

- Using the script on the class site, install the `rstan` package and run the FTO example from class. Make sure you are using R version 4.3 or higher. Windows users: you will need to install the Rtools toolchain, which may take some time.

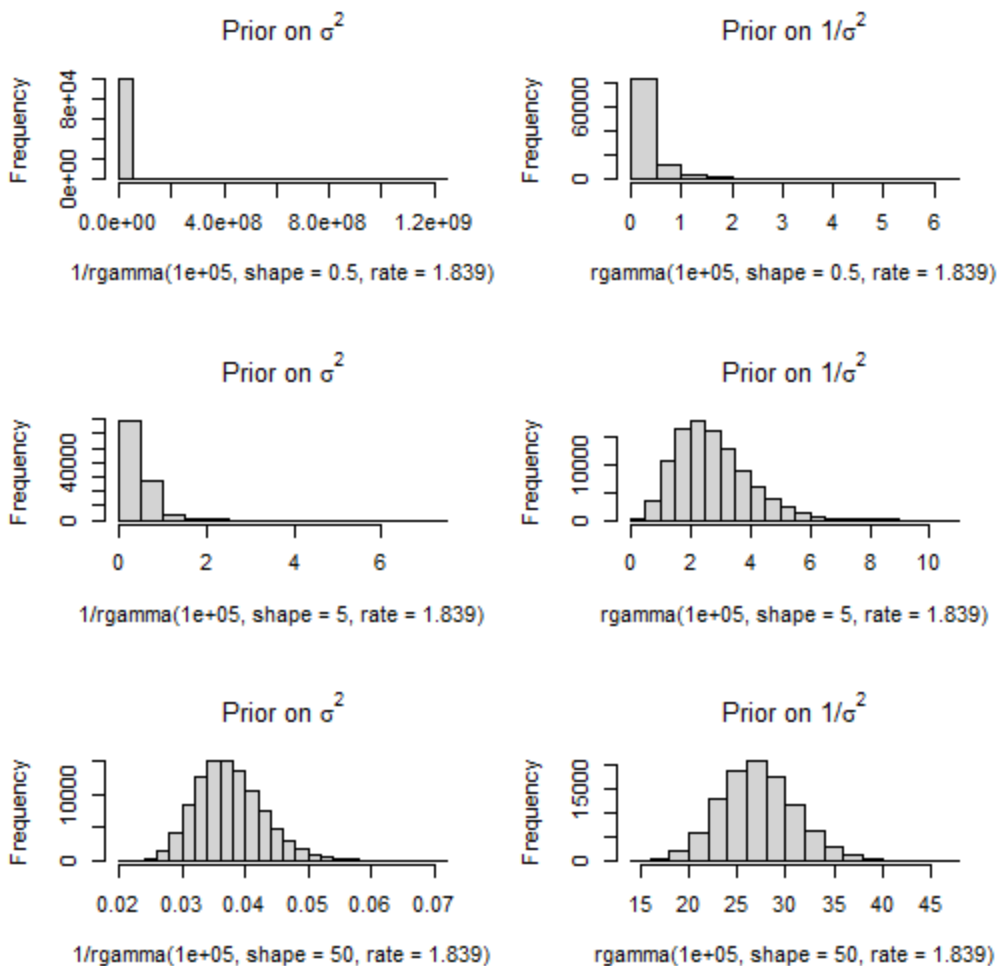
Do check with us if this doesn't work on your machine!

- Using the model and code in Q1, explore the impact of the prior on σ^2 in the FTO example. What happens to the posterior intervals for the β_j parameters if the prior for σ^2 gives more support to smaller values? Or to larger values?

