

Modeling when people quit: Bayesian censored geometric models with hierarchical and latent-mixture extensions

Kensuke Okada¹ · Joachim Vandekerckhove² · Michael D. Lee²

Published online: 31 March 2017
© Psychonomic Society, Inc. 2017

Abstract People often interact with environments that can provide only a finite number of items as resources. Eventually a book contains no more chapters, there are no more albums available from a band, and every Pokémon has been caught. When interacting with these sorts of environments, people either actively choose to quit collecting new items, or they are forced to quit when the items are exhausted. Modeling the distribution of how many items people collect before they quit involves untangling these two possibilities. We propose that censored geometric models are a useful basic technique for modeling the quitting distribution, and, show how, by implementing these models in a hierarchical and latent-mixture framework through Bayesian methods, they can be extended to capture the additional features of specific situations. We demonstrate this approach by developing and testing a series of models in two case studies involving real-world data. One case study deals with people choosing jokes from a recommender system, and the other deals with people completing items in a personality survey.

Keywords Modeling quitting · Censored geometric distribution · Search termination · Hierarchical models · Latent-mixture models · Bayesian methods

Introduction

The world is full of environments in which people can collect some or all of a finite set of items, and it is a basic psychological question to ask when and why they quit collecting. Many people have read every chapter in the Lord of the Rings, but some stopped part way through. Many people own some Fleetwood Mac albums, and a few people own all of them. Most people gave up on playing Pokémon GO when creatures were still available, but a few heroic individuals persevered to catch them all. In all of these cases, there is a distribution of the number of items collected across all people, and that distribution has a hard limit at the total number of available items.

A key feature of these distributions is that they combine two fundamentally different reasons for people finalizing their collection. The first involves people choosing to quit collecting items. The other involves them being forced to quit because all available items are exhausted. The first is a property of the individual, involving their preferences, choices, and decision making, while the second is a property of the environment. One way to understand the situation is that the limited environmental resources limit the measurement of people's quitting behavior. People may have been willing to read more chapters, buy more albums, or collect more Pokémon, but, if they have already collected them all, that willingness cannot be quantified because of the limitations of the environment.

In this paper, we propose that censored geometric models form a useful starting point for modeling people's quitting behavior, and demonstrate how they can be extended to capture the more detailed features of specific situations. Statistical methods for censored data are well-established,

✉ Michael D. Lee
mdlee@uci.edu

¹ Department of Psychology, Senshu University, Tokyo, Japan

² Department of Cognitive Sciences, University of California, Irvine, CA 92697-5100, USA

and widely used in some area of empirical sciences, especially for problems that involve survival or adoption (e.g., Amemiya 1984; Cox & Oakes 1984; Klein & Moeschberger 2005). Previous applications of censored geometric distributions range from biological applications like bird nesting success (Green, 1977) and species abundance (Norris & Pollock, 1998), to economics applications like technology adoption (McWilliams et al., 1998). The basic idea is that there is some probability of a terminating event—whether a success in the context of adoption, or a failure in the context of quitting—at every step, and a bound on how many steps are possible. The process giving the probability of termination corresponds to the geometric distribution, while the bound of how many steps can be observed corresponds to the censoring. These properties are a natural characterization of our problem of understanding the distribution of quitting as a mixture of people’s choices and environmental limits.

We present two case studies. In the first, we model how many jokes people read from an on-line recommender system that has a total of 70 jokes. In the second, we model how many questions people answer in an on-line personality survey that has a total of 322 relevant questions. In both case studies, we start with the basic censored geometric model, and then develop and evaluate extensions of the basic model—using hierarchical and latent-mixture approaches, and Bayesian methods—trying to achieve descriptive adequacy and insight into the underlying cognitive processes involved in quitting.

First case study: Quitting reading jokes

The first case study is based on data from the Jester on-line system (Goldberg et al., 2001). This system requires users to read and rate a small set of jokes, and then allows them to continue reading and rating the other jokes in the system. Since its inception in 1999, the Jester system has evolved, with the stage-wise addition of jokes, and the continual collection of data from new users.

