SCRIPT `MOD6S2A`: CONJUGATE PRIOR EXAMPLE

INSTRUCTOR: KLAUS MOELTNER

Examle 1: Vague priors for σ^2 and β

Consider a conjugate normal regression model with a single slope coefficient β. We first draw precision h from a gamma distribution, then $\boldsymbol{\beta}$ from a normal, conditional on h.
In our lecture notes the gamma density is parameterized in terms of its mean and degrees of freedom as

$$p(h) = \left(\frac{2\tau_0}{\nu_0}\right)^{-\nu_0/2} \Gamma\left(\frac{\nu_0}{2}\right)^{-1} h^{\frac{\nu_0-2}{2}} exp\left(\frac{-h\nu_0}{2\tau_0}\right)$$

where τ_0 is the mean and ν_0 is the D.o.F. parameter.

R, on the other hand, parameterizes the gamma distribution with shape a and scale s as

$$p(h) = \frac{1}{s^a \Gamma(a)} h^{a-1} exp(-h/s)$$

Thus, we have $a = \nu_0/2$ and $s = (2\tau_0/\nu_0)$.

Examle 1: Vague priors for σ^2 and β

You can obtain a relatively vague prior for σ^2 and β by setting both ν_0 and τ_0 relatively low, and choose a large variance for $\boldsymbol{\beta}$:

```
R> v0<-5 #prior for D.o.f. parameter for gamma density
R> tau0<-0.1 #prior for mean of gamma distribution
R> mu0<-0 #prior mean for beta
R> V0<-100 #prior variance for beta
R> #
R> R<-100000 #number of draws
R> #
R> hprior<-rgamma(R,v0/2,scale=(2*tau0/v0))
R> sig2prior<-1/hprior #get implied prior for error variance
R> betaprior<-rnorm(R,mu0,sqrt(sig2prior*V0))
R> #
R> # evaluate kernel densities (for plot)
R> hdens<-density(hprior,kernel="epanechnikov",n=1000)
R> sig2dens<-density(sig2prior,kernel="epanechnikov",n=10000)
R> betadens<-density(betaprior,kernel="epanechnikov",n=1000)
```

```
R> par(mfrow=c(3,1),mar=c(2,4,2,0),oma=c(0,0,0,0) )
R> plot(hdens,type="l",main = "prior for h",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1)
R> plot(sig2dens,type="l",main = "prior for sig2",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(0,100))
R> plot(betadens,type="l",main = "prior for beta",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(-100,100),)
R> #note: it's a good idea to let "plot" find the limits for x and y in this case,
R> # where we have many extreme draws. Else "plot" may produce misleading lines.
```

