

Class 3: An introduction to Bayesian Statistics

Andrew Parnell
andrew.parnell@mu.ie



PRESS RECORD

Learning outcomes

- ▶ Understand the terms prior, likelihood and posterior
- ▶ Know what a posterior probability distribution is, and why we take samples from it
- ▶ Know how to formulate a linear regression model in a Bayesian format
- ▶ Be able to suggest appropriate prior distributions for different situations

Who was Bayes?

An essay towards solving a problem on the doctrine of chances (1763)

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$



What is Bayesian statistics?

- ▶ Bayesian statistics is based on an interpretation of Bayes' theorem
- ▶ All quantities are divided up into *data* (i.e. things which have been observed) and *parameters* (i.e. things which haven't been observed)
- ▶ We use Bayes' interpretation of the theorem to get the *posterior probability distribution*, the probability of the unobserved given the observed
- ▶ Used now in almost all areas of statistical application (finance, medicine, environmetrics, gambling, etc, etc)

Why Bayes?

The Bayesian approach has numerous advantages:

- ▶ It's easier to build complex models and to analyse the parameters you want directly
- ▶ We automatically obtain the best parameter estimates and their uncertainty from the posterior samples
- ▶ It allows us to get away from (terrible) null hypothesis testing and p -values

Bayes theorem in english

Bayes' theorem can be written in words as:

posterior is proportional to likelihood times prior

... or ...

posterior \propto likelihood \times prior

Each of the three terms *posterior*, *likelihood*, and *prior* are *probability distributions* (pdfs).

In a Bayesian model, every item of interest is either data (which we will write as x) or parameters (which we will write as θ). Often the parameters are divided up into those of interest, and other *nuisance parameters*

Bayes theorem in maths

Bayes' equation is usually written mathematically as:

$$p(\theta|x) \propto p(x|\theta) \times p(\theta)$$


or, more fully:

$$p(\theta|x) = \frac{p(x|\theta) \times p(\theta)}{p(x)}$$

- ▶ The *posterior* is the probability of the parameters given the data
- ▶ The *likelihood* is the probability of observing the data given the parameters (unknowns)
- ▶ The *prior* represents external knowledge about the parameters

A very simple linear regression example

Suppose you had some data that looked like this:



What you are used to doing

```
model = lm(log(earn) ~ height_cm, data = dat)
summary(model)

##
## Call:
## lm(formula = log(earn) ~ height_cm, data = dat)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -4.4351 -0.3705  0.1615  0.5761  2.3302
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  5.895057   0.489107  12.053 < 2e-16 ***
## height_cm    0.022555   0.002866   7.869 8.84e-15 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.9066 on 1057 degrees of freedom
## Multiple R-squared:  0.05533,    Adjusted R-squared:  0.05444
## F-statistic: 61.91 on 1 and 1057 DF,  p-value: 8.836e-15
```

What you will now get instead

```
print(model_run)
```

```
## Inference for Bugs model at "4", fit using jags,  
## 3 chains, each with 2000 iterations (first 1000 discarded)  
## n.sims = 3000 iterations saved  
##           mu.vect sd.vect   2.5%   25%   50%   75%   97.5% Rhat  
## intercept    5.885  0.501  4.901  5.553  5.883  6.226  6.852 1.003  
## residual_sd  0.907  0.019  0.870  0.894  0.907  0.920  0.947 1.003  
## slope        0.023  0.003  0.017  0.021  0.023  0.025  0.028 1.004  
## deviance    2798.545  2.472 2795.738 2796.701 2797.903 2799.639 2805.077 1.002  
##           n.eff  
## intercept      700  
## residual_sd    680  
## slope          620  
## deviance      1700  
##  
## For each parameter, n.eff is a crude measure of effective sample size,  
## and Rhat is the potential scale reduction factor (at convergence, Rhat=1).  
##  
## DIC info (using the rule, pD = var(deviance)/2)  
## pD = 3.1 and DIC = 2801.6  
## DIC is an estimate of expected predictive error (lower deviance is better).
```

Using prior information

- ▶ The Bayesian model in the previous slide divided up everything into *parameters* (the intercept, slope and residual standard deviation), and data (the height and log(earnings) values)
- ▶ The software in the background created a posterior probability distribution of the parameters given the data
- ▶ The model I fitted used no *prior information*. However, if we had done a previous experiment that suggested the intercept should be around 9 with standard deviation 0.5 we can put this in the model

