

Notes on Exam

The final exam is worth 50% of the grade for this course. This set of notes is intended to give you an outline of the structure of the exam.

The exam will be made up of 8 questions of which you must do four. Each question is allocated equal weight. Below I provide more detail on the types of questions you might expect to find in each of the parts.

When you arrive at the exam you will be provided with the exam paper at the end of which is a section labelled "Some Key Results". These will provide you with the definitions of some key distributions (so you do not have to memorize these). To be precise they will provide you with definitions of the Normal, Gamma and Normal-Gamma distributions and the form for the prior and posterior for the Normal linear regression model with natural conjugate prior.

The exam will contain various types of questions.

Type 1: Understanding Bayesian Theory

These questions will be essay-orientated and focussed on the lecture material. The questions will be fairly broad in scope and allow you to reveal your level of understanding of general concepts in Bayesian analysis.

Type 2: Doing Bayesian Econometrics in Practice

The questions in this section relate to doing Bayesian econometrics in practice. To do well in these questions, you should know the lecture material well and remember or be able to figure out the general outlines of the relevant formula, but it is not necessary to reproduce the exact formulae. To give an example of what is meant by the preceding sentence, consider the Normal linear regression model with independent Normal-Gamma prior. I would expect you to know that the prior for β is Normal and the prior for h is Gamma. The posterior does not have a convenient form by the posterior conditionals for $\beta|h$ and $h|\beta$ are Normal and Gamma, respectively. But I would not expect you to know the precise formulae for the mean of the Normal conditional posterior, etc.. I would expect you to talk about posterior simulation using Gibbs sampling in this model. I would expect you to be able to about extensions to this model (e.g. allowing for general error covariance matrix), but not memorize exact formulae.

To give you some flavour for the types of questions which might be asked, the general structure of an answer to a question might be:

“Sketch out the form for a posterior for a certain model covered in the lectures, argue in favour of a particular way of doing posterior inference (e.g. argue that Gibbs sampling is a good thing to do), then sketch out the form of the posterior conditionals for the Gibbs sampler”.

Or

“For a new model that you have not seen before, recognize that it is related to a model covered in lectures. Discuss how methods covered in the lectures can be extended to handle the new model”

Or

“Recognize a particular model as being one you have learned about in the lectures, talk about its properties and discuss problems/issues (e.g. when different priors are used) which might arise when doing empirical work.”

Type 3: Bayesian Derivations

There will be only one question of this type. This type of question will require you to do derivations and/or proofs. Some such were done in the lectures, but the questions in the theoretical problem sets are most relevant. I have put two theoretical problem sets of the course website. The question on the exam will be similar to those in Problem Set 1: Bayesian Theory and not as difficult as those in Problem Set 2. The question of this type will not require knowledge of matrix algebra.