We focus on a small subset of the available data, from a time when there were 70 jokes in the system, and people had to read and rate the same 15 jokes. The behavior of 2,607 users is available for this snapshot of the system. The distribution of the number of jokes read for these people is shown in Fig. 1, and displays a clear pattern. After the required 15 jokes have been read, there is a group of people who gradually give up, and a group of people who read all 70 jokes. The research problem is to understand this distribution of quitting.

A simple censored geometric model

A simple plausible model says that people have a probability θ of making the current joke their last one, and assumes that this probability is the same for all people. It also assumes that this quitting probability does not change regardless of how many jokes have already been read. According to this model, the number of jokes α_i that the i th person would read from a system with endless jokes is

$$\alpha_i \sim 15 + \text{Geometric}(\theta).$$

For the current data, however, the maximum number of jokes that can be read is 70. Thus, the Jester on-line system, viewed as a measuring instrument of people’s desire to read jokes, censors the measurement of α_i to a maximum of 70. This means that the observed data y_i shown in Fig. 1, counting the number of jokes seen by the i th person, is:

$$y_i = \begin{cases} \alpha_i & \text{if } \alpha_i \leq 70 \\ 70 & \text{if } \alpha_i > 70. \end{cases}$$

To allow Bayesian inference, the model is completed by a prior $\theta \sim \text{Uniform}(0, 1)$ on the probability of termination after each joke. We implemented this model in JAGS (Plummer, 2003), which uses Markov-chain Monte-Carlo (MCMC) methods to return samples from the joint posterior parameter distribution conditional on the model and data.

The results of applying this model are shown in Fig. 2. The black line shows the posterior predictive distribution of the model, which provides an indication of its descriptive adequacy (Shiffrin et al., 2008). It is clear that it describes the distribution of the observed data well, and captures both the gradual decrease from 15 to 69 jokes, and the large number of people who read all 70 jokes. The posterior expectation of the quitting probability θ is about 3%. This means there is about a 3% chance, as a person reads a joke, that it will be their last.

Figure 2 provides evidence that the simple model of how many jokes people read is a very useful one. It accounts for the gradual drop-off in people reading 15 to 69 jokes, as well as the large group of people who read all of the jokes. The psychological explanation for the spike at 70 is simply that many people wanted to keep reading jokes, but the system ran out. In this way, the data are explained as a combination of people’s quitting behavior, and the limitations of the system.

An extended model with individual differences

While the simple censored geometric model is able to describe the Jester data well, it makes a number of strong psychological assumptions that are worth testing. In particular, it




Fig. 1 The distribution of the number of jokes read by 2,607 people with access to a total of 70 jokes in the Jester on-line recommendation system

assumes there are no individual differences between people. It seems plausible that there are differences in the quitting probabilities for different people. It could also be that some people use a fundamentally different strategy, reading all of the jokes, regardless of how many the system can provide. In the same way people usually read all the chapters in a book, exhausting its content, it could be that some people aim to read all of Jester's jokes.

It is possible to allow for these sorts of individual differences using hierarchical and latent-mixture extensions of the basic model (Bartlema et al., 2014). The hierarchical part of the extension allows for the possibility that each person has their own termination probability, denoted θ_i for the i th

person, drawn from a group distribution. We assume the group distribution is a truncated Gaussian between 0 and 1, so that

$$\theta_i \sim \text{Gaussian}_{(0,1)}\left(\mu, \frac{1}{\sigma^2}\right),$$

and give the mean and standard deviation relatively uninformative priors $\mu \sim \text{Uniform}(0, 1)$ and $\sigma \sim \text{Uniform}(0, 1)$.

The latent-mixture part of the extension allows for the existence of two qualitatively different groups of people. People in the first group use the hierarchical version of the simple model, terminating with probability θ_i or being censored to read all of the jokes. People in the second group




Fig. 2 The distribution of the number of jokes read by people, and the posterior predictive agreement of the censored geometric model

always read all of the jokes. Psychologically, the two trends seen in the data are now accounted for in terms of two groups of people, one group giving the gradual decline from 15 to 69, and the other giving the spike at 70.