A model with prior information

```
print(model_run_2)
```

```
## Inference for Bugs model at "5", fit using jags,  
## 3 chains, each with 2000 iterations (first 1000 discarded)  
## n.sims = 3000 iterations saved  
##           mu.vect sd.vect   2.5%   25%   50%   75%   97.5% Rhat  
## intercept    7.422  0.352  6.718  7.184  7.421  7.657  8.107 1.003  
## residual_sd  0.911  0.020  0.873  0.898  0.911  0.924  0.951 1.001  
## slope        0.014  0.002  0.010  0.012  0.014  0.015  0.018 1.002  
## deviance    2807.755  4.964 2799.551 2804.190 2807.232 2810.798 2819.251 1.001  
##           n.eff  
## intercept    980  
## residual_sd 3000  
## slope        1200  
## deviance    2900  
##  
## For each parameter, n.eff is a crude measure of effective sample size,  
## and Rhat is the potential scale reduction factor (at convergence, Rhat=1).  
##  
## DIC info (using the rule,  $pD = \text{var}(\text{deviance})/2$ )  
##  $pD = 12.3$  and  $DIC = 2820.1$   
## DIC is an estimate of expected predictive error (lower deviance is better).
```

An early example of a Bayesian model

- ▶ To create the Bayesian version of this model I used the following JAGS code:

```
model_code = '  
model {  
  # Likelihood  
  for(i in 1:N) {  
    y[i] ~ dnorm(intercept + slope*x[i], residual_sd^-2)  
  }  
  # Priors  
  intercept ~ dnorm(9,0.5^-2)  
  slope ~ dnorm(0,100^-2)  
  residual_sd ~ dunif(0,100)  
}  
'
```

Understanding the different parts of a Bayesian model

- ▶ The likelihood is the probability of observing the data given the parameters. It represents the *data generating process*
- ▶ The prior is the probability distribution of the parameters independent from the current data you have been generated. It often requires care (and philosophy) to choose. More on this later
- ▶ The posterior is the probability distribution of the parameters given the data. It is always the target of our inference
- ▶ The two software packages we will explore (Stan and JAGS), create the posterior distribution for us

How do I specify the prior distribution?

There are several choices when it comes to specifying prior distributions:

- ▶ *Informative*, when there is information from a previous study, or other good external source, e.g. $\text{intercept} \sim N(-30, 5^2)$
- ▶ *Vague*, when there is only weak information, perhaps as to the likely range of the parameter e.g. $\text{intercept} \sim N(0, 100^2)$
- ▶ *Flat*, when there is no information at all about a parameter (very rare). In JAGS, write $\text{intercept} \sim \text{dflat}()$

In fact, choosing the prior and choosing the likelihood are very similar problems

Practical differences between frequentist statistics and Bayes

- ▶ In frequentist statistics you tend to get a single best estimate of a parameter and a standard error, often assumed normally distributed, and a p-value
- ▶ In Bayesian statistics you get a large set of samples of the parameter values which match the data best. You get to choose what you do with these
- ▶ In frequentist statistics if the p-value is less than 0.05 you win. If not you cry and try a different model
- ▶ In Bayesian statistics you try to quantify the size of an effect from the posterior distribution, or find a particular posterior probability, e.g. $P(\text{slope} > 0 \text{ given the data})$.

More generally choosing a likelihood and a prior

- ▶ The key to choosing a likelihood is to pick a probability distribution which matches your data, e.g. if it's a continuous measurement and is unbounded then a normal distribution. If it's count data bounded above and below then a Binomial might be appropriate
- ▶ The key to choosing a prior distribution is to choose values which you believe represent the reasonable range that the parameter can take, or come from a related study in the literature
- ▶ Again, use an *informative prior* if possible

Note: the shape of the distribution of the response variable is usually completely unimportant when choosing the likelihood!

Stan vs JAGS

We will be using two different software tools to calculate posterior distributions. These represent the state of the art for user-friendly, research quality Bayesian statistics.

- ▶ Stan positives: very flexible, uses sensible distribution names, everything is declared, lots of documentation support, written by people at the top of the field
- ▶ Stan negatives: cannot have discrete parameters, some odd declaration choices, slower to run code, code tends to be longer
- ▶ JAGS positives: very quick for simple models, no declarations required, a bit older than Stan so more queries answered online
- ▶ JAGS negatives: harder to get complex models running, not as fancy an algorithm as Stan, crazy way of specifying normal distributions

Posterior computation in Stan

A Stan model looks like this:

```
stan_code = '  
data {  
  int N;  
  vector[N] x;  
  vector[N] y;  
}  
parameters {  
  real intercept;  
  real slope;  
  real<lower=0> residual_sd;  
}  
model {  
  // Likelihood  
  y ~ normal(intercept + slope * x, residual_sd);  
  // Priors  
  intercept ~ normal(0, 100);  
  slope ~ normal(0, 100);  
  residual_sd ~ uniform(0, 100);  
}  
'
```

Running a model in Stan

To run the Stan model:

```
stan_run = stan(data = list(N = nrow(dat),  
                             y = log(dat$earn),  
                             x = dat$height_cm),  
                model_code = stan_code)  
print(stan_run)  
plot(stan_run)
```