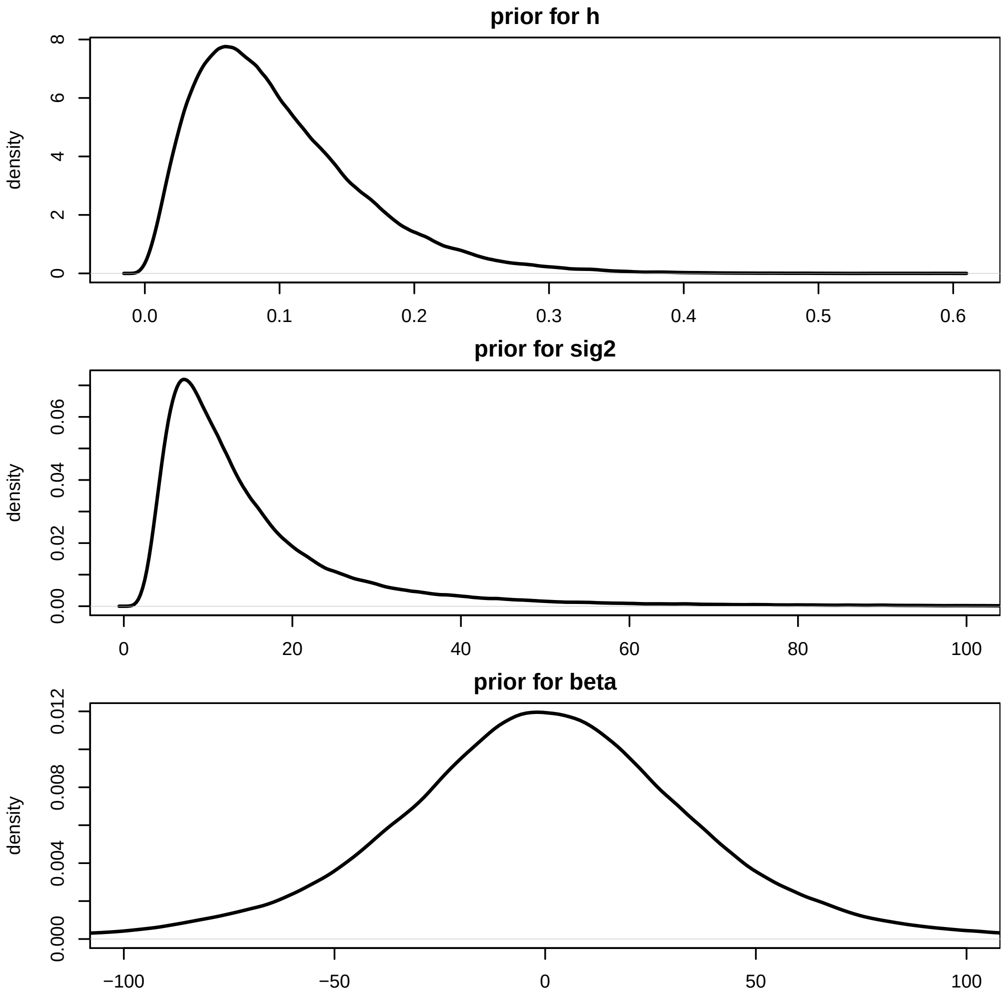

Figure 1. Prior plots, v.1

Example 2: Informed prior for σ^2, vague prior for β

Boosting the mean parameter for the gamma density produces a fairly informed prior for σ^2, and a relatively vague prior for β.

```
R> v0<-5 #prior for D.o.f. parameter for gamma density
R> tau0<-1 #prior for mean of gamma distribution
```

```
R> mu0<-0 #prior mean for beta
R> V0<-100 #prior variance for beta
R> #
R> R<-100000 #number of draws
R> #
R> hprior<-rgamma(R,v0/2,scale=(2*tau0/v0))
R> sig2prior<-1/hprior #get implied prior for error variance
R> betaprior<-rnorm(R,mu0,sqrt(sig2prior*V0))
R> #
R> # evaluate kernel densities (for plot)
R> hdens<-density(hprior,kernel="epanechnikov",n=1000)
R> sig2dens<-density(sig2prior,kernel="epanechnikov",n=10000)
R> betadens<-density(betaprior,kernel="epanechnikov",n=1000)
```

Example 3: Informed priors for σ^2 and β

If we now lower the variance parameter for β and - possibly - change its mean, we obtain an "informed" prior for β as well.

```
R> v0<-5 #prior for D.o.f. parameter for gamma density
R> tau0<-1 #prior for mean of gamma distribution
R> mu0<-3 #prior mean for beta
R> V0<-0.01 #prior variance for beta
R> #
R> R<-100000 #number of draws
R> #
R> hprior<-rgamma(R,v0/2,scale=(2*tau0/v0))
R> sig2prior<-1/hprior #get implied prior for error variance
R> betaprior<-rnorm(R,mu0,sqrt(sig2prior*V0))
R> #
R> # evaluate kernel densities (for plot)
R> hdens<-density(hprior,kernel="epanechnikov",n=1000)
R> sig2dens<-density(sig2prior,kernel="epanechnikov",n=10000)
R> betadens<-density(betaprior,kernel="epanechnikov",n=1000)

R> proc.time()-tic
   user  system elapsed
   0.87    0.11    1.11
```

```
R> par(mfrow=c(3,1),mar=c(2,4,2,0),oma=c(0,0,0,0) )
R> plot(hdens,type="l",main = "prior for h",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1)
R> plot(sig2dens,type="l",main = "prior for sig2",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(0,10))
R> plot(betadens,type="l",main = "prior for beta",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(-100,100))
R> #note: it's a good idea to let "plot" find the limits for x and y in this case,
R> # where we have many extreme draws. Else "plot" may produce misleading lines.
```

