

Bayesian Nonparametric Modeling of Individual Differences: A Case Study Using Decision-Making on Bandit Problems

Matthew D. Zeigenfuse (mzeigenf@uci.edu)

Michael D. Lee (mdlee@uci.edu)

Department of Cognitive Sciences
University of California, Irvine

Abstract

We develop and compare two non-parametric Bayesian approaches for modeling individual differences in cognitive processes. These approaches both allow major discrete differences between groups of people to be modeled, without making strong prior assumptions about how many groups are required. Instead, the number of groups can naturally grow as more information about the behavior of people becomes available. One of our models extends previous work by allowing continuous differences between people within the same group to be modeled. We demonstrate both approaches in a case study using a classic heuristic model of human decision-making on bandit problems, applied to previously reported behavioral data from 451 participants. We conclude that the ability to model both discrete and continuous aspects of individual differences in cognition is important, and that non-parametric approaches are well suited to inferring these types of differences from empirical data.

Keywords: Individual differences, non-parametric Bayesian modeling, bandit problems, win-stay lose-shift

Introduction

Individual differences in cognitive processes are basic, ubiquitous, and important. Almost all aspects of cognition, ranging from the simplest reaction time task, to the most involved problem-solving task, reveal systematic and meaningful variation in how different people perform. Entire fields are devoted to studying individual differences: the measurement and understanding of individual variation is the basic goal for research in psychometrics, including particularly the assessment of how people co-vary in their cognitive abilities.

There has, however, been less consideration of individual differences in experimental cognitive psychology, in the sense that it is rare to see theories of how people differ in a cognitive process directly incorporated into formal models. There is a general recognition that averaging across participants can be problematic when there are individual differences (e.g., Ashby, Maddox, & Lee, 1994; Estes, 1956), and often cognitive models of fit on an individual-participant basis, so that variations in model parameters can be observed and interpreted. But accounts of this variation are rarely formalized within the modeling, and so theories of individual differences are not yet fully incorporated in the modeling of cognition.

There are, nevertheless, a number of useful approaches—directly adapted from the statistics literature—that have been applied in the cognitive sciences to model individual differences. These include finding clusters of participants with different model parameters (e.g. Lee & Webb, 2005; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003) and hierarchical modeling approaches that make assumptions about the distri-

bution of model parameters across participants (e.g. Peruggia, Van Zandt, & Chen, 2002; Rouder, Sun, Speckman, Lu, & Zhou, 2003; Shiffrin, Lee, Kim, & Wagenmakers, 2008). The clustering approach focuses on discrete differences between people, capturing major or ‘qualitative’ differences between people, whereas the hierarchical approach focuses on continuous or ‘quantitative’ differences. In this sense, these two approaches are complementary, and could be combined in a natural way to allow for distinct groups of people who also show smaller individual variation within their groups.

One problematic property of these existing approaches, however, is that they require relatively strong prior assumptions about how many groups are needed to model the individual differences in observed performance. In contrast, such strong assumptions are not required by a newer approach to modeling individual differences—also borrowed from statistics—involving ‘non-parametric’ (also known as ‘infinite dimensional’) Bayesian modeling (Navarro, Griffiths, Steyvers, & Lee, 2006).¹ In non-parametric approaches, how a model represents individual differences can change as additional information from additional peoples become available. Intuitively, one might believe it is quite likely the second participant tested on an experimental task will be different from the first, but the likelihood the 41st participant will be different from all of the first 40 is far smaller. In this way, we expect the detail needed to express individual differences will depend upon the empirical evidence that is available, and grow with the number of participants. This means the complexity needed to model individual differences is not fixed, but inherently flexible. Non-parametric methods naturally have this flexibility, and so provide an intuitive and interesting perspective for modeling individual differences in cognition.

In this paper, we develop two non-parametric approaches for modeling individual differences, extending the previous work of Navarro et al. (2006). In particular, we develop a new non-parametric approach that can grow representations of both the discrete and individual aspects of individual differences. We develop and evaluate both of the modeling approaches in terms of a case study, involving human decision-making on bandit problems. This helps make our modeling approaches concrete, subjects them to a first empirical test,

¹The name ‘non-parametric’ suggests there are no parameter involved in modeling, which is not true. The name ‘infinite dimensional’ suggests that there are potentially infinitely many parameters, which is a much better conception. Nevertheless, the name ‘non-parametric’ is more widely used.

and also makes a contribution to understanding how people vary on an interesting and well-studied decision-making task.

