Computational Statistics. Chapter 5: MCMC. Solution of exercises

Thierry Denoeux

2024-11-14

```
set.seed(2021)
```

Exercise 1

Question a

As the density of ϵ is symmetric, the MH ratio is the ratio of the densities at x^* and $x^{(t-1)}$, i.e., we have

$$R(x^{(t-1)}, x^*) = \frac{f(x^*)}{f(x^{(t-1)})} = \exp(|x^{(t-1)}| - |x^*|).$$

The following function `MH_Laplace` implements the random walk MH algorithm for this problem. It returns a chain of length N and the acceptance rate

```
RW_Laplace <- function(N,sig){
  x<-vector(N,mode="numeric")
  x[1]<-rnorm(1,mean=0,sd=sig)
  accept <- 0
  for(t in (2:N)){
    epsilon<-rnorm(1,mean=0,sd=sig)
    xstar<-x[t-1]+ epsilon
    U<-runif(1)
    R<-exp(abs(x[t-1]) - abs(xstar))
    if(U < R){ #accept the candidate value
      x[t]<-xstar
      accept <- accept + 1
      } else x[t]<-x[t-1]
  }
  accept<-accept/(N-1)
  return(list(x=x,accept=accept))
}
```

Let us generate a sample of size 10^5 with $\sigma = 10$:

```
res1 <- RW_Laplace(100000,10)
print(res1$accept)
```

```
## [1] 0.1535215
```

We can see that tbe acceptance rate is a bit low. Let us make the proposal distribution more focussed by setting $\sigma = 2$:

```r
res1 <- RW_Laplace(100000,2)
print(res1$accept)
```

```
## [1] 0.5236052
```

This time, the acceptance rate is close to 50/%. The sample path and correlation plots show good mixing (the chain quickly moves away from its starting value, and the autocorrelation decreases quickly as the lag between iterations increases):

```r
plot(res1$x,type="l",xlab='t',ylab=expression(x[(t)]))
```