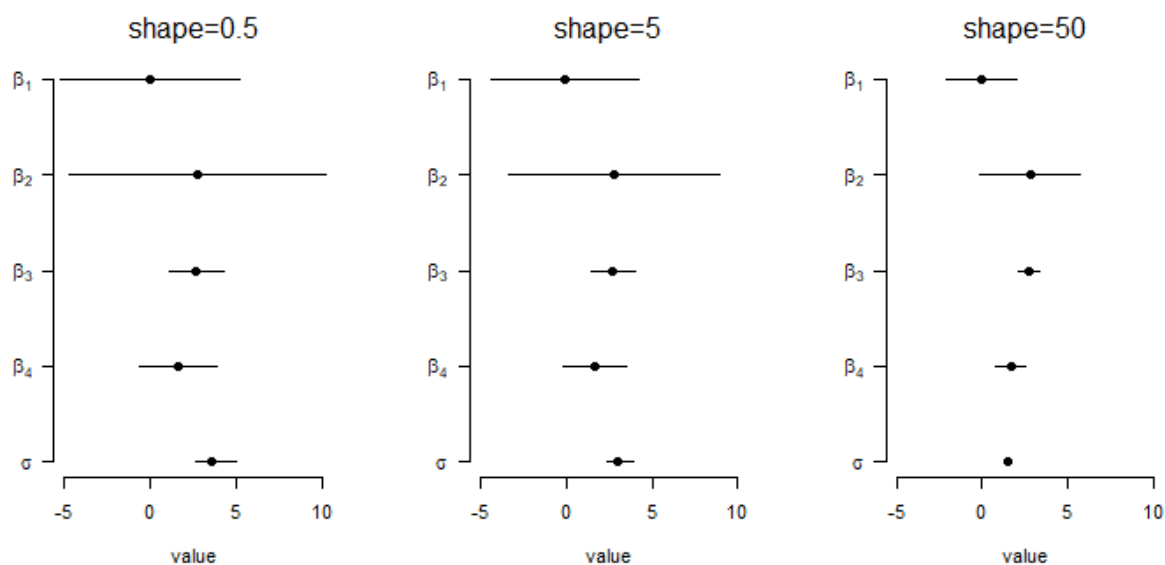
Hint: use the Monte Carlo method from Session 2 to check what prior on σ^2 is given by different values of ν_0, σ_0 .

If you have trouble with the first part – getting `rstan` to run – please ask for help on Slack, as this package will be used again in some later sessions.

For the second part, we illustrate the posteriors changing the “shape” parameter of the Gamma prior on $1/\sigma^2$. We use shape=0.5, 5 and 50, and the impact on the prior for σ^2 and $1/\sigma^2$ – obtained using simple Monte Carlo methods – are shown below:



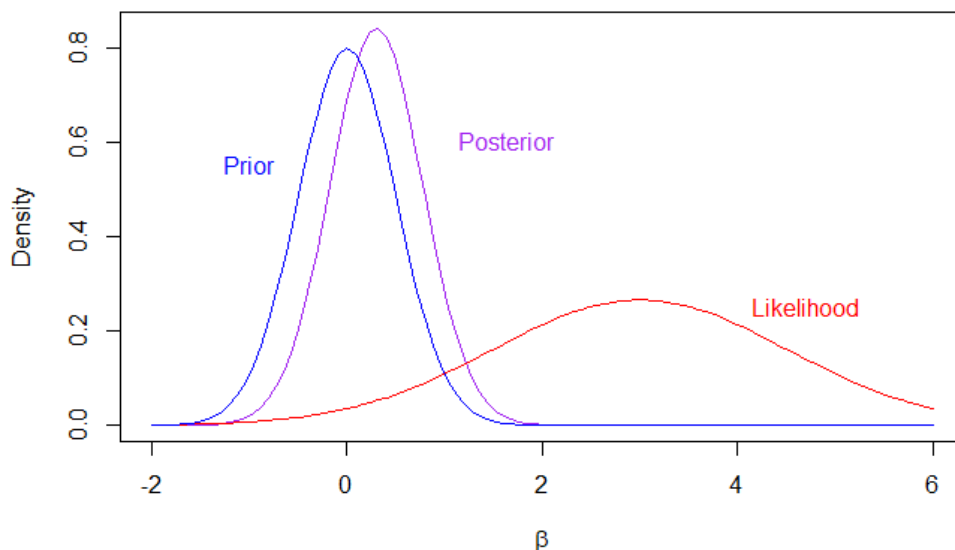
The impact on the posteriors for the regression parameters can be seen from the following plots of their 95% credible intervals – the posterior medians stay at similar values but the intervals shrink as our prior is more confident that each data point is highly informative.



3. Suppose with a small and noisy study of a parameter β , you obtain an estimate $\hat{\beta}=3$ with standard error 1.5. (This gives a Z-score of $3/1.5=2$, and hence a p-value just below 0.05.) Suppose your scientific colleagues have a skeptical $N(0,0.5^2)$ prior, with very little support for values of as large as the observed estimate.

Assuming (reasonably) that the likelihood is well-approximated by $N(3,1.5^2)$, what is the posterior? Do you think your p-value below 0.05 should be the start of a Nobel-worthy revolution of all that your colleagues foolishly believed?