A binary indicator parameter z_i is used for the i th person to indicate to which of the two groups they belong. These parameters are assumed to follow a base-rate ϕ representing the proportion of people in the “read all the jokes” group, so that

$$z_i \sim \text{Bernoulli}(\phi),$$

where the base-rate is given the prior $\phi \sim \text{Uniform}(0, 1)$.

The observed data y_i are now distributed according to which group the i th person belongs. If they belong to the first group, the data are distributed according to the simple model, but with termination probability θ_i . If they belong to the second group, they always read all 70 jokes.

$$y_i \sim \begin{cases} 15 + \text{Geometric}(\theta_i) & \text{if } z_i = 0 \text{ and } \alpha_i \leq 70 \\ 70 & \text{if } z_i = 0 \text{ and } \alpha_i > 70 \\ 70 & \text{if } z_i = 1. \end{cases}$$

Once again, we implemented this model in JAGS, and applied it to the Jester data. The important part of the posterior distribution, to test for the possibility of individual differences, is the joint posterior of the σ and ϕ parameters. When $\sigma = 0$, there is no variability in the termination probabilities over people, and so that type of individual difference does not exist. When $\phi = 0$, the base-rate of people in the “read all the jokes” group is zero, and so that type of individual difference does not exist. When both $\sigma = 0$ and $\phi = 0$, both types of possible individual difference are removed, and the extended model reduces to the original simple one.

Figure 3 shows the joint posterior distribution over σ and ϕ , as a two-dimensional plot. The joint posterior is a three-dimensional surface, and Fig. 3 can be thought of as viewing this surface from directly above as a grid. The area of each black circle is proportional to the density of the prior for each combination of σ and ϕ , and shows the uniform priors that were assumed. The area of each gray circle is proportional to the density of the posterior in that region of the parameter space. It is clear that essentially all of the posterior density is near the point $\sigma = 0$ and $\phi = 0$.

According to the Savage-Dickey method, the ratio of the posterior to the prior at a point in the parameter space approximates a Bayes factor. This Bayes factor compares a specific nested model—the one obtained by setting the parameters of the more general model to the specific values—to the more general model itself (Wagenmakers et al., 2010). Here, the interest is in the Bayes factor

comparing the original simple model to the hierarchical latent-mixture extension. We approximated this by taking a small region ϵ around the point ($\sigma = 0, \phi = 0$), and counting the number of samples within this region in the posterior and prior. The logarithm of the Bayes factor was estimated as about 10 in favor of the simple model, which we interpret as very strong evidence.

In other words, the current data do not provide evidence for different people having different termination probabilities, nor for there being a different group of people who want to read all of the jokes. The data are well described by a simple censored geometric model of termination.

Second case study: Quitting answering questions

The Synthetic Aperture Personality Assessment (SAPA) test collects people’s responses to a series of personality-assessment questions through an on-line system (Condon & Revelle, 2015; Revelle et al., 2012). The system includes a total of 696 questions, collated from 92 public-domain personality scales. People can complete up to 14 pages of these questions. Navigation through the system involves clicking a “next” button at the end of each page to move to the next, or a “quit” button at the end of the page to quit answering questions.

The first 4 pages of the SAPA system each have 18 personality questions, and the remaining 10 pages each have 25 personality questions. Thus, people can complete up to 322 personality questions over the 14 available pages. The instructions given by the system encourage people to complete the first 4 pages, corresponding to a total of 72 personality questions.

The data we analyze come from Condon and Revelle (2015), and detail the responses to personality questions made by 23,680 people who answered at least one question using the SAPA system. As in the Jester case study, we are not interested in the content of people’s answers, but in how many pages and questions they complete. The distribution of the number of questions answered across all people is shown in Fig. 4, with the tick marks on the x -axis corresponding to the total number of questions available by accessing each successive survey page. The distribution of the number of questions people answer has an interesting structure, with peaks corresponding to question counts near page boundaries, and an especially large peak around 72 questions, which corresponds to completing the first 4 pages.

