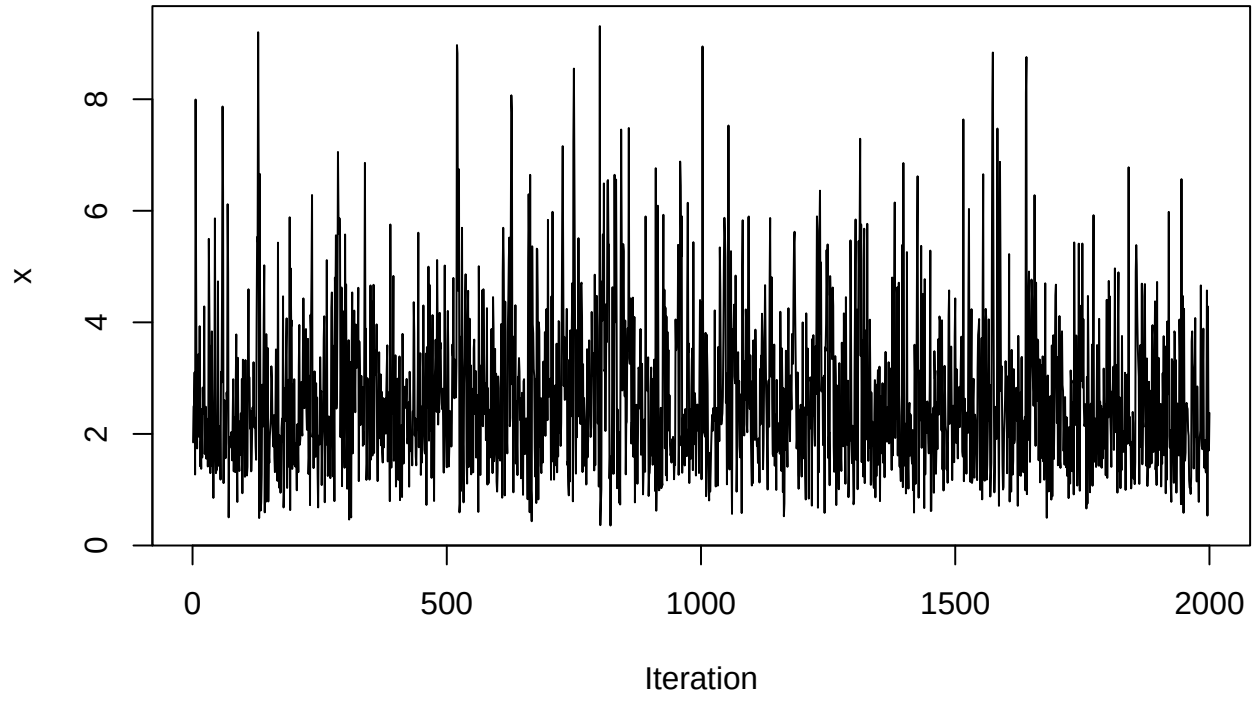
```
fn<-function(x,a,b,g,u){
    return(-(a-1)*log(x)+b*x+g/x+log(u))
}
sample.x<-function(a,b,g,u){
    lower<-uniroot(fn,interval=c(0,hminx),a=a,b=b,g=g,u=u)$root
    upper<-uniroot(fn,interval=c(hminx,100),a=a,b=b,g=g,u=u)$root
    return(runif(1,lower,upper))
}
x<-2
h<-function(x,a,b,g){return(-(a-1)*log(x)+b*x+g/x)}
u<-runif(1,0,exp(-h(x,al,be,gam)))
par(mar=c(4,4,3,2))
plot(x,u,xlim=range(0,10),ylim=range(0,0.25))
nits<-2000
xsamp<-rep(0,nits)
for(iter in 1:nits){
        x<-sample.x(al,be,gam,u)
        u<-runif(1,0,exp(-h(x,al,be,gam)))
        points(x,u,pch=19,cex=0.75)
        xsamp[iter]<-x
}
title('Sampled (x,u) points')
```