We have Prior $N(0,0.5^2)$, Likelihood $N(3,1.5^2)$. Hence the prior precision is $1/0.5^2 = 4$, and the precision from the likelihood is $1/1.5^2=4/9$. The posterior precision is therefore $4+4/9=40/9$, and the posterior variance is $9/40=0.225$. To calculate the posterior mean, we take a weighted average of the prior mean (zero) and the mean from the likelihood. The weight on this value is $4/9/(40/9) = 1/10$, so the posterior mean is $3/10$. In a picture:



Code

```
# Question 1 - most of the code cut-and-pasted from class example:

# install.packages("rstan")
setwd("C:/Users/kenrice/Desktop/SISGBayes")
load("yX_FTO.Rdata")
y <- yX$y
X <- yX$X
n <- nrow(X)
p <- ncol(X)

cat(file="FTOexample.stan", "
data {
  int n; //the number of observations
  int p; //the number of columns in the model matrix
  real y[n]; //the response
  matrix[n,p] X; //the model matrix
  real g; // Zellner scale factor
  vector[p] mu; // Zellner prior mean (all zeros)
  matrix[p,p] XtXinv; // information matrix
}
parameters {
  vector[p] beta; //the regression parameters
  real invsigma2; //the precision, a.k.a. inverse-variance
}
transformed parameters {
  vector[n] linpred;
  cov_matrix[p] Sigma;
  real sigma;
  linpred = X*beta;
  sigma = 1/sqrt(invsigma2);
  for (j in 1:p){
    for (k in 1:p){
      Sigma[j,k] = g*sigma^2*XtXinv[j,k];
    } }
}
model {
  beta ~ multi_normal(mu, Sigma);
  y ~ normal(linpred, sigma);
  invsigma2 ~ gamma(5, 1.839); // change the first argument here, can be 0.5/5/50
}
")
# do the MCMC, store the results
library("rstan")
stan2 <- stan(file = "FTOexample.stan",
data = list(n=n,p=p, y=y, X=X, g=n, mu=rep(0,p), XtXinv=solve(crossprod(X)) ),
iter = 100000, chains = 1, pars=c("beta","sigma"))

print(stan2)
plot(stan2)
traceplot(stan2)
```

```

stan3 <- stan(file = "FTOexample.stan",
data = list(n=n,p=p, y=y, X=X, g=n, mu=rep(0,p), XtXinv=solve(crossprod(X)) ),
iter = 100000, chains = 1, pars=c("beta","sigma"))

stan4 <- stan(file = "FTOexample.stan",
data = list(n=n,p=p, y=y, X=X, g=n, mu=rep(0,p), XtXinv=solve(crossprod(X)) ),
iter = 100000, chains = 1, pars=c("beta","sigma"))

par(mfrow=c(3,2))
hist(1/rgamma(1E5, shape=0.5, rate=1.839), main=expression("**sigma^2"))
hist(rgamma(1E5, shape=0.5, rate=1.839), main=expression("Prior on 1/**sigma^2"))

hist(1/rgamma(1E5, shape=5, rate=1.839), main=expression("**sigma^2"))
hist(rgamma(1E5, shape=5, rate=1.839), main=expression("Prior on 1/**sigma^2"))

hist(1/rgamma(1E5, shape=50, rate=1.839), main=expression("**sigma^2"))
hist(rgamma(1E5, shape=50, rate=1.839), main=expression("Prior on 1/**sigma^2"))

#tables of what we'll plot
signif(summary(stan2)$summary[1:5,c(6,4,8)],2)
signif(summary(stan3)$summary[1:5,c(6,4,8)],2)
signif(summary(stan4)$summary[1:5,c(6,4,8)],2)

# plotting function
do.one <- function(mmm) {
  plot(0,0,xlim=c(-5,10), ylim=c(1, 5), axes=FALSE, xlab="value", ylab="")
  segments(mmm[,2], 5:1, mmm[,3], 5:1)
  points(y=5:1, x=mmm[,1], pch=19)
  axis(side=1)
  axis(side=2, las=1, at=5:1, c(
    expression(beta[1]),
    expression(beta[2]),
    expression(beta[3]),
    expression(beta[4]),
    expression(sigma))
  )
}

#make the plots
par(mfrow=c(1,3))
do.one(summary(stan2)$summary[1:5,c(6,4,8)])
mtext(side=3, line=1, "shape=0.5")
do.one(summary(stan3)$summary[1:5,c(6,4,8)])
mtext(side=3, line=1, "shape=5")
do.one(summary(stan4)$summary[1:5,c(6,4,8)])
mtext(side=3, line=1, "shape=50")

# Question 2

curve(dnorm(x, 3/10, sqrt(9/40)), -2,6, col="purple", ylab="Density",
xlab=expression(beta))
curve(dnorm(x, 3, 1.5), -2,6, col="red", add=TRUE)
curve(dnorm(x, 0, 0.5), -2,6, col="blue", add=TRUE)

```

```
text(x=c(-1, 1, 4), y=c(0.6, 0.6, 0.25), c("Prior","Posterior","Likelihood"),  
col=c("blue","purple","red"), pos=c(1,4,4))
```