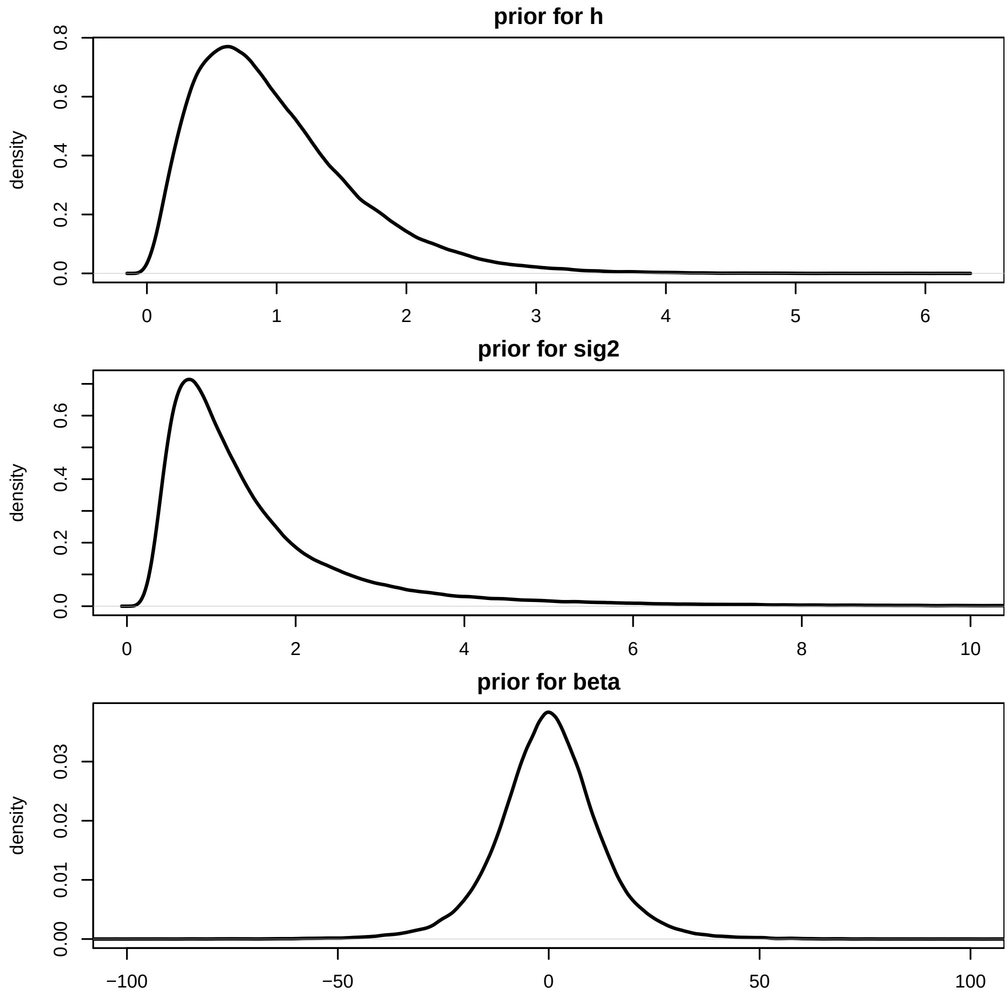

Figure 2. Prior plots, v.2

```
R> par(mfrow=c(3,1),mar=c(2,4,2,0),oma=c(0,0,0,0) )
R> plot(hdens,type="l",main = "prior for h",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1)
R> plot(sig2dens,type="l",main = "prior for sig2",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(0,10))
