Re: Anyone found an Elementary Bayesian Book yet?

Source: http://sci.tech−archive.net/Archive/sci.stat.math/2006−11/msg00072.html

- *From*: "Reef Fish" <large_nassua_grouper@xxxxxxxxx>
- *Date*: 2 Nov 2006 05:37:45 −0800

David Jones wrote:

> No. Beta(r+1+alpha, n−r+1+beta) as a posterior density would contain
> p to the power of (r+alpha), not (r+alpha−1) as required.
> Beta(alpha+r, beta+n−r) is the answer given in cookbooks.

It's my turn to say No. And this is going to be the last round for me.

The cookbook says n'' = n + n' and r'' = r + r'

where ALL the n's and r's are just powers, in p^r (1−p)^(n−r), the likelihood form, without making an index shift to the alpha's and beta's. I did that to make the point about Beta being the conjuate prior, and I regret doing so now, because it introduced the confusion of mixing one parametrization with a different one.

So, I'll go back the original book from which Russell Martin cited:

RM> Chapter 3 treats Bernoulli
RM> processes, and beta distributions as conjugate priors.
RM> Interestingly (if I'm reading it correctly) he suggests using
RM> Chapter 3 treats Bernoulli
RM> processes, and beta distributions as conjugate priors.
RM> Interestingly (if I'm reading it correctly) he suggests using
RM> r=n=0 as "vague" prior parameters. He acknowledges
RM> that this gives a prior beta density that is undefined, but
RM> writes, "Nevertheless, if we ignore this deficiency and
RM> apply Eq, (3.5) using r'=n'=0 as prior parameters, then
RM> the posterior parameters become r''=r and n''=n.

Russell never showed Eq. (3.5), but from the use of r'=n'=0 for the prior parameters, I assume the author was using the "standard cookbook" form as I explained:

RF> He even had the primes according to the usual cookbook
RF> conventions. The r and n denote the SAMPLE r and n,
RF> those in the likelihood function. r' and n' denote the
RF> parameters in the prior distribution Beta(r',n') rather than
RF> alpha and beta. That's because then you have the
RF> "no brainer" of using the Beta as the conjugate prior,
RF> because the posterior is given by (double primes)

RF> r" = r + r' and n" = n + n'.

Russell can check me on that. I was inferring from Russell's cited paragraph that the line above is the Eq. (3.5) with my explanation of Beta(r',n') above.

Then Russell's B(0.0) is the r'=n'=0 of 1/(p(1−p)) improper prior. The uniform prior is B(1,1), a proper prior, in that notion.

The r and n remain the sample data and likelihood r number of successs out of n trial.

So, the improper prior, denoted by B(0.0) = 1/(p(1−p)), is the one that yields the original likelihood function as the identical posterior distribution.

Now we get back to David Jones.

> and that's why B(0,0) will give the likelihood function for the posterior.

> No Beta(1,1) (uniform) gives (a scaled version of) the likelihood function for the posterior. As noted by others, this doesn't necessarily mean that the uniform can be counted as non−informative ... one of my books says that even Bayes had doubts about making such a claim.

What do you mean by "(a scaled version of) the likelihood"?

B(0,0) gives back the original likelihood function for the posterior.

B(1,1) gives a function that bumps the powers of p and (1−p) in the original likelihood by 1, to yield a posterior distribution DIFFERENT from the original likelihood.

> Another check is to see what happens when there is no data from the
> experiment, so that the posterior is the same as the prior. "No

experiment" would be equivalent to n=r=0

The fly in the ointment is that with no data, you DON'T have a likelihood function and you don't have a posterior density.

But now, reverting back to Russell Martin's book and the convention used THERE, your check works. Because then n" = n' and r" = r' if you take the data n=r=0.

David Jones

I think David Jones's book and Russell Martin's book should now agree with the same recipe in the cookbooks from which I've taught:

RF> r" = r + r' and n" = n + n'.

If not, then I'll just stand by the present version as the version that is consistent in ALL the books I've used, which give

r" = r + 1 and n" = n + 1

for the uniform prior, U(0,1), denoted as B(1,1) in these books.

−− Reef Fish Bob.

.