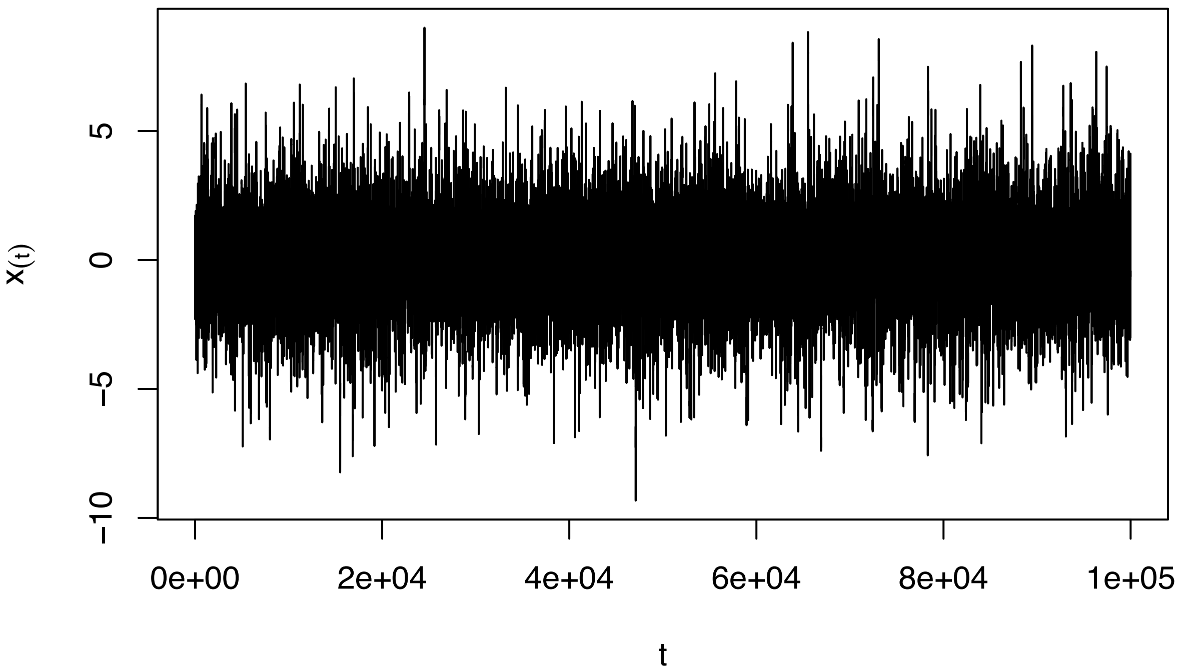


```r
acf(res1$x,lag.max=100)
```

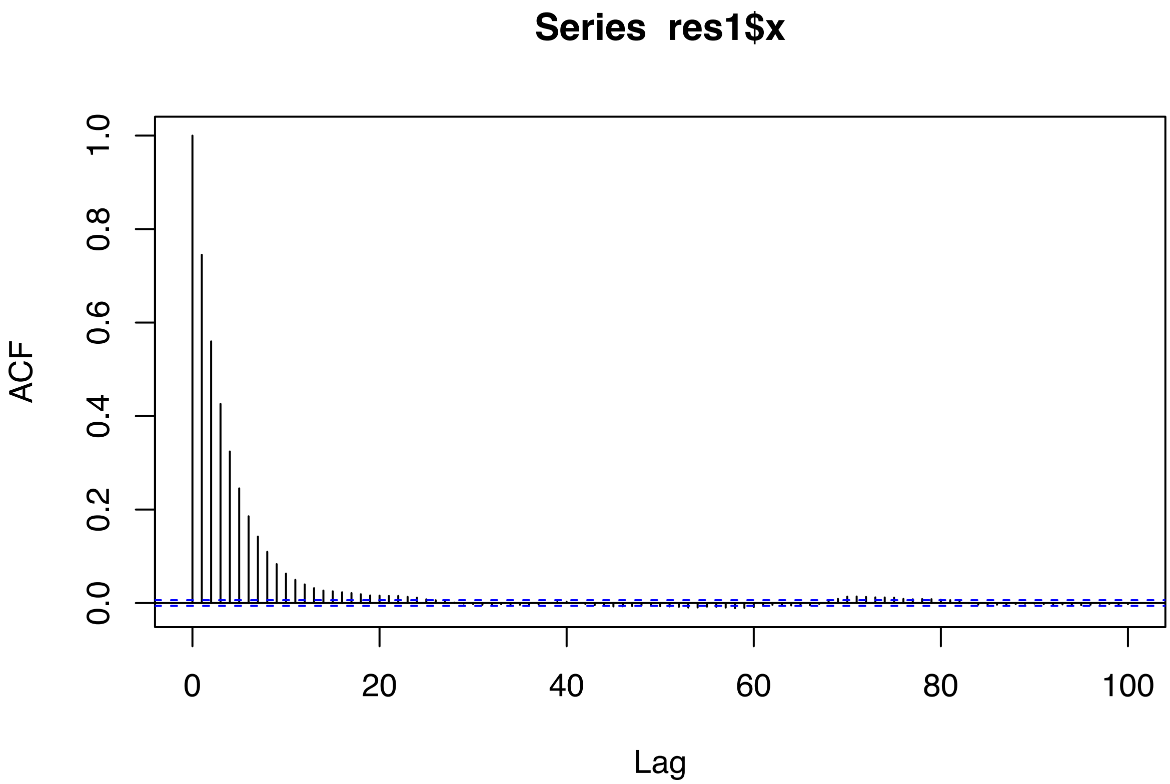

Plot of the histogram with the Laplace density:

```
u<-seq(-10,10,0.01)
fu<-0.5*exp(-abs(u))
hist(res1$x,breaks=50,freq=FALSE,ylim=range(fu))
lines(u,0.5*exp(-abs(u)))
```