The structure of this paper is as follows: First, we develop, in a general way, the infinite discrete and discrete and continuous approaches to modeling individual differences. We then apply these approaches to the specific case of modeling variation in a model of human decision-making on bandit problems. Using previously reported data from 451 participants on bandit problems, we evaluate both of the approaches. We conclude that it is important to be able to capture both discrete and continuous individual differences, and discuss the merits of our non-parametric models in this light.

Infinite Individual Differences Models

Model-based approaches to individual differences make explicit assumptions about the ways people can vary and covary. As described in the Introduction, the literature on model-based individual differences is dominated by two disjoint, but quite compatible, points-of-view, which we now formalize. The discrete individual differences (DID) modeling approach holds that there are $M < \infty$ distinct groups of individuals, where M is fixed, each possessing a unique value θ_z of the parameter of the cognitive model. If we let θ_i denote the cognitive model parameter of participant i , then $\theta_i = \theta_z$ for each participant in group z . Thus, we can think of DID models as picking out which individuals are alike and aggregating across only these individuals.

The continuous individual differences (CID) modeling approach holds that individuals are related through a continuous, typically uni-modal probability density P , specifying the relative probabilities of values of the parameters in a cognitive model. If we again define θ_i to be the cognitive model parameter of participant i , CID models of individual differences assume that $\theta_i \sim P$.

Navarro et al. (2006) developed what we call the *infinite discrete individual differences* (iDID) model². This model assumes there are an infinite number of groups to which individuals can be potentially belong, but that only a finite number of these will be manifest in finite data. In other words, $M = \infty$ for the population at large, but in a given experiment, we will only observe $m < \infty$ of these groups. Just as in the finite DID model, within any one of these groups, the parameter value for each individual is the group value.

To extend this approach to allow for individual differences, we take the following approach: Suppose we have $M < \infty$ groups of individuals. To each of these groups we associate a continuous, unimodal distribution P_z specifying how probable values of the parameter are in group z . Within group z , we take the parameters of each participant to be given by a CID model. In other words, for participant i in group z , $\theta_i \sim P_z$. We refer to this as the discrete and continuous individual differences (D&CID) model.

Unfortunately, this model suffers the same drawback as the

DID model, which is that the number of groups is a fixed rather than function of available information. Fortunately, however, it can be rectified in the same way. Suppose in the D&CID model we assume that $M = \infty$ for the entire population, but that in a given experiment, only $m < \infty$ of these groups will be observed. Then, as in that model, the cognitive model parameter for a particular participant is distributed as the group level distribution of the group to they belong. We term this the *infinite discrete and continuous individual differences* (iD&CID) model.

Application to Bandit Problems

In this section we apply the general iDID and iD&CID approaches to the specific problem of modeling human decision-making on bandit problems. First we describe bandit problems themselves, and the basic cognitive model we will use, and then we describe the individual differences models in detail.

Bandit Problem Decision-Making

Since their original mathematical formulation (Robbins, 1952), bandit problems have been studied extensively in the machine learning and statistics literatures (e.g. Berry, 1972; Brezzi & Lai, 2002; Gittens, 1989; Macready & Walpert, 1998), as a classic example of reinforcement learning, and in psychology as a task requiring people to balance the competing demands of exploration and exploitation (e.g. Banks, Olson, & Porter, 1997; Cohen, McClure, & Yu, 2007; Daw, O’Doherty, Dayan, Seymour, & Dolan, 2006; Meyer & Shi, 1995; Steyvers, Lee, & Wagenmakers, in press).

In a K -armed bandit problem, there is a sequence of N trials, on each of which a participant chooses one of K possible alternatives. Each arm $k \leq K$ offers the participant a reward with probability ϕ_k , which is fixed over the trial sequence, but not known by the participant. The goal for the participant is to make choices that maximize the total number of rewards attained over the N trials.

The Win-Stay Lose-Shift (WSLS) heuristic is a classic account of decision-making on bandit problems (Robbins, 1952), and can be used as a basic cognitive model to understand human behavior (Steyvers et al., in press). In its deterministic form, it says people choose a bandit arm k so long as that arm continues to give them a reward. Thus, if a participant received a reward choosing arm k on the previous trial, WSLS says the participant will choose k again. If they did not receive a reward, WSLS says they will move from arm k to arm $k' \neq k$.