A simple censored geometric model

Applying a simple censored geometric model to this problem is naturally done at the level of survey pages, since it is




Fig. 3 The prior and posterior distributions over individual differences in termination rate, σ , and the base-rate of group differences, ϕ . The prior is shown by black circles, while the posterior is shown by gray circles

at the end of each page that people can quit. If the probability of quitting at the end of a page is θ , then the number of pages the i th person attempts in a system with limitless pages is given by

$$\alpha_i \sim 1 + \text{Geometric}(\theta).$$

Because the system provides a maximum of 14 pages, however, the number attempted by the i th person, p_i is censored as

$$p_i = \begin{cases} \alpha_i & \text{if } \alpha_i \leq 14 \\ 14 & \text{if } \alpha_i > 14. \end{cases}$$




Fig. 4 The distribution of the number of personality items completed by 23,680 people with access to a total of 322 items

Due to the structure of the survey, the total number of questions q_i that the i th person could answer, over the p_i pages they attempt, is given by

$$q_i = \min(p_i, 4) \times 18 + \max(p_i - 4, 0) \times 25.$$

It is clear from the distribution of the data in Fig. 4, people do not always answer all of the questions on a page. At the question counts that correspond to page boundaries, there are usually some people who answered slightly fewer questions than the total for all of the pages to that point. For example, in the group who accessed 4 survey pages, many people complete slightly fewer than 72 questions. This observation suggests that people sometimes skip questions, whether by choice or oversight. A simple way to model this skipping is to assume there is a probability ψ of answering any available question. This means that the observed number of questions y_i answered by the i th person is

$$y_i \sim \text{Binomial}(\psi, q_i).$$

The model is completed with uniform priors on the page-quitting and question-skipping probabilities

$$\theta \sim \text{Uniform}(0, 1),$$

$$\psi \sim \text{Uniform}(0, 1).$$

We implemented this model in JAGS and applied it to the SAPA data. The results are shown in Fig. 5, with the posterior predictive distribution of the model overlaid on the distribution of the data. It is clear that the simple model fails the basic test of descriptive adequacy. The posterior predictive distribution is unable to match the data on which it is based.




Fig. 5 The distribution of the number of items read by people, and the posterior predictive agreement of the censored geometric model

An extended model with an instruction-following group

At least one reason for the failure of the simple censored geometric model is obvious. This reason is that it does not incorporate the relevant information regarding the system's instruction to complete 4 pages. It is clear from the distribution of the number of questions answered that many people quit after 4 pages.

To incorporate this information, we extend the model to consider two groups of people, using a latent-mixture approach. The first group of people continue to be modeled by the original simple censored geometric account, and can attempt any number of survey pages from 1 to 14. The second group of people follow the default instructions and attempt exactly 4 pages. In this extended model, a binary indicator parameter z_i indicates to which of these groups the i th person belongs. The number of pages attempted by the i th person now becomes

$$p_i = \begin{cases} \alpha_i & \text{if } z_i = 0 \text{ and } \alpha_i \leq 14 \\ 14 & \text{if } z_i = 0 \text{ and } \alpha_i > 14 \\ 4 & \text{if } z_i = 1. \end{cases}$$

The indicator z_i is assumed to follow a base-rate ω , which represents the proportion of people who follow the “just answer four pages and quit” instructions

$$z_i \sim \text{Bernoulli}(\omega),$$

and the base-rate is given a uniform prior

$$\omega \sim \text{Uniform}(0, 1).$$

The remainder of the model, generating the observed number of questions completed q_i , is the same as the original model.

The results for the extended latent-mixture model are shown in Fig. 6. The posterior predictive distribution now matches the distribution of the data much more closely. In particular, the large number of people answering 72-or-slightly-fewer questions is now much better described.