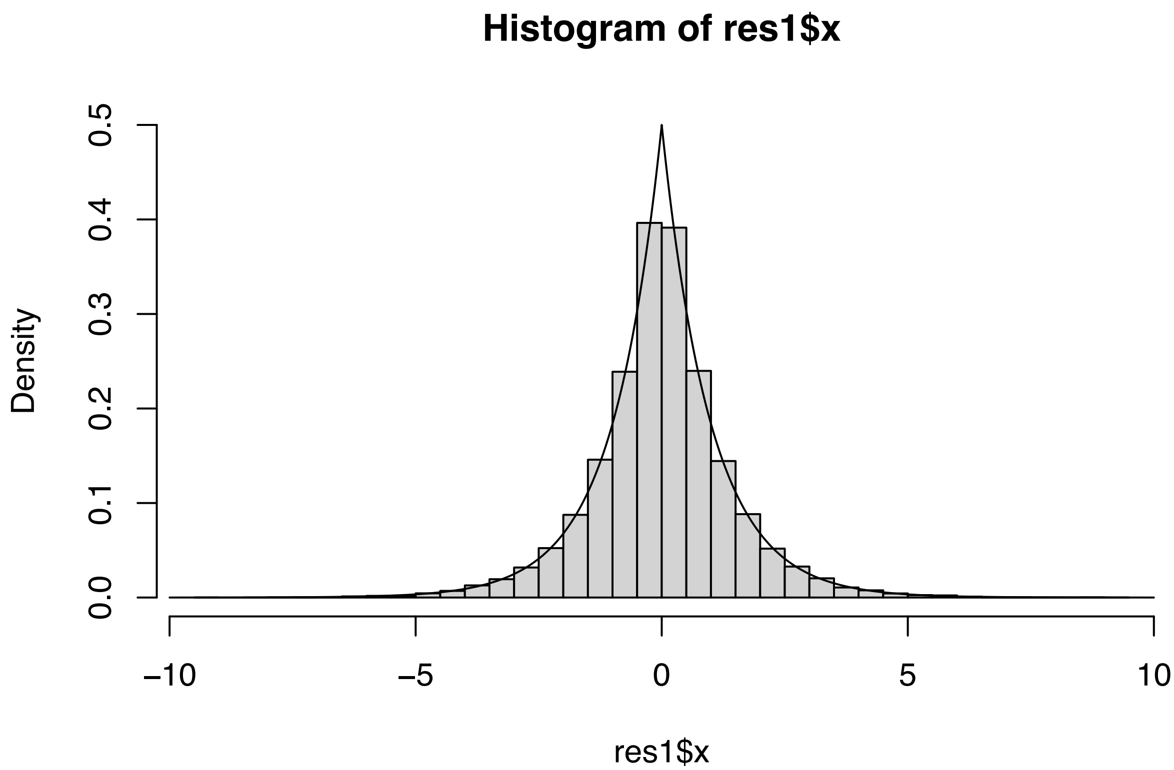

Let us now generate another sample of the same size, this time with $\sigma = 0.1$:

```
res11<-RW_Laplace(100000,0.1)
print(res11$accept)
```

```
## [1] 0.9610096
```

This time, the sample path and correlation plots show poor mixing (the chain remains at or near the same value for many iterations, and the autocorrelation decays very slowly):

```
plot(res11$x,type="l",xlab='t',ylab=expression(x[(t)]))
```