We employ a probabilistic generalization of the WSLS heuristic. Instead of always staying with arm k after receiving a reward, and always switching to k' after not receiving one, we assume people will stay on arm k with (high) probability θ after receiving a reward, and switch with the same probability θ after not receiving one. In this way, the parameter θ can be conceived as an ‘accuracy of execution’ parameter that measures how faithfully the basic deterministic WSLS heuristic is applied in practice.

²In their paper, Navarro et al. (2006) refer to this as the *infinite groups model*.

Non-parametric Models

In this section we develop iDID and iD&CID models of individual differences in human bandit problem behavior using the WSLs heuristic. The two models share the same cognitive model, WSLs, but differ in how they model the variation between individuals in the parameter of that model. Both models assume there are an infinite number of groups; however, the iDID model assumes that within one these groups individuals do not vary whereas the iD&CID model assumes that, within a group, individuals vary continuously with respect to a unimodal distribution.

In order to understand better the two models and their relation to each other, it is helpful describe how each generates data. Both can be thought of as probabilistic generative models of human decision-making on the bandit problem task, operating in three stages. First, an assignment of participants to groups is sampled. Then, each participant’s cognitive model parameter is sampled given this assignment. Finally, each participant’s observed data is sampled given their model parameter. The remainder of this section is devoted to discussing each of these steps in greater detail.

Cognitive Model On trial $t \geq 2$ of a bandit problem (WSLs cannot be applied on the first trial, and the model itself assumes guessing on the first trial) a participant chooses the same alternative as on the previous trial with probability θ_i , given that a reward was received on that trial. If a reward was not received, they will switch to another alternative with the same probability.

For each participant, let the observed data, x_i , be the number of times that participant applied the WSLs heuristic. This means the data follow a binomial distribution, with

$$x_i \sim \text{Bin}(N, \theta_i). \quad (1)$$

Parameters Next we specify how the θ_i are generated. In the iDID case, a participant’s probability of following WSLs is simply the probability for the group of which they are a member. Thus, if participant i is a member of group z , $\theta_i = \theta_z$, where θ_z is the probability of following WSLs in group z . We take the θ_z to be independently uniformly distributed on the unit interval. In other words, *a priori* we know nothing about the value of θ_z for each group z except that it lies between 0 and 1.

The iD&CID case is slightly more complicated. In this case, each participant’s WSLs probability is a random draw from some unimodal group distribution. If participant i is a member of group z , $\theta_i \sim P_z$, where P_z is the group distribution for group z , which we take to be a Beta distribution with shape parameters a_z and b_z . The Beta distribution is a commonly used density on the unit interval as it has a number of desirable statistical properties. Moreover, in this case, the parameters have an intuitive interpretation. Suppose we ran some experiment before running the current one. Then a_z can be thought of as the number of times in this previous experiment participants from group z used WSLs and b_z can be thought of the number of times they did not.

It remains to specify a prior distribution on the pair (a_z, b_z) . In this we follow Gelman, Carlin, Stern, and Rubin (2003) and Steyvers et al. (in press), defining flat prior distributions over the mean, $a_z/(a_z + b_z)$, and the square root of the inverse ‘sample size’, $(a_z + b_z)^{-1/2}$, of the Beta distribution. Converting to a joint distribution over (a_z, b_z) , we get

$$p(a_z, b_z) \propto (a_z + b_z)^{-5/2}. \quad (2)$$

Assignments Finally, we specify how groups assignments are generated. Both the iDID and iD&CID modeling approaches rest on the assumption that there are an infinite number of ways in which individuals can potentially vary, only a finite number of which will ever manifest in finite data. The *Chinese restaurant process* (Aldous, 1985; see Navarro et al., 2006, for an introduction aimed at cognitive scientists) is a prior distribution which implements this idea in a probabilistic way.

The Chinese restaurant process operates as follows. Suppose we have a Chinese restaurant containing an infinite number of tables each with an infinite capacity. These tables are assumed to be distinguishable only by which customers are seated at them. When the first customer walks in they are seated at the first table (we can of course pick a table arbitrarily to be first since the tables are indistinguishable). For each subsequent customer one of two things happens: (i) the customer is seated at a previously seated table, or (ii) the customer is seated at a new table. For each previously seated table, the new customer is seated at that table with probability proportional to the number of customers already seated at that table. The new customer is seated at a new table with probability proportional to a constant $\alpha > 0$. In this analogy, the customers are the participants in our experiment, and the tables represent groups with individual differences to which they may belong.