R> plot(betadens,type="l",main = "prior for beta",xlab = "",ylab = "density",
 lwd=2,col=1,lty=1,xlim=c(2,4))
R> #note: it's a good idea to let "plot" find the limits for x and y in this case,
R> # where we have many extreme draws. Else "plot" may produce misleading lines.
```

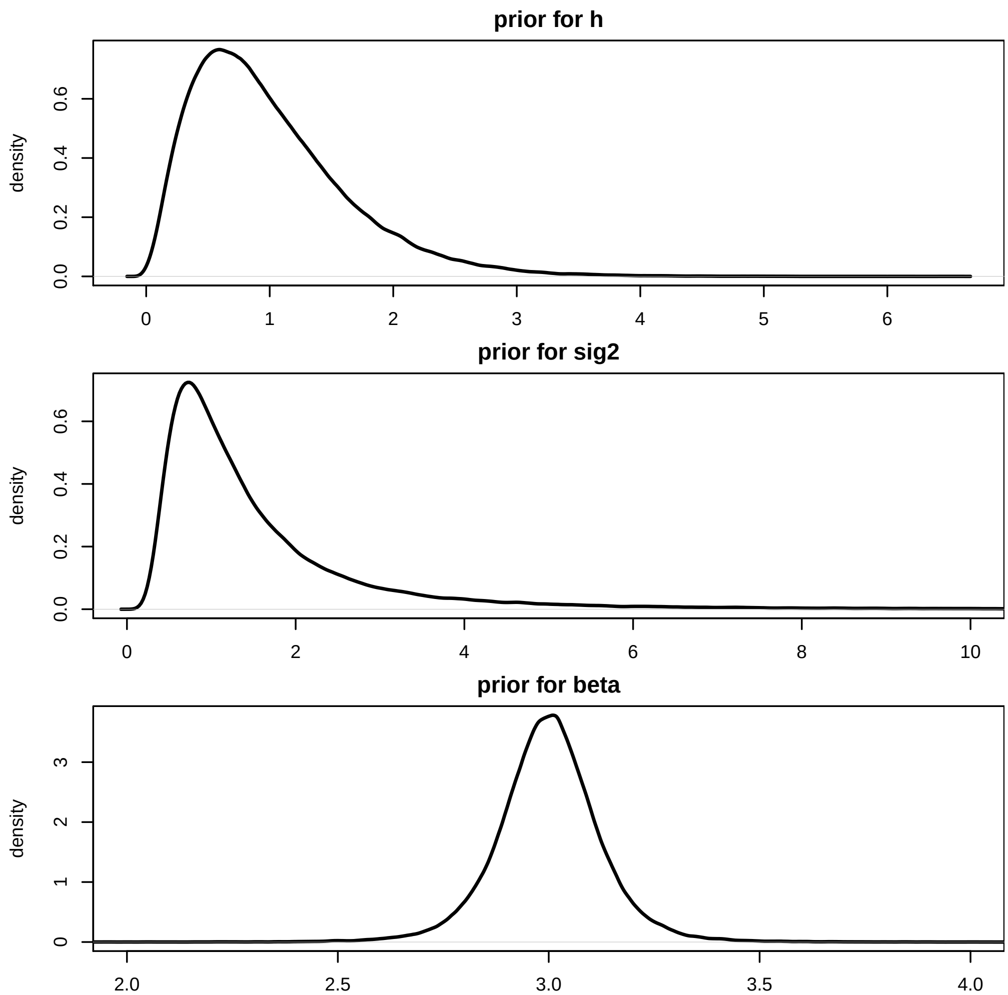

FIGURE 3. Prior plots, v.3