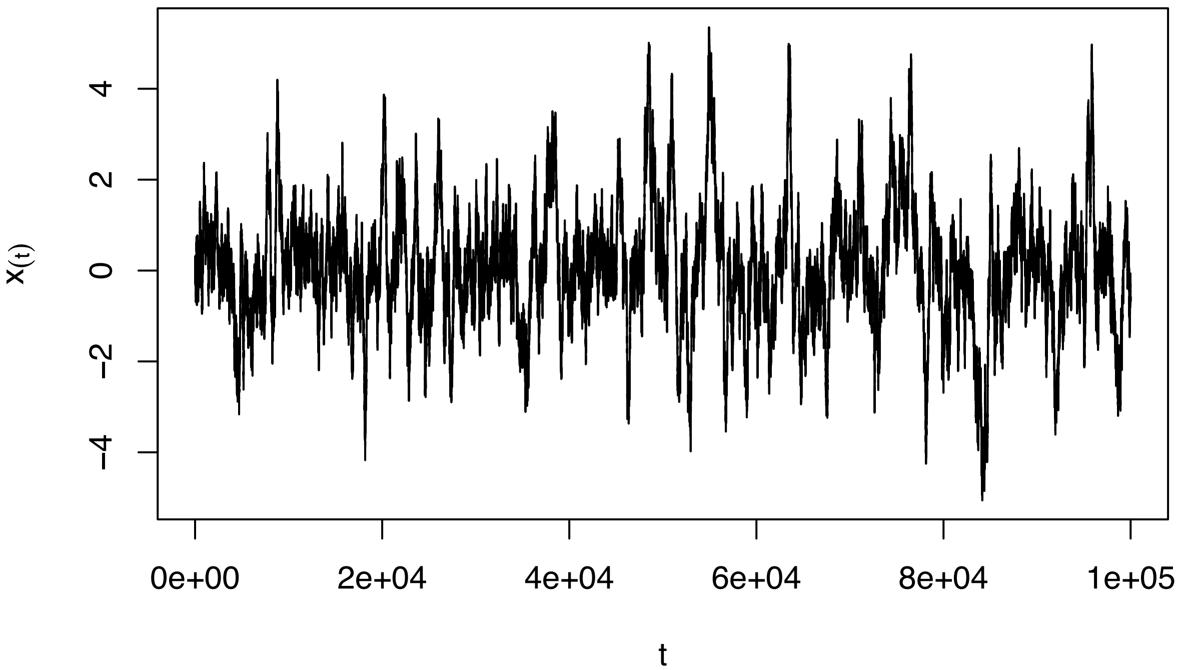


```
acf(res11$x,lag.max=1000)
```

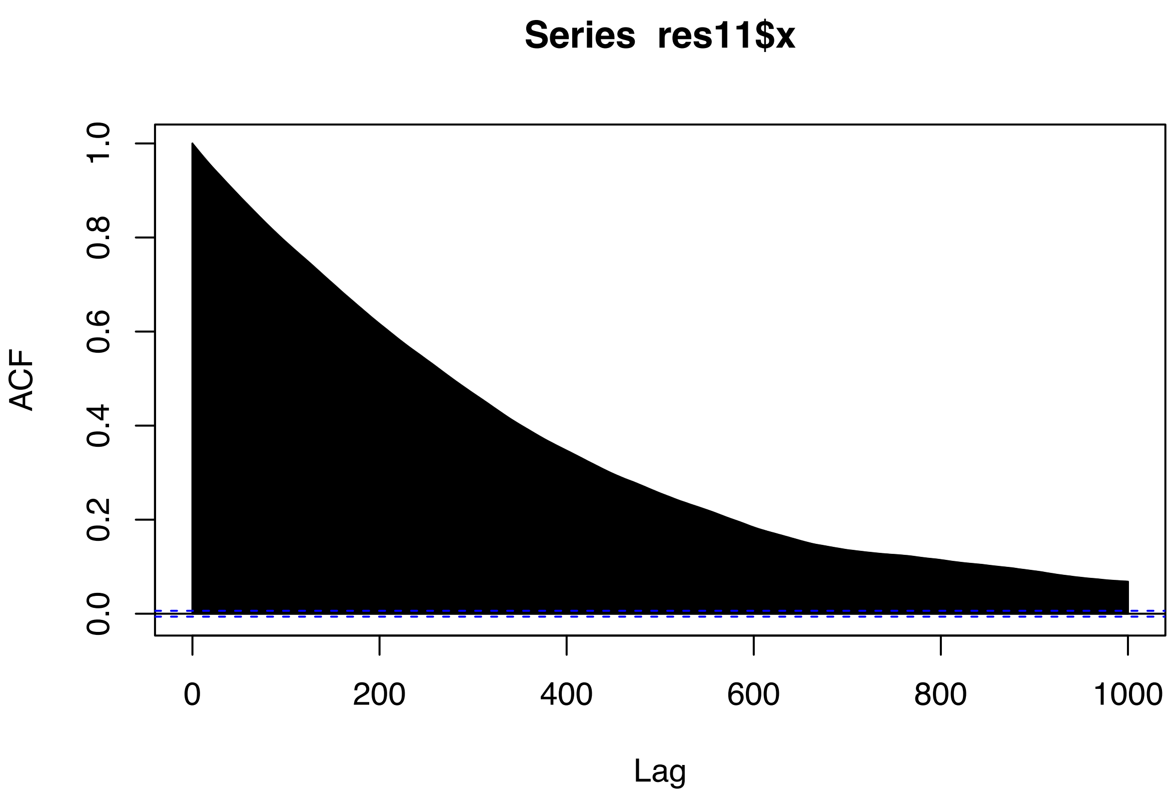

Question b

We can check that the standard Cauchy density is somewhat similar to the Laplace density, but is has heavier tails:

```
x<-seq(-5,5,0.01)
f<-0.5*exp(-abs(x))
g<-dcauchy(x)
plot(x,f,type="l",col="blue")
lines(x,g,col="red")
```