One issue remains with regard to assignment. Clearly, the magnitude of α affects the number of tables seated since increasing the magnitude of α increases the probability of seating a new table. Hence, increasing the value of α increases the number of tables we expect to seat *a priori*, which may in turn affect the number of tables we see *a posteriori*. To deal with this, Antoniak (1974) suggests placing a prior distribution on α . Following Escobar and West (1995) and Navarro et al. (2006), we use an inverted Gamma distribution.

Inference Methods

Inference on the model was performed numerically using Markov Chain Monte Carlo (MCMC) posterior sampling methods in two stages. In the first, the posterior distribution over assignments was sampled using Gibbs sampling, for the iDID model, and a Gibbs sampling scheme with a Metropolis-Hastings step, for the iD&CID model. In the second, parameter values for the models were sampled given particular assignments. For the iDID model, the posterior was sampled exactly using beta-binomial conjugacy. For the iD&CID model, posterior sampling required another Gibbs sampler to integrate across the group level distributions.

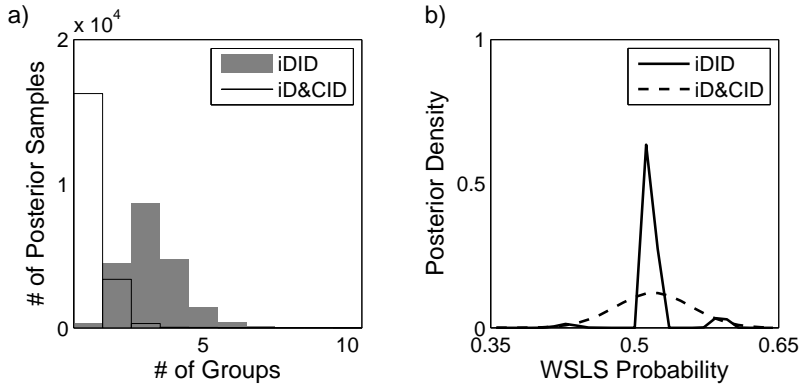


Figure 1: a) Posterior distributions over number of groups for iDID and iD&CID models. b) Posterior distributions over WSLs probability for iDID and iD&CID models conditional upon their MAP group estimates of 3 and 1, respectively.

Results for Bandit Problem Data

We applied the two models to data collected by Steyvers et al. (in press). Their experiment consisted of 451 participants who each completed 20 bandit problems with 15 trials and 4 alternatives. For each problem, both reward rates were chosen from a Beta distribution with mean 1/2 and sample size 4. This resulted in a 20 pairs of reward rates, used in a randomly permuted order for each of the 451 participants.

In their analysis, Steyvers et al. (in press) compared four models of human bandit performance: a guessing model that assumed people chose at random; WSLs, as we have described; a success ratio model that assumed people chose alternatives based on their ratios of successes to failures; and the optimal decision process, which can be found by standard recursive programming methods (e.g., Kaelbling, Littman, & Moore, 1996). Steyvers et al. (in press) found using, Bayes Factors, WSLs fit 47% of participants better than the other three models. This suggests that, despite its simplicity, WSLs is a good model of these participants on the task.

Our analysis proceeds in three stages. In the first, we perform inference on the number of groups using the iDID and iD&CID models. In the second, we fix the number of groups for each model, and perform inference over the most likely assignment of people to groups and the probabilities of applying WSLs given this assignment. Finally, we generate data given these distributions WSLs probabilities, and use these predictive data to assess the fit of each model.

Number of Groups

The marginal posterior distributions over the number of groups inferred by each of the iDID and iD&CID models are shown in Figure 1a. The former has a mode at three groups, indicating the assignments sampled by the iDID model consist of three groups more often than any other number. The latter has a mode at one group. In fact, the iD&CID model samples the assignment placing all participants in a single group about 75% of the time.

Figure 1a illustrates an important distinction between the two models, which is that the number of groups inferred by the iDID model is stochastically greater than that of the iD&CID model. If k_S and k_H are random variables denoting the number of groups present in an assignment under the iDID and iD&CID models respectively, k_S is stochastically greater than k_H means $p(k_H > K) \leq p(k_S > K)$ for any $K \geq 1$. This occurs because the iD&CID model allows within group variability in the probability of applying WSLs that the iDID model does not. Within group variability in this parameter leads to group distributions in the iD&CID model giving non-negligible mass to a wider range of data values than the iDID model, making distinct individuals more likely to belong to the same group under the iD&CID model.