Posterior computation in JAGS

The same model in JAGS:




```
jags_code = '  
model{  
  # Likelihood  
  for(i in 1:N) {  
    y[i] ~ dnorm(intercept + slope * x[i],  
                residual_sd^-2)  
  }  
  # Priors  
  intercept ~ dnorm(0, 100^-2)  
  slope ~ dnorm(0, 100^-2)  
  residual_sd ~ dunif(0, 100)  
}  
'
```

Note the $\hat{-2}$ everywhere - JAGS uses precision rather than standard deviation in the normal distribution

Running the JAGS mdoel

```
jags_run = jags(data = list(N = nrow(dat),  
                           y = log(dat$earn),  
                           x = dat$height_cm),  
               parameters.to.save = c('intercept',  
                                       'slope',  
                                       'residual_sd'),  
               model.file = textConnection(jags_code))  
  
print(jags_run)  
plot(jags_run)
```

Bugs model at "6", fit using jags, 3 chains, each with 2000 iterations (first 1000 discarded)



Calculating the posterior vs sampling from it

- ▶ There are two ways to get at a posterior:
 1. Calculate it directly using hard maths
 2. Use a simulation method
- ▶ Number 1 is impractical once you move beyond a few parameters, so number 2 is used by almost everybody
- ▶ This means that we create *samples* from the posterior distribution. Here are three samples from the earnings example:

```
##      deviance intercept residual_sd      slope
## [1,] 2800.902   5.453393   0.8680240 0.02525612
## [2,] 2799.329   5.588173   0.9064387 0.02465221
## [3,] 2797.616   5.583262   0.9320569 0.02438453
```

- ▶ We often create thousands of posterior samples to represent the posterior distribution

Things you can do with posterior samples

- ▶ Create histograms or density plots:
- ▶ Individual summaries such as means, standard deviations, and quantiles (e.g. 95% confidence intervals)
- ▶ Joint summaries such as scatter plots or correlations
- ▶ Transformations such as logs/exponents, squares/square roots, etc

The posterior distribution will usually be stored in a matrix where each row is a sample, and each column is a different parameter. Having the posterior distribution enables you to get at exactly the quantities you are interested in

Summary so far: for and against Bayes

For:

- ▶ A Bayesian model can be simply displayed as a likelihood and a prior. Everything is explicit
- ▶ JAGS/Stan finds the posterior distribution for us so we don't need to do any maths
- ▶ We can get exactly the quantity we are interested in, the probability distribution of our unknowns given our knowns

Against:

- ▶ It can be hard to create a prior distribution (and a likelihood)
- ▶ Not having p-values can make papers harder to publish (but this is changing)

Checking assumptions (e.g. residuals)

- ▶ Sometimes people, because Bayesian modelling seems much richer and incorporates more information, think that their model is perfect
- ▶ In reality we need to check our assumptions. This may include:
 1. Checking residuals in a linear regression model
 2. Checking whether the parameter values actually match the data
 3. Seeing whether a simpler or richer model might fit the data better

Some of this we will cover in later classes, but e.g. traditional residual analysis for linear regression is still important here

The danger of vague priors

- ▶ Suppose you use the prior distribution $\text{intercept} \sim \text{dnorm}(0, 100^{-2})$ in JAGS because you had very little information about the intercept. Is this a reasonable thing to do? Do you honestly believe that the intercept might be as small as -200 or as big as 200?
- ▶ If you fit a model and the parameters do not match your views about the data, there must be some information you have not encoded in the prior, go back and change it!
- ▶ In more complex models, you need the prior to constrain some of the parameters so that the model will fit. These are known as *regularisation* priors
- ▶ Use an *informative prior* when you can

General tips

- ▶ If you have lots of disparate data, try to build one model for all it. You'll be able to *borrow strength* from the data (e.g. in a hierarchical model) and reduce the uncertainty in your parameter estimates
- ▶ Try your hardest to use informative priors, and always justify the values you use (especially when trying to publish). In this course we're presenting generic examples so have almost always used vague priors
- ▶ Check your model. Many of the usual requirements from traditional statistics (e.g. residual checks) are still relevant in the Bayesian world. There are some extra Bayesian tricks we can also do; discussed in later classes

Summary

- ▶ Bayesian statistical models involve a *likelihood* and a *prior*. These both need to be carefully chosen. From these we create a posterior distribution
- ▶ The likelihood represents the information about the data generating process; the prior information about the unknown parameters
- ▶ We usually create and analyse samples from the posterior probability distribution of the unknowns (the parameters) given the knowns (the data)
- ▶ From the posterior distribution we can create means, medians, standard deviations, credible intervals, etc, from samples we take using e.g. JAGS or Stan