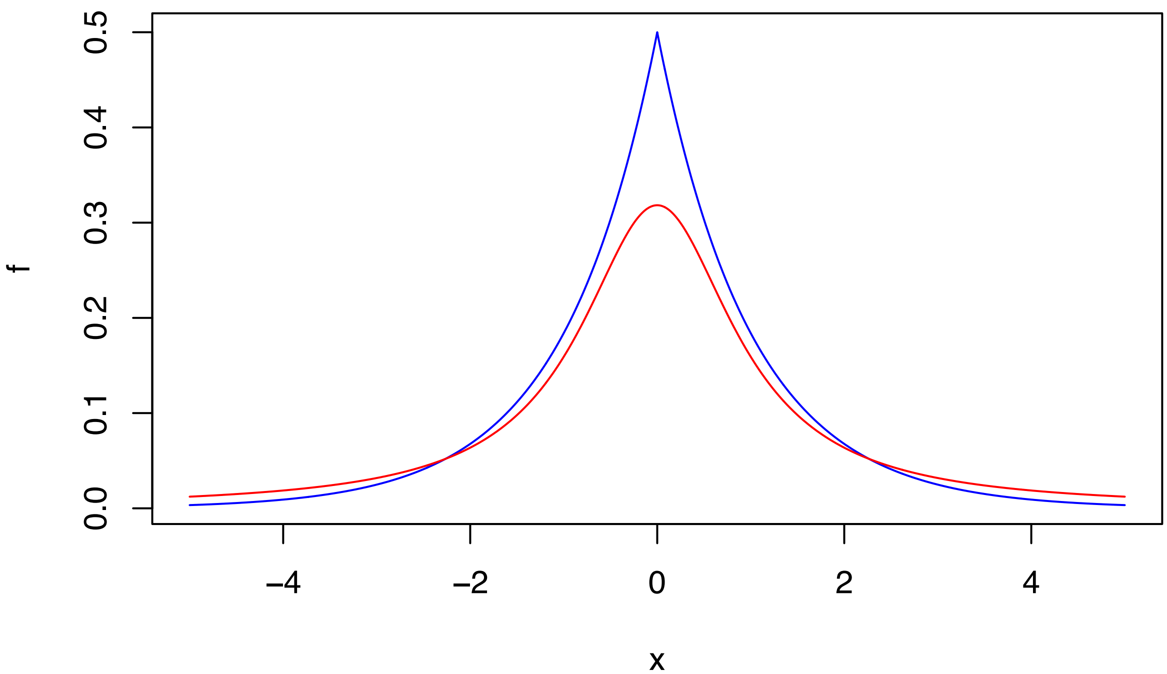

Let us write a function that implements an independence chain sampler:

```
IC_Laplace<-function(N,s=1){
  x<-rep(0,N)
  x[1]<-rcauchy(1,scale=s)
  accept<-0
  for(t in (2:N)){
    xstar<-rcauchy(1,scale=s)
    U<-runif(1)
    R <-  exp(-abs(xstar))*dcauchy(x[t-1],scale=s)/
          (exp(-abs(x[t-1]))*dcauchy(xstar,scale=s))
    if(U < R){
      x[t]<-xstar
      accept<-accept+1
    } else x[t]<-x[t-1]
  }
  accept<-accept/(N-1)
  return(list(x=x,accept=accept))
}
```

Let us run this algorithm to generate a chain of length $N = 10^5$:

```
res2<-IC_Laplace(10^5)
res2$accept
```

```
## [1] 0.7882079
```

The generated chain has good mixing:

```r
plot(res2$x,type="l",xlab='t',ylab=expression(x[(t)]))
```