The extended model still, however, has a number of systematic failures in its descriptive adequacy. It does not include as many people in some of the “peaks” as are observed in the data, especially those associated with quitting after 1 or 4 pages. Even more tellingly, it mis-describes the distribution of questions answered for the cases where people viewed many pages, and this mis-description appears to become progressively worse for larger numbers of pages. It is especially clear for the distribution of the observed data around 300 questions, which is not well captured by the posterior predictive distribution.

An extended model with changing question-completion rates

The reason for the failure of the latent-mixture model is subtle. The key insight is given by the progressively-worsening alignment of predicted and observed distributions of question completions for the peaks as the number of pages viewed progresses. The observed data progressively fall further short of the model’s posterior predictive distribution. This means that people complete fewer questions than the latent-mixture model expects, as the number of pages increases. One possible cause of this mismatch is that the question-completion rate is not constant, but decreases as the number of pages viewed increased. This seems psychologically plausible, corresponding to the idea that people

skip or miss more questions as they view more pages of the survey.

To test this idea, we extended the latent-mixture model hierarchically, to allow for the question-completion rate to change as a function of the number of pages viewed. We experimented with a number of different functional forms for the relationship between the question-completing rate ψ_k on the k th page, including a bounded-linear function. To describe the data well, it seems to be important to allow for a non-linear relationship allowing for the decrease in question completion rates to fall more quickly over pages. To allow for this, we used a standard logistic function

$$\psi_k = \frac{1}{1 + \exp(-\beta_1 k - \beta_0)},$$

with standard priors

$$\beta_1 \sim \text{Cauchy}(0, 1),$$

$$\beta_0 \sim \text{Cauchy}(0, 1).$$

The remainder of the model, generating the observed number of questions completed q_i , is the same as in the extended latent-mixture model.

We implemented this hierarchically-extended model allowing for changing question-completion rates in JAGS. The results are shown in Fig. 7. The posterior predictive distribution now matches the distribution of the data very closely. In particular, the distribution of people completing around 300 questions, consistent with viewing all 14 survey pages, but failing to complete many questions over the pages, now agrees well with the data.

The inferred form of the relationship between question-completion rate and number of pages read is shown in the




Fig. 6 The distribution of the number of items read by people, and the posterior predictive agreement of an extended censored geometric model including in a latent-mixture a group of people who attempt exactly 4 survey pages




Fig. 7 The distribution of the number of items read by people, and the posterior predictive agreement of the censored geometric model with changing question-completion rates

left-hand panel of Fig. 8. This shows the posterior distribution of the logistic function based on the joint posterior of the β_1 and β_0 parameters. The probability of completing a question is inferred to decrease from about 0.99 to about 0.93 over the course of the 14 pages, with an accelerating decrease for the last few pages.

The finding of a decreasing question-completion rate is psychologically plausible, and leads to improved posterior predictive agreement. It can also be tested formally as a model-selection problem, using Bayes factors. The relevant comparison is between a model that assumes a constant question-completion rate with the model that assumes a logistic decrease. To estimate this Bayes factor, we again relied on the Savage-Dickey method, which is applicable

because when $\beta_1 = 0$, the logistic model reduces to the constant model.

The right-hand panel of Fig. 8 summarizes this analysis, showing the prior and posterior distributions for β_1 . It is clear the posterior has negligible density relative to the prior at the critical point $\beta_1 = 0$, and we estimate the logarithm of the Bayes factor in favor of the logistic model to be about 600, which is overwhelming evidence. To test the robustness of this conclusion to the prior, we repeated the analysis for $\beta_1 \sim \text{Cauchy}(0, 2.5)$ prior, as recommended by Gelman et al. (2008), and got a similarly overwhelming result.

Our conclusion from this sequence of model building and testing is that the basic censored geometric model needs to




Fig. 8 The posterior distribution of the logistic function based on the joint posterior of the β_1 and β_0 parameters (*left-hand panel*) and the prior and the posterior distributions for β_1 used to estimate the Bayes factor using the Savage-Dickey method (*right-hand panel*)

be extended in two ways to describe the distribution of questions answered in the SAPA data. The first extension is to include a group of people who follow the default instructions and view 4 pages. This is naturally achieved using latent-mixture modeling. The second extension is to allow for a non-linear decrease in the rate at which questions are answered as more pages are viewed. This is naturally achieved using hierarchical modeling. The resultant hierarchical latent-mixture model provides an accurate and interpretable description of the data.