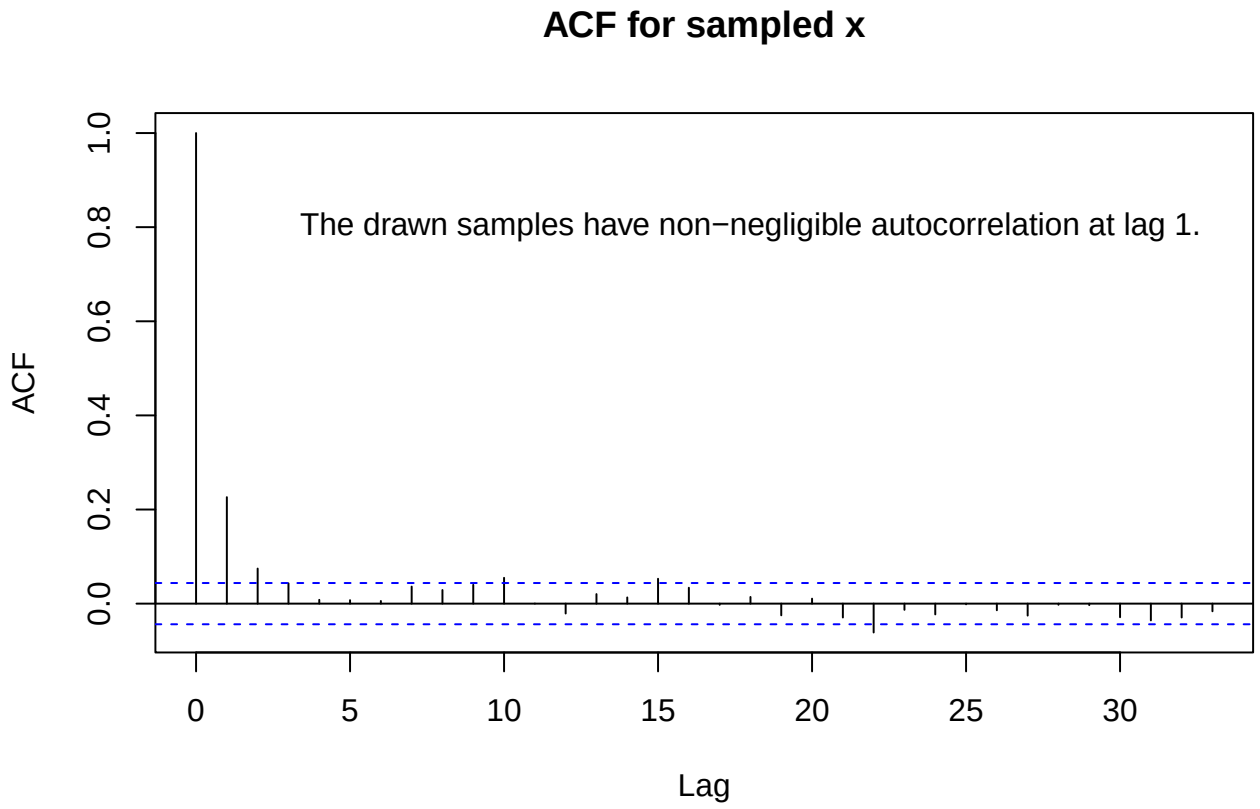

```
library(GeneralizedHyperbolic) #This distribution is the Generalized Inverse Gaussian
hist(xsamp,freq=FALSE,br=seq(0,10,by=0.2),
        ylim=range(0,0.5),main='',xlab='x');box()
xv<-seq(0,10,by=0.01);lines(xv,dgig(xv,lam=al,chi=2*gam,psi=2*be),col='red')
```




```
plot(xsamp,type='l',xlab='Iteration',ylab='x')
```



```
acf(xsamp,main='ACF for sampled x')
text(18,0.8,'The drawn samples have non-negligible autocorrelation at lag 1.')
```



The effective sample size is a measure of the extent to which the autocorrelation affects the variance of any estimator based on the drawn samples. Here we have run the Gibbs sampler for $N_{iter} = 2000$ iterations:

```
library(coda)
effectiveSize(xsamp)

##     var1
## 1261.264
```

This result tells us that instead of having a sample of size N_{iter} to estimate distribution quantities, the autocorrelation reduces the size to $N_{eff} = 1261.264000$.