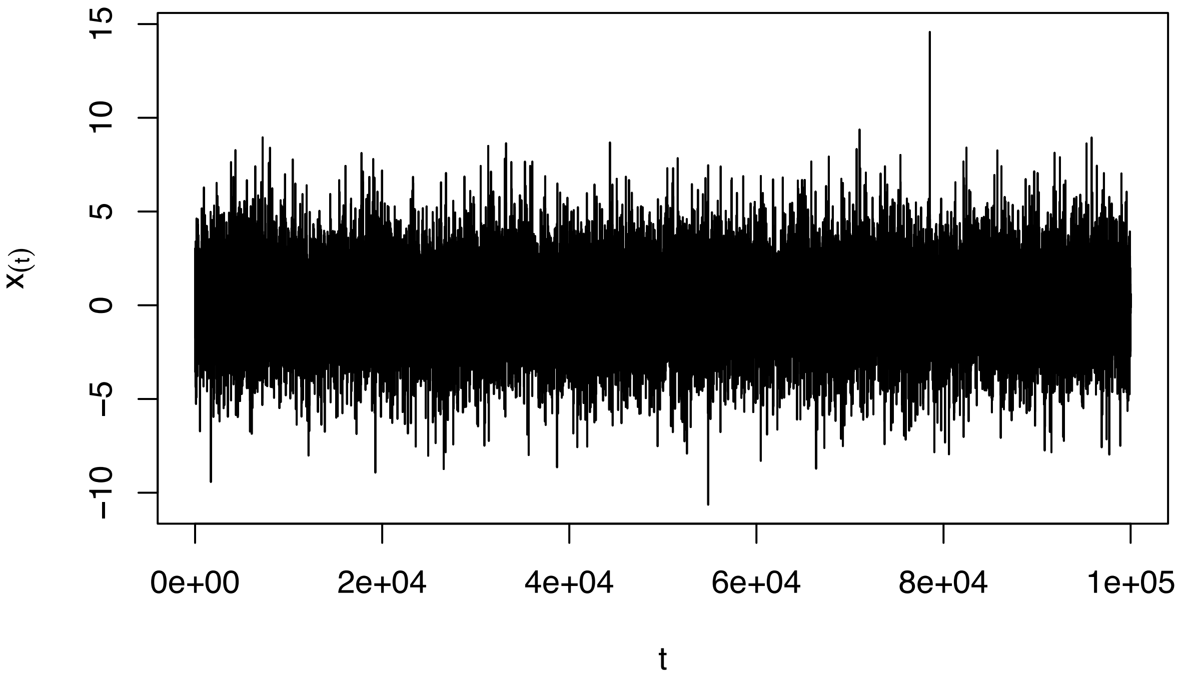


```r
acf(res2$x,lag.max=1000)
```

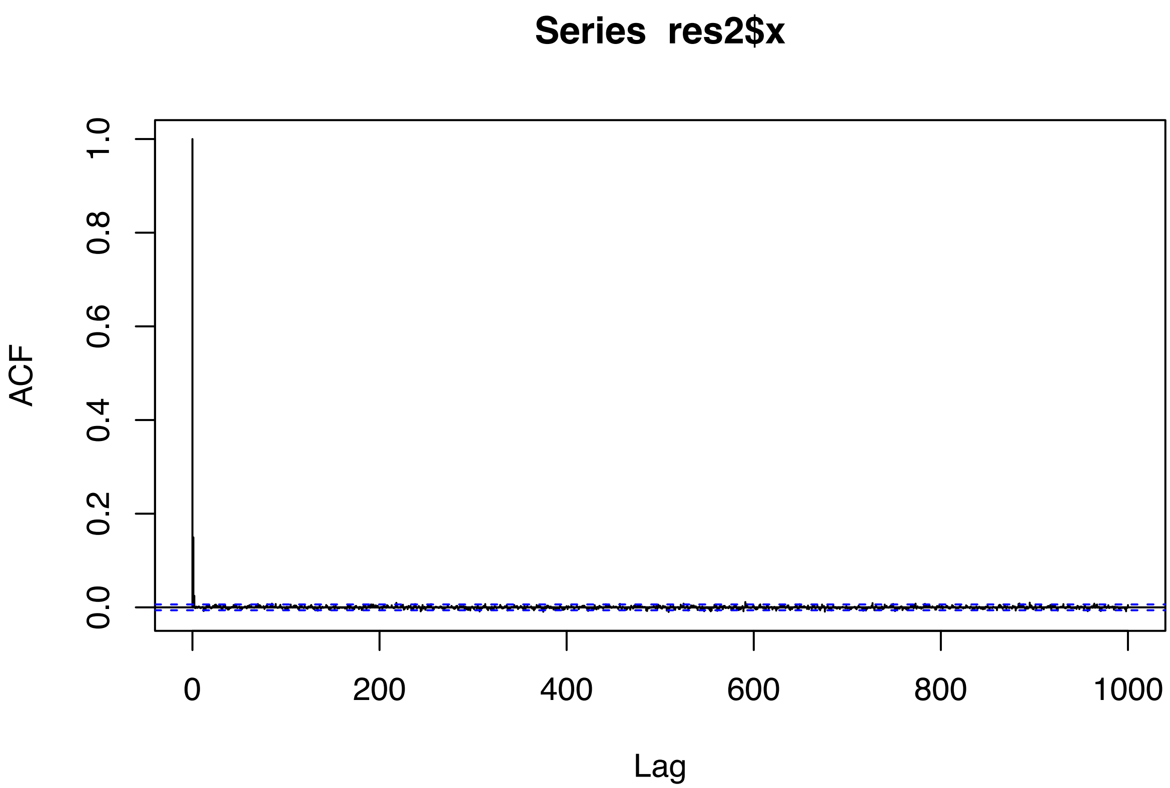

Plot of the histogram with the Laplace density:

```r
u<-seq(-10,10,0.01)
fu<-0.5*exp(-abs(u))
hist(res2$x,breaks=50,freq=FALSE,ylim=range(fu))
```

```
lines(u,0.5*exp(-abs(u)))
```