Parameter Inference

We now focus our attention on inference about the distribution of WSLs probabilities given a fixed number of groups. For each model, we fix the number of groups to its modal value (3 for the iDID model, 1 for the iD&CID model, see Figure 1a) and determine the maximum *a posteriori* (MAP) assignment of individuals to groups. Finally, we infer the conditional distributions over the probabilities of applying WSLs given this MAP assignment. Figure 1b shows the distribution over the conditional probabilities of applying WSLs. In the figure, each mode corresponds to a single group³; thus, the three groups of the iDID model have modes

³Though similar, the densities shown in Figure 1b are not, in fact, with respect to the same quantity. In the iDID model, each group has a single WSLs probability θ_z and the observed data for every member of that group follow a binomial distribution with rate θ_z . Thus, the density in the iDID plot should be thought of as depicting the uncertainty as to where the three modes lie rather than a sampling distribution for the individual θ_j . In the iD&CID model, individuals within a group are not constrained to use exactly the same θ_z , but instead to follow the same unimodal distribution. That being the case, the distribution shown in the figure should be interpreted as the expected sampling distribution the θ_i (since it is averaged across the full joint posterior distribution of its parameters).

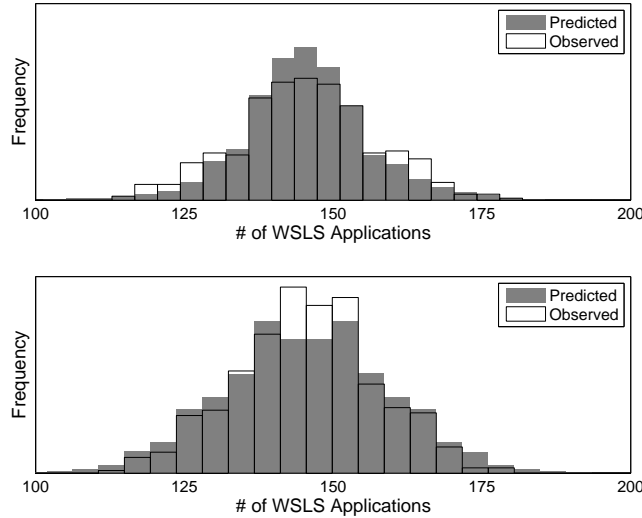


Figure 2: Posterior predictive distribution over the number of applications of WSLs conditional on the groups shown in Figure 1 plotted against the observed data for the iDID (top) and iD&CID (bottom) models.

near 0.4, 0.51, and 0.6 and the single group of the iD&CID model has a mode near 0.51. This shows how the iDID model subdivides participants into groups: there is an “average” probability of applying WSLs group, into which most participants fall, a “high” probability of applying WSLs group, into which participants well-described by WSLs fall, and a “low” probability of applying WSLs group, into which participants poorly described by WSLs fall. In contrast, Figure 1b shows why the iD&CID model does not require these groups to account for the data. Because it allows individuals within a group to vary continuously across a range of values, the iD&CID model is able to capture the “average” group with the center of its single group distribution and the “high” and “low” groups with the tails.

In addition to illustrating the group structure in parameter space of the two models given their respective conditional MAP assignments, Figure 1b shows how each model fits the observed data using the groups and their WSLs probability distributions. The central mode of the iDID model is approximately equal to the mode of the iD&CID model and the densities of the upper and lower modes of the iDID model correspond to the densities of tails of the iD&CID model. For the iDID, intuitively, this tells us that the iDID is doing its best to capture variation in the data—perhaps more naturally captured by the continuous iD&CID model—by finding the best placement of group modes. For the iD&CID model, the model tries to balance peakedness near the mode of the data against the rate at which mass falls off away from the mode.

Model Fit

Figure 2 shows the posterior predictive distributions over the number of applications of WSLs for the iDID (top) and iD&CID (bottom) models conditional on the MAP assignment of individuals to groups. For both models, the posterior

predictive distribution is the distribution of the number of applications of WSLs averaged across the posterior distribution of the probability of applying WSLs. Comparing these distributions to the observed data offers a standard Bayesian posterior predictive check (e.g. Gelman, Carlin, Stern, & Rubin, 2004) for the models.

Figure 2 shows both models are able to capture general features of the observed data, such as the shape of the histogram, but not all of the specifics. The iDID is too peaked, as it overestimates the masses of points near the mode and underestimates those of points near the tails. Alternatively, the iD&CID is too flat. Though it predicts the masses of the observed data away from the mode well, it is not peaked enough at the mode. Overall, however, we would argue that the iD&CID model fits the data better, because none of the observed data are given low predictive probability. There is a sense in which the iDID is too confident in its assessment of the variation in human performance, because it specifies too narrow a range in human performance. This could be regarded as a form of over-fitting. The iD&CID, in contrast, is too vague, and so, while giving high probability to the modal data values, slightly under-estimates their magnitude. This could be regarded as a form of under-fitting. A general rule in modeling is that over-fitting is dangerous, because it makes you think you know more than you really know, while under-fitting is relatively harmless (Grünwald, 2007).