Discussion

Collectively, the two case studies we have presented highlight the usefulness of censored geometric models for understanding people's quitting behavior. The modeling challenge in the joke-reading and question-answering applications is that the basic psychological questions about people's preferences and individual differences are not immediately answerable from their observed behavior. The limitations of the Jester and SAPA systems as measurement instruments mean that it is not obvious when people quit of their own accord, and when the system forces them to. This interaction between psychology and measurement requires the development of models that incorporate the censoring properties of the environment. Our case studies show how a basic censored geometric model can be applied, and then extended and tested in accordance with the details of a specific problem.

One of the interesting features in the two case studies is that multiple meaningful parameters are estimated from a single observed variable y_i , which is the observed number of jokes read or questions answered. In all of the models we considered, the quitting rate parameter θ of the geometric distribution is inferred. Models that included latent-mixture extensions also inferred a base-rate parameter for group membership—the ϕ parameter in the joke-reading case study and the ω parameter in the question-answering case study—as well as membership parameters z_i for each person. Models that included hierarchical extensions involved still other parameters. The joke-reading case study allowed for individual differences using the σ parameter. The question-answering case study included the number of answered pages p_i and presented questions q_i as parameters, and the ψ_k parameter representing the probability of a question on the k th page being answered. All of these parameters are psychologically interpretable, and potentially important in answering scientific and applied questions. Our modeling approach allows them all to be inferred simultaneously from a simple distribution of counts over the number of jokes read or questions answered. This capability to make rich inferences from limited data comes the

incorporation of detailed psychological assumptions about the data-generation process, and from the ability of Bayesian methods to make inferences for highly-structured models and sparse data.

The scientific approach to developing models we adopted in the case studies is a very pragmatic one. We did not use strong guiding theory to make prior predictions about the distribution of the data, and then test those predictions (Feynman, 1994). Instead, we used knowledge of the data to guide the incremental extension of the basic censored geometric model, seeking descriptive adequacy. We used more stringent model evaluation at key junctures, to test the most important conclusions. In the joke-reading case study, we used Bayes factors to quantify the evidence the data provided against individual differences in quitting probabilities, and against the possibility that there was a sub-group of people determined to read all of the jokes. In the question-answering case study, we used Bayes factors to provide evidence for the need for a non-linear decrease in the question-completion rate as more pages were accessed. We think this scientific process of model extension to improve descriptive adequacy and model evaluation to find evidence for successful extensions is a general and useful one.

It should be clear that our claims are not about the generality or predictive ability of the models of joke-reading and question-answering presented. Our claims are about the usefulness of the censored geometric as a base model, and the possibility of tailoring this model to specific problems. In the joke-reading case study, a simple censored geometric model proved descriptively adequate, and extensions allowing for various individual differences did not improve the descriptive adequacy. In the question-answering case study, a simple censored geometric model was descriptively inadequate. Extending the model to include a group of people who followed specific instructions to access 4 pages, and allowing for the question-completion rate to change as more pages were accessed, cumulatively resulted in a descriptively adequate model. In both case studies, the process of model development led to useful psychological insights into people's quitting behavior, in terms of individual differences, the existence of subgroups, or other properties of the data-generation process.

Our approach to modeling and inference is entirely Bayesian. We think this is the right choice conceptually, since Bayesian theory handles censoring in a principled way, and Bayesian parameter estimation and model selection are similarly principled. This means the Bayesian theoretical framework has capabilities the frequentist one does not. The ability to find evidence in favor of a “null” model, for example, is necessary to quantify the evidence against individual differences in the joke-reading case study, and requires a Bayesian approach. Just as importantly, we think Bayesian methods are the right choice pragmatically. The

ability to extend models in flexible and creative ways, using latent-mixture and hierarchical approaches, is an essential ingredient to the exploratory improvement of models demonstrated in our case studies. The JAGS implementation of the various models is straightforward, and makes inference based on computational Bayesian methods easy. Proposing an extended model requires little additional work in model implementation, which facilitates the creative and exploratory process of model development.

