

POLS 506: Advanced Topics in Political Methodology
Bayesian Statistics
562 Sewall Hall, Thursday 2:30p-5:20p

Dr. Justin Esarey
Associate Professor of Political Science
Rice University

Syllabus Version: 8/21/2017

Office: 108 Herzstein Hall
Office Hours: W R 9:00 – 10:00 AM
E-mail: justin@justinesarey.com
Phone: 678-383-9629

COURSE OBJECTIVES AND LEARNING OUTCOMES

This course introduces students to advanced statistical techniques and their application to problems in political science.

Students will be able to:

1. understand and apply basic concepts of Bayesian inference
2. conduct a Monte Carlo analysis to discover model characteristics and diagnose problems
3. understand and implement computational estimation techniques for sampling from an analytically unknown distribution (e.g., MCMC and HMC)
4. explain, implement, interpret, and diagnose problems with Bayesian statistical models (including hierarchical models and item response theory estimates)
5. conduct hypothesis tests in the Bayesian framework
6. explain the consequences of missing data in a sample, and implement corrective statistical techniques

GRADING POLICIES AND ASSIGNMENT DETAILS

Grade Components:

- Homework: 25%
- Exam 1: 25%
- Exam 2: 25%
- Final Exam: 25%

Grading Scale:

100%-97%: A+
96.9%-93%: A
92.9%-90%: A-
89.9%-87%: B+
86.9%-83%: B
82.9%-80%: B-
79.9%-77%: C+

76.9%-73%: C
72.9%-70%: C-
69.9%-67%: D+
66.9%-63%: D
62.9%-60%: D-
>59.5%: F

Exams: There will be three exams in this class, two midterms and a final. All exams are cumulative, but will focus on material learned since the last exam. You must complete each exam within the allotted time period, and must submit a typed LaTeX answer sheet. The exams are open book and open note, but you may not consult anyone for advice on the exam. The rough timing of the exams is indicated on the course outline, and specific times will be scheduled in consultation with the class.

Homework: Homework problem sets will be distributed during class. I encourage collaborative work on problem sets: the goal of a homework problem set is to help you learn the material and enable you to perform well on the (non-collaborative!) exams. With that said, simply copying another student's homework answers is not permitted and will be treated as academic dishonesty.

All homeworks must be typed in LaTeX.

Attendance: Attendance is mandatory in this class, and as graduate students I expect that attendance will not be a problem for you. Every class you fail to attend (without an acceptable excuse—see below) will result in a 2.5 percentage point deduction from your final grade. (I expect that this will never happen.)

Attendance penalties may be waived in the event of death in the immediate family (parent, spouse, sibling, or child) within 2 weeks before the due date, in the event of an unforeseeable medical emergency affecting yourself, your spouse, or your child, if you are participating in a pre-approved academic activity (e.g., a conference), or for other unforeseeable exigencies; all waivers are at the discretion of the instructor. Supporting documentation may be required to support an attendance penalty waiver.

COURSE POLICIES

Late Work: Assignments are due at the date and time I specify for the assignment. Late homeworks will be marked off at 5 percentage points for the first 24 hours late, and an additional 10 percentage points for every subsequent 24 hours late. For exams, the first hour late incurs a 5 percentage point penalty and each additional hour incurs a 10 percentage point penalty.

Late work penalties may be waived in the event of death in the immediate family (parent, spouse, sibling, or child) within 2 weeks before the due date, in the event of an unforeseeable medical emergency affecting yourself, your spouse, or your child, or for other unforeseeable exigencies; all waivers are at the discretion of the instructor. Supporting documentation may be required to support an attendance penalty waiver.

Honor Code/Academic Misconduct: All forms of academic misconduct will be handled according to the Rice University Honor Code. Details on the Honor Code are available at <http://honor.rice.edu/honor-system-handbook/>.

If you ever have any questions about what you should do to stay within the honor code on a particular assignment, PLEASE contact me with your question and I can assist you. I cannot guarantee a timely response unless you contact me at least 24 hours in advance of the time the assignment is due.

Students with Disabilities: If you have a disability and require accommodation in this class, please contact me as soon as possible (within the first two weeks of class) to discuss these accommodations. You will also need to contact the Disability Support Services Office (telephone extension: 5841) in the Allen Center.

Syllabus Change Policy: The policies of this syllabus may be changed by Prof. Esarey with advance notice.

COURSE MATERIALS

Required Texts:

- Jackman, Simon. 2009. *Bayesian Analysis for the Social Sciences*. Wiley.
- Gelman, Andrew, and Jennifer Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge.

Recommended Texts:

- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2014. *Bayesian Data Analysis, Third Edition*. Chapman and Hall/CRC.
- Gill, Jeff. 2008. *Bayesian Methods: A Social and Behavioral Sciences Approach, 2nd ed.* Chapman and Hall/CRC.

Other readings are available on the web or the OWL-Space website.