Discussion

Our results suggest that both the iDID and iD&CID are good accounts of individual differences with respect to the WSLs model employed here, but that the iD&CID is better. Both models are able to fit the data reasonably well, as measured by posterior predictive distributions. But the iD&CID model is able to capture the single group structure,

and fit the pattern in the observed data better.

In terms of bandit problem performance, this paper builds upon the results of Steyvers et al. (in press) by showing that, among those individuals applying WSLS, the winning and staying or losing and shifting is not the same for all people, or even subsets of people. Rather, the suggestion is that people exhibit a wide range of behaviors on bandit problems, and that multiple models will probably be necessary to explain human behavior fully. Future bandit problem work should focus on evaluating numbers of different heuristic models, and partitioning participants into groups to capture variations in the way those models are applied, using accounts of individual differences similar to those presented here.

More generally, we have presented a new approach to modeling of individual differences, iD&CID, and compared this to an existing model, iDID in a concrete way. We found that the iD&CID model was better able to account for both the group structure and distributional pattern of the data, suggesting the larger-scale applicability of the iD&CID model to the general problem of modeling individual differences in human cognitive processes.

Acknowledgments

This work is was supported by Air Force Office of Scientific Research award FA9550-07-1-0082 to Michael Lee and Mark Steyvers. We thank Dan Navarro for helpful discussions.

References

- Aldous, D. J. (1985). Exchangeability and related topics. In P. L. Hennequin (Ed.), *École d'été probabilités de saint-flour XII*. Berlin: Springer.
- Antoniak, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Annals of Statistics*, 2, 1152-1174.
- Ashby, F. G., Maddox, W. T., & Lee, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science*, 5, 144-151.
- Banks, J., Olson, M., & Porter, D. (1997). An experimental analysis of the bandit problem. *Economic Theory*, 10, 55-77.
- Berry, D. A. (1972). A Bernoulli two-armed bandit. *The Annals of Mathematical Statistics*, 43(3), 871-897.
- Brezzi, M., & Lai, T. L. (2002). Optimal learning and experimentation in bandit problems. *Journal of Economic Dynamics & Control*, 27, 87-108.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? Exploration versus exploitation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 933-942.
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876-879.
- Escobar, M. D., & West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90, 577-588.
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, 53(134-140).
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2003). *Bayesian data analysis*. NY: Chapman and Hall.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian data analysis (2nd ed.)*. Boca Raton (FL): Chapman & Hall/CRC.
- Gittens, J. C. (1989). *Multi-armed bandit allocation indices*. New York: Wiley.
- Grünwald, P. D. (2007). *The minimum description length principle*. Cambridge, MA: MIT Press.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237-285.
- Lee, M. D., & Webb, M. R. (2005). Modeling individual differences in cognition. *Psychonomic Bulletin & Review*, 12(4), 605-621.
- Macready, W. G., & Walpert, D. H. (1998). Bandit problems and the exploration/exploitation tradeoff. *IEEE Transactions on Evolutionary Computation*, 2(1), 2-22.
- Meyer, R. J., & Shi, Y. (1995). Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Science*, 41(5), 817-834.
- Navarro, D. J., Griffiths, T. L., Steyvers, M., & Lee, M. D. (2006). Modeling individual differences using Dirichlet processes. *Journal of Mathematical Psychology*, 50, 101-122.
- Peruggia, M., Van Zandt, T., & Chen, M. (2002). Was it a car or a cat I saw? An analysis of response times for word recognition. In C. Gatsonis et al. (Eds.), *Case studies in Bayesian statistics* (Vol. 6, p. 319-334). New York: Springer.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 55, 527-535.
- Rouder, J. N., Sun, D., Speckman, P. L., Lu, J., & Zhou, D. (2003). A hierarchical Bayesian statistical framework for skewed variables with an application to response time distributions. *Psychometrika*, 68.
- Shiffrin, R. M., Lee, M. D., Kim, W.-J., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32(8), 1248-1284.
- Steyvers, M., Lee, M. D., & Wagenmakers, E. J. (in press). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453-487.