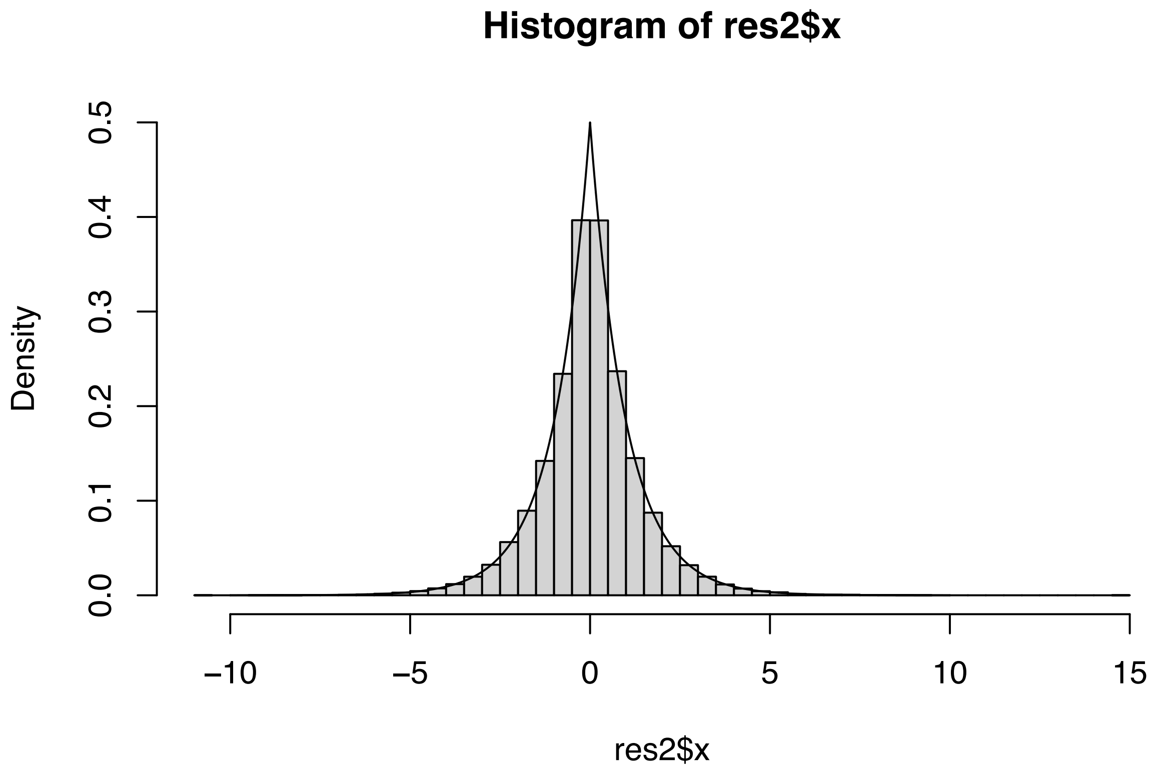

Question c

Efficient sample size calculation for the random walk sampler:

```
rho1<-acf(res1$x,lag.max=100)$acf
```

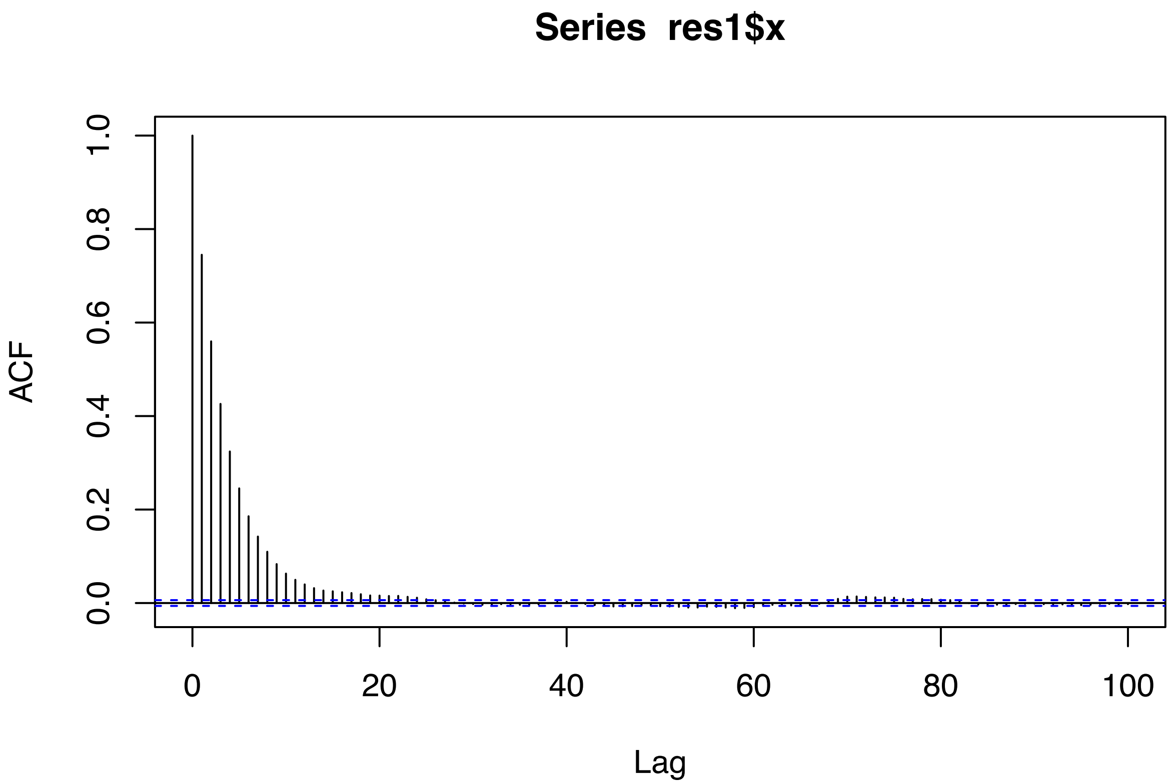


```
tau1<-1+2*sum(rho1[which(rho1>=0.1)])
print(10^5/tau1)
```

```
## [1] 11798.02
```