Software: This course will teach material primarily through R and WinBugs. R is free and available from <http://cran.r-project.org/>. WinBUGS and OpenBUGS are free and available from <http://www.mrc-bsu.cam.ac.uk/software/bugs/>. JAGS is an alternative Gibbs Sampler that works on Mac OS X; it is free and available from <http://mcmc-jags.sourceforge.net/>. We will also study Hamiltonian Monte Carlo and STAN; STAN is available at <https://github.com/stan-dev/rstan/wiki/RStan-Getting-Started>.

All homeworks, papers and presentation slides must be typed in LaTeX. If you choose to do so, you may use LyX (<http://www.lyx.org/>), a WYSIWYG editor, in combination with MiKTeX on Windows (<http://miktex.org/>), MacTeX on Macintosh (<http://www.tug.org/mactex/>) or TeXLive on Linux (<http://www.tug.org/texlive/>).

All students must have a valid Rice e-mail address and login (and access to the OWL-space website) to participate in this course.

COURSE OUTLINE AND ASSIGNED READINGS

0) Software and Preliminaries

Questions: What are the basic tools required for this course?

Skills and concepts: using R, LaTeX, and WinBUGS/OpenBUGS/JAGS/STAN

Readings: none.

1) Basic Concepts of Bayesian Inference

Questions: What is Bayesian inference, and how is it different from inference derived from other procedures (e.g., ML)? What is a prior, and how do we specify it? How do we test a hypothesis in a Bayesian sense?

Skills and concepts: Bayes' rule; conjugate prior and analytically closed posterior probability densities; Bayesian p-values and credible intervals.

Readings:

- Jackman, Chapter 1
- Sections 1-2 of Jackman, Simon. 2004. "Bayesian Analysis for Political Research." *Annual Review of Political Science* 7: 483-505. URL: <http://www.annualreviews.org/doi/abs/10.1146/annurev.polisci.7.012003.104706>.
- *Recommended:* Gill, Chapter 1

2) Simple Bayesian Models

Questions: How can we use Bayesian models for simple models (mean/variance estimates and regressions)?

Skills and concepts: Estimating a mean and variance with Bayesian methods; Bayesian regression.

Readings:

- Jackman, Chapter 2
- Sections 3 of Jackman, Simon. 2004. "Bayesian Analysis for Political Research." *Annual Review of Political Science* 7: 483-505.
- (recommended) Gill, Chapters 2-3

3) Basic Monte Carlo Procedures and Sampling Algorithms

Questions: How does one conduct a Monte Carlo simulation? How does the idea of Monte Carlo analysis apply to sampling from a target distribution? What is a Markov Chain?

Skills and concepts: Monte Carlo simulation methods; sampling algorithms; Markov chains.

Readings:

- Jackman, Chapters 3-4
- (recommended) Gill, Chapter 8

4) The Metropolis-Hastings algorithm and the Gibbs Sampler

Questions: How do we produce estimates of a posterior density, and how do these depend on a prior? How do we implement such a technique in R? How do we ensure that our model is behaving as expected?

Skills and concepts: MCMC; the Metropolis-Hastings algorithm; the Gibbs sampler; model checking.

Readings:

- Jackman, Chapter 5
- Section 4 onward of Jackman, Simon. 2004. "Bayesian Analysis for Political Research." *Annual Review of Political Science* 7: 483-505.
- (recommended) Gill, Chapter 9

***** CUTOFF FOR EXAM 1 *****

5) Practical MCMC for Estimating Models

Questions: What software packages exist to automate MCMC estimation in R? How can we apply these techniques to substantively relevant statistical models?

Skills and concepts: MCMCpack, WinBUGS, and R2WinBugs; estimating simple GLM-family models with Bayesian methods.

Readings:

- Jackman, Chapter 6
- The WinBUGS manual, particularly the Tutorial section. URL: <http://cran.r-project.org/web/packages/R2WinBUGS/vignettes/R2WinBUGS.pdf>
- Sturtz et al., “R2WinBUGS: A Package for Running WinBUGS from R.” URL: <http://cran.r-project.org/web/packages/R2WinBUGS/vignettes/R2WinBUGS.pdf>
- The MCMCpack manual. URL: <https://cran.r-project.org/web/packages/MCMCpack/MCMCpack.pdf>
- (recommended) Gill, Chapter 12

6) Bayesian Hierarchical Linear Models and GLMs

Questions: What is a hierarchical model? How do we estimate a hierarchical model? How do we report results from a hierarchical model? How do hierarchical models relate to standard OLS, random effects, and fixed effects models? In what situations would we expect a hierarchical model to improve inference, and in which would we *not* expect improvement?

Skills and Concepts: review of random effects/fixed effects; estimating and interpreting a hierarchical GLM model in R using the lme4 package; comparing RE/FE results to HLM/hierarchical GLM results

Readings:

- Gelman and Hill: Chapters 11-12 and 14-15
- (recommended) Gelman et al. *Bayesian Data Analysis, Third Edition:*
 - Chapter 15: Hierarchical Linear Models
 - Chapter 16: Generalized Linear Models
- (recommended) Jackman, Chapters 7-8

7) Fitting Hierarchical Models with BUGS

Questions: How do you fit and interpret a hierarchical model using WinBUGS, OpenBUGS, and/or JAGS?

Skills and Concepts: WinBUGS/OpenBUGS, R2WinBUGS, and JAGS use for building and interpreting hierarchical models.

Readings:

- Gelman and Hill: Chapters 16-17

8) Hamiltonian Monte Carlo and STAN

Questions: What is Hamiltonian Monte Carlo, and how can it be used to sample from posterior distributions using STAN?

Skills and Concepts: Hamiltonian Monte Carlo; fitting and interpreting models using STAN.

Readings:

- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2014. "Computation in R and Stan." Appendix C in Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin, *Bayesian Data Analysis, Third Edition*: 589-606. Boca Raton: CRC Press.
- Betancourt, M. 2017. "A Conceptual Introduction to Hamiltonian Monte Carlo." arXiv preprint arXiv:1701.02434. URL: <https://arxiv.org/abs/1701.02434>.
- Baldwin, Breck. 2017. "RStan Getting Started." Online: <https://github.com/stan-dev/rstan/wiki/RStan-Getting-Started> accessed 8/21/2017.
- *Recommended:* Neal, Radford M. 2011. "MCMC using Hamiltonian Dynamics." Chapter 5 in Brooks, Gelman, Jones, and Meng, eds., *Handbook of Markov Chain Monte Carlo*: 113-162. URL: <http://www.mcmchandbook.net/HandbookChapter5.pdf>. Boca Raton: CRC Press.
- *Reference:* Stan Development Team. 2017. "Stan Modeling Language: User's Guide and Reference Model, version 2.16.0." URL: <https://github.com/stan-dev/stan/releases/download/v2.16.0/stan-reference-2.16.0.pdf> accessed 8/21/2017.

***** CUTOFF FOR EXAM 2 *****

9) Item Response Theory and the Scaling of Latent Dimensions

Questions: What is item response theory? How does it relate unobservable (latent) dimensions to observable behavior? What are potential applications of IRT to political science (e.g., what can we measure and how can the measurements be used)? When would we expect IRT to do a good job of measuring latent concepts, and when would we expect problems?

Skills and Concepts: origins and historical uses of IRT; estimation of IRT ideal points via MCMC; applications of IRT to Political Science; assessing the robustness of IRT estimates.

Readings:

- Ayala, R. J. 2009. *The Theory and Practice of Item Response Theory*. Guilford.
 - Chapters 1-2 and 5: “Introduction to Measurement,” “The One-Parameter Model,” and “The Two-Parameter Model.”
- Curtis., S., McKay. 2010. “BUGS Code for Item Response Theory.” *Journal of Statistical Software* 36: Code Snippet 1. URL: www.jstatsoft.org/v36/c01/paper/.
- Clinton, Joshua, Simon Jackman, and Douglas Rivers. “The Statistical Analysis of Roll Call Data.” *American Political Science Review* 98(2): 355-370.

10) Model Checking, Validation, and Comparison

Questions: How do we evaluate whether our procedures are working correctly? How do we detect and combat specification and convergence problems in our model? How do we compare models to one another?

Skills and concepts: posterior predictive checks; DICs; Bayes factors.

Readings:

- Gelman and Hill, Chapter 19
- Gelman et al. *Bayesian Data Analysis, Third Edition*:
 - Chapter 6: Model checking
 - Chapter 7: Evaluating, comparing, and expanding models

11) Missing Data Imputation

Questions: If observations are missing in your data set, what are the consequences for inference? What can be done to improve the situation? When does missing data imputation help, and when does it hurt?

Skills and Concepts: types of missingness in data; casewise deletion; simple imputation; multiple imputation; multiple imputation via chained equations; Bayesian data augmentation.

Readings:

- Gelman and Hill: Chapter 25
- van Buuren and Groothuis-Oudshoorn. 2011. “mice: Multivariate Imputation by Chained Equations in R.” *Journal of Statistical Software* 45(3). URL: <http://www.jstatsoft.org/v45/i03/>.

12) Multilevel Regression and Poststratification

Questions: How can we use Bayesian methods (specifically, hierarchical models and their “strength borrowing” characteristics) to estimate disparate relationships in small subsamples of a larger data set?

Skills and Concepts: multilevel regression and poststratification; survey analysis with hierarchical Bayesian models.

Readings:

- Park, David, Andrew Gelman, and Joseph Bafumi. 2004. “Bayesian Multilevel Estimation with Poststratification: State-Level Estimates from National Polls.” *Political Analysis* 12(4): 375-385.
- Lax, Jeffrey and Justin Phillips. 2009. “How Should We Estimate Public Opinion in the States?” *American Journal of Political Science* 53(1): 107-121.