This paper dealt with the problem of modeling when people quit, and one particular class of models, in the form of censored geometric models, to understand quitting behavior. There are, of course, many other statistical and modeling contexts in which censoring is important. Being able to model censored data is a core capability in a statistician's toolkit, but it is hard to find applications in psychology. We do not think this is because psychology is bereft of modeling challenges involving censored data. The motivation for our case studies—environments with only a finite number of items, limiting the possibility of observing people's underlying preferences—is a good example. In these situations, basic psychological questions can only be answered by including the censoring properties of the environment in modeling and inference. Psychology as a field seems not to have considered these sorts of problems very often. It is possible the field regards them as uninteresting problems. It is more likely the field has ignored these problems because of ignorance of the statistical methods required to deal with censoring.

In this context, the primary goal of this paper is to demonstrate the feasibility of applying censored models to psychological problems. Censored statistical models are important in statistics, and in fields like biology, medicine, epidemiology, engineering, economics, and demography. We think they are equally important in the psychological sciences. Whenever the measurement of people's behavior is limited in some way, the modeling problem involves censoring. We hope that the pragmatic approach to model development and evaluation—and the conceptual clarity and practicality of Bayesian methods—that we have demonstrated will lead to greater exploration of psychological models involving quitting in censored environments.

Acknowledgments We thank David Condon for helpful information on the data collection method in the SAPA dataset. All of the JAGS code, the raw behavioral data, and other supplementary material, are available at a project page for this paper on the Open Science Framework at <https://osf.io/v6x8b/>. JV was supported by grants NSF#1534472, JTF#48192, and NSF#1230118.

References

- Amemiya, T. (1984). Tobit models: A survey. *Journal of Econometrics*, 24, 3–61.
- Bartlema, A., Lee, M. D., Wetzels, R., & Vanpaemel, W. (2014). A Bayesian hierarchical mixture approach to individual differences: Case studies in selective attention and representation in category learning. *Journal of Mathematical Psychology*, 59, 132–150.
- Condon, D., & Revelle, W. (2015). Selected personality data from the sapa-project: On the structure of phrased self-report items. *Journal of Open Psychology Data*, 3.
- Cox, D. R., & Oakes, D. (1984). *Analysis of survival data*. CRC Press.
- Feynman, R. (1994). *The character of physical law*. Modern Press.
- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2, 1360–1383.
- Goldberg, K., Roeder, T., Gupta, D., & Perkins, C. (2001). Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4, 133–151.
- Green, R. F. (1977). Do more birds produce fewer young? A comment on mayfield's measure of nest success. *The Wilson Bulletin*, 89, 173–175.
- Klein, J. P., & Moeschberger, M. L. (2005). *Survival analysis: Techniques for censored and truncated data*. Springer Science & Business Media.
- McWilliams, B., Tsur, Y., Hochman, E., & Zilberman, D. (1998). Count-data regression models of the time to adopt new technologies. *Applied Economics Letters*, 5, 369–373.
- Norris, J. L., & Pollock, K. H. (1998). Non-parametric mle for poisson species abundance models allowing for heterogeneity between species. *Environmental and Ecological Statistics*, 5, 391–402.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In Hornik, K., Leisch, F., & Zeileis, A. (Eds.) *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vienna, Austria.
- Revelle, W., Wilt, J., & Rosenthal, A. (2012). Individual differences in cognition: New methods for examining the personality-cognition link. In A. Gruszka, G. Matthews & B. Szymura (Eds.) *Handbook of individual differences cognition* (pp. 27–49). Springer.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cognitive Science*, 32(8), 1248–1284.
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: a tutorial on the savage-dickey method. *Cognitive Psychology*, 60(3), 158–189.