Efficient sample size calculation for the independence chain sampler:

```
rho2<-acf(res2$x,lag.max=100)$acf
```

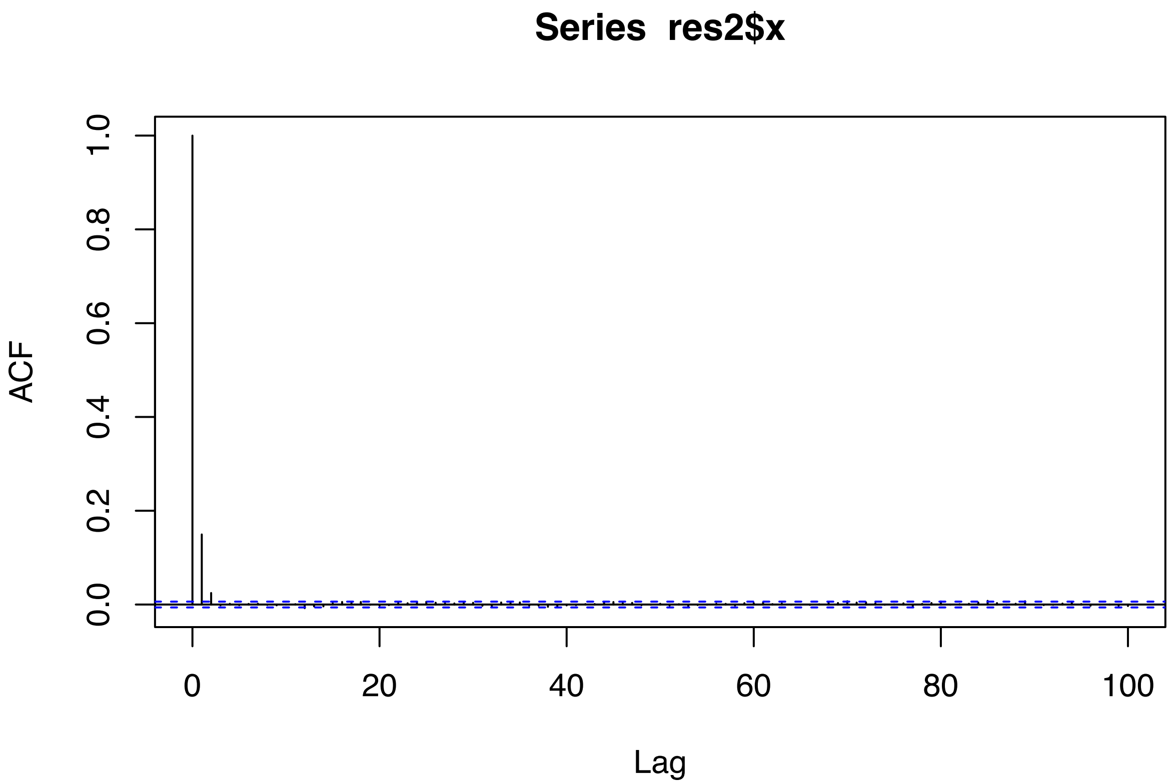


```
tau2<-1+2*sum(rho1[which(rho2>=0.1)])
print(10^5/tau2)
```

```
## [1] 22271.29
```

The latter seems more efficient. (This could be confirmed by generating several chains and averaging the results).

Exercise 2

Question a

The likelihood function is

$$L(\beta; y_1, \dots, y_n) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(y_i - \beta x_i)^2\right) = (2\pi)^{-n/2} \exp\left(-\frac{1}{2}\sum_{i=1}^{n}(y_i - \beta x_i)^2\right).$$

The density of the Gamma distribution with shape parameter a and rate b is $f(\beta) \propto \beta^{a-1}\exp(-b\beta)I(\beta > 0)$. Here $a = 2$ and $b = 1$, so $f(\beta) \propto \beta\exp(-\beta)I(\beta > 0)$. Consequently, the posterior density is

$$f(\beta \mid y_1, \dots, y_n) \propto \beta\exp(-\beta)\exp\left(-\frac{1}{2}\sum_{i=1}^{n}(y_i - \beta x_i)^2\right) I(\beta > 0).$$

Question b

We first write a function that computes the log-likelihood: