

Bayesian Cultural Consensus Theory

Zita Oravecz¹,
Joachim Vandekerckhove¹,
and William H. Batchelder¹

Field Methods
2014, Vol. 26(3) 207-222
© The Author(s) 2014
Reprints and permission:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/1525822X13520280
fm.sagepub.com



Abstract

In this article, we present a Bayesian inference framework for cultural consensus theory (CCT) models for dichotomous (True/False) response data and provide an associated, user-friendly software package along with a detailed user's guide to carry out the inference. We believe that the time is ripe for Bayesian statistical inference to become the default choice in the field of CCT. Unfortunately, a lack of publications presenting a practical description of the Bayesian framework in the context of CCT models as well as a dearth of accessible software to apply Bayesian inference to CCT data has so far prevented this from happening. We introduce the Bayesian treatment of several CCT models, focusing on the various merits of Bayesian parameter estimation and interpretation of results, and also introduce the Bayesian Cultural Consensus Toolbox software package.

Keywords

cultural consensus theory, Bayesian inference, user-friendly software, dichotomous response data

¹ Department of Cognitive Sciences, University of California, Irvine, CA, USA

Corresponding Author:

William H. Batchelder, Department of Cognitive Sciences, University of California, Irvine, 2129 Social Sciences Plaza A, Irvine, CA 92697, USA.

Email: whbatche@uci.edu

Introduction

Cultural consensus theory (CCT; Romney and Batchelder 1999) refers to a family of models that enable researchers to learn about informants' shared cultural knowledge by taking advantage of modern techniques in signal detection theory and psychometrics. Since its publication in a seminal 1986 paper by Romney et al. (1986), CCT has seen hundreds of applications. An overview of these studies, including possible related concerns, can be found in Weller (2007). Most of these applications have concerned dichotomous (True/False) response data. Central to the success of CCT as an option for ethnographic researchers is the availability of a software tool to perform statistical inference for CCT models. Such tools are currently available in two separate software packages: ANTHROPAC (Borgatti 1996) and UCINET (Borgatti et al. 2002). Unfortunately, these programs have several limitations in their ability to handle certain CCT models. Also, as they do not rely on the Bayesian inference framework, their output is limited as to how they represent the knowledge about the model parameters obtained from the data.

CCT investigates shared cultural knowledge and belief systems by administering questionnaires to informants who share this knowledge. Typically, several informants are asked multiple questions about a cultural subject and they respond True or False. The standard CCT model for such data was introduced by Romney et al. (1986), and named the General Condorcet Model (GCM) in Batchelder and Romney (1988). The basic and still widely used version of this model (recent examples are Hopkins 2011 and Miller 2011) provides us with estimates of the informants' cultural competence (level of cultural knowledge) and with the culturally correct (consensus) answers. The standard CCT model is based on a formal mathematical model from signal detection theory (Macmillan and Creelman 2005) that takes into account the possibility of a guessing mechanism as well.

The GCM has been augmented in several ways. Batchelder and Romney (1988) improve its parameter estimation method and in most of the derivations a respondent-specific guessing bias parameter is incorporated. However, both software packages currently used for CCT analysis (ANTHROPAC and UCINET) make the same simplifying neutral guessing bias assumption to enable estimation of competencies and culturally correct answers (see equation 1). As a result, it has not been possible to investigate variability in guessing bias in ethnographic studies. A further model extension introduced by Batchelder and Romney (1988) allows item difficulty to be heterogeneous. This involves specifying cultural competence as a function of the respondent's

ability and the item's difficulty by the Rasch (1960) model, which is a central model in psychometric test theory. A general estimation framework for such extended models was developed much later by Karabatsos and Batchelder (2003). They take advantage of the Bayesian framework (Gelman et al. 2004) and derive statistical inference for the most general GCM: one that allows heterogeneity in the competencies, informant abilities, and item difficulties, as well as for some constrained versions.

The Bayesian GCMs described in Karabatsos and Batchelder (2003) have so far not been featured in any substantive ethnographic investigation. Likely one reason for that is the lack of practical exposure to the Bayesian framework. This article presents a Bayesian inference framework for CCT models for dichotomous (True/False) response data and provides an associated, user-friendly software package to carry out the inference. The software, called Bayesian Cultural Consensus Toolbox (BCCT), is freely accessible on the first author's website (bayesian.zitaoravec.net) along with the *User's Guide for Bayesian Cultural Consensus Toolbox* document, also available online as an ancillary material on the *Field Methods*' website.

We start with summarizing the properties of the GCM, then we introduce the Bayesian parameter estimation framework and highlight the novelty of the results that can be obtained from BCCT.

Properties of the General Condorcet Model

CCT assumes that cultural truth or shared knowledge resides in the agreement of the members of a culture. A typical data set analyzed by CCT is collected from informants (respondents) who are assumed to constitute a single culture. As part of a standard analysis, one should always investigate whether this assumption holds, and a later section provides a way to do this in the context of Bayesian inference.

The type of data set that we analyze in this article consists of "true" and "false" answers from N respondents on M items. The CCT model for dichotomous responses is called the General Condorcet Model. Respondents are indexed with i , and items with k . Then the data set is coded as:

$$Y_{ik} = \begin{cases} 1 & \text{if } i \text{ responds 'true' to item } k \\ 0 & \text{if } i \text{ responds 'false' to item } k \end{cases}$$

and exhibited as a N by M array, $\mathbf{Y} = (Y_{ik})_{N \times M}$.

Accordingly, the culturally correct answers to be estimated are denoted as:

$$Z_k = \begin{cases} 1 & \text{if item } k \text{ is 'true'} \\ 0 & \text{if item } k \text{ is 'false'} \end{cases}$$

and the entire answer key by $\mathbf{Z} = (Z_k)_{1 \times M}$. The GCM relies on a formal model called the two high-threshold model from signal detection theory. The rationale behind it is that when an informant i is faced with a question, they either know the (culturally) correct answer with probability D_i or they do not. In the latter case, they guess “true” with a certain probability denoted by g_i for informant i . This parameter is defined on the interval $[0, 1]$ and the closer g_i is to 1, the more likely it is that informant i will guess “true.” However, the most basic model used in ANTHROPAC and UCI-NET sets this parameter to 0.5 for all informants, as described in Romney et al. (1986). According to this model, the probability of respondent i answering “true” to question k can be written as:

$$p(Y_{ik} = 1) = Z_k[D_i + (1 - D_i) 0.5] + (1 - Z_k)(1 - D_i) 0.5 = Z_k D_i + (1 - D_i) 0.5, \quad (1)$$

where D_i is commonly labeled the “informant specific competence.” As a first step toward extending the simplified model, the constraint on the guessing bias can be relaxed. Additionally, the competence parameter D_i can be indexed by item k to facilitate further extensions. These two simple modifications of equation 1 are represented as:

$$p(Y_{ik} = 1) = Z_k D_{ik} + (1 - D_{ik}) g_i. \quad (2)$$

The informant-specific guessing bias g_i allows respondents to differ in their response tendencies.

As a next step, we suppose that the questionnaire items that tap some domain of shared cultural beliefs might have differential cultural saliency; that is, not be of equal difficulty. Hence, we fit a Rasch measurement model to the competence parameter, as described in Karabatsos and Batchelder (2003). Consequently, the competence of an informant i for an item k can be modeled as the function of the informant’s ability parameter θ_i and the item’s difficulty δ_k :

$$D_{ik} = \frac{\theta_i(1 - \delta_k)}{\theta_i(1 - \delta_k) + \delta_k(1 - \theta_i)}. \quad (3)$$

In this specification of the Rasch model, θ_i and δ_k can take values between 0 and 1. As can be seen from equation 3, the probability of knowing the culturally correct answer (i.e., D_{ik}) increases with increasing ability (θ_i) and decreases with increasing item difficulty (δ_k).

As a summary, the extended GCM has four classes of parameters. Three of them, namely the informant-specific ability θ_i , the informant-specific guessing bias g_i , and the item-specific difficulty δ_k parameters, can take values anywhere between 0 and 1. The fourth class, namely the answer key parameters Z_k , can take only the values of 0 and 1. Together, the most complex model has $2N + 2M$ parameters and $N \times M$ data points, and as CCT is typically not run on data where $N < 4$ or $M < 5$, the model will have more data points than parameters.

Constraints can be introduced for the first three classes of parameters. As mentioned previously, the basic GCM assumes homogeneous item difficulty and, most often, neutral guessing bias. If the guessing bias is assumed to be neutral (consequently homogeneous among respondents), then:

$$g_i = g = 0.5, \quad (4)$$

and informants are equally likely to guess “true” or “false.” We can constrain the item difficulty in the same manner to have a neutral, homogeneous value across all items:

$$\delta_k = \delta = 0.5. \quad (5)$$

If we substitute 0.5 for δ_k into equation 3, most of the terms cancel out and we are left only with $D_{ik} = \theta_i$, that is, $D_i = \theta_i$. Therefore, in models where we do not allow for heterogeneous item difficulty (i.e., $\delta = 0.5$), we keep the term “competence” (for D). In models with heterogeneous item difficulties, the term “ability” will be used (for θ). By applying both constraints (equations 4 and 5), we arrive back to the basic version of the model as described in equation 1.

Finally, the model can be further constrained by assuming that informants show no variability in their ability parameters, formally:

$$\theta_i = \theta. \quad (6)$$

Equations 4, 5, and 6 describe constraints for three classes of parameters. Therefore, combining the heterogeneous and constrained parameter settings leaves us with $2^3 = 8$ models. All of these models are implemented in the BCCT software package that we introduce later in this article.

Introduction to Bayesian Parameter Estimation from the Perspective of CCT Models

In this section, we describe the essence of the Bayesian approach to model inference that is becoming more and more popular in the social and

behavioral sciences. Although nowadays the majority of statistical methods for the social sciences are based on the classical frequentist statistical viewpoint, this is changing rapidly (e.g., Gill 2002; Jackman 2009; Kruschke 2011; Lee and Wagenmakers 2013). In the next paragraphs, we argue that social science researchers would achieve great benefits by adopting the Bayesian framework for drawing statistical inferences for CCT models.

Let us consider the case where we want to draw inferences regarding some particular informant's competence parameter. The classical estimation methods as implemented in ANTHROPAC or UCINET would only provide us with a point estimate (a single numerical value) for this parameter. In the Bayesian framework, parameters are considered random variables that have a probability distribution. Bayes' theorem capitalizes directly on the axioms of probability theory. Henceforth, for example, when deriving an estimate for a particular informant's competence parameter (D_i), the Bayesian framework provides us with a posterior probability distribution of the parameter values that reflects our knowledge of the parameter given the data. Based on this posterior probability distribution, we can report a point estimate for the parameter such as its mean or median. Additionally, we can make straightforward inferences about the probability that an informant's specific competence parameter lies in a certain region. To illustrate this, see Figure 1, which is part of the output of our Bayesian inference program (BCCT). It shows that for each parameter not only point estimates like the mean and median, but also a standard deviation of the posterior distribution and 2.5% and 97.5% percentiles are provided (the "R-hat" (\hat{R}), and AC values are discussed in detail in the user's guide). These percentiles can simply be interpreted as the region that contains the parameter with probability 0.95 and we call this interval a 95% Bayesian credible interval (BCI). We will discuss Figure 1 in more detail in a later section.

We are able to make straightforward probability statements about parameter estimates because in the Bayesian framework, the knowledge about parameters is always described in terms of probability distributions. The starting point in Bayesian inference is our prior knowledge, or the *prior probability distribution*, of the model parameters. In our example, we have to decide what we consider a reasonable distribution of the competence parameter of an informant before we factor in the actual data. Most often, we will have no relevant information about particular respondents' ability parameters, so any a priori statement about their abilities must be necessarily vague. Indeed, this is in the spirit of goals of ethnographic research, which is not to impose prior beliefs on the process of discovering cultural

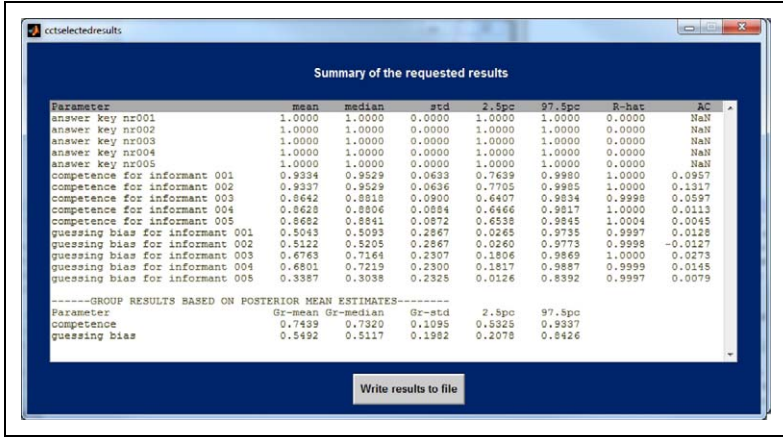


Figure 1. Screenshot of the BCCT window displaying summary of posterior statistics on selected parameters.

knowledge. Fortunately, there exist formal mathematical expressions of uncertainty for the Bayesian framework. For example, as the ability parameter in the CCT framework is, by definition, constrained to the range [0, 1], and if we have no a priori preference for any particular value, the probability density to express indifference between possible parameter values is the uniform distribution on the [0, 1] interval. We can, of course, make the same assumption independently about every informant. In the same vein, it makes sense to express similar uncertainty about each participant’s guessing bias parameter. Moreover, we would typically begin with noninformative prior distributions for the item difficulties and answer key items as well.¹ When we collect consensus data on a certain topic, that information is used to update all these initial prior distributions. The more data one acquires, the less influential the priors on the model parameters become.

Bayes’ theorem is the mathematical formalism used for updating prior knowledge regarding model parameters with the information provided by the observed data. Knowledge about any particular parameter will still be expressed as a probability distribution. The probability distribution of a parameter conditional on the data (i.e., after having observed the data) is known as the posterior distribution mentioned previously. Let us start with the example of how to derive the answer key items when competencies and guessing bias are already known. As described in Romney et al. (1986), by using Bayes’s theorem, the probability that a particular answer key Z is correct is given by the following equation:

$$p(\mathbf{Z}|\mathbf{Y}) = \frac{p(\mathbf{Y}|\mathbf{Z})p(\mathbf{Z})}{p(\mathbf{Y})}, \quad (7)$$

where \mathbf{Z} is the vector answer key and \mathbf{Y} is the data, as defined previously. On the left-hand side, $p(\mathbf{Y}|\mathbf{Z})$ is the posterior distribution of the possible answer keys: for example, if there are M items, there are 2^M possible values of \mathbf{Z} . On the right-hand side, $p(\mathbf{Y}|\mathbf{Z})$ is the *likelihood* of the data given the answer key, while $p(\mathbf{Z})$ expresses our prior knowledge of the answer key values. Most often, no relevant a priori information is at hand, so it is typically assumed that all consensus answer keys are a priori equally likely. The denominator $p(\mathbf{Y})$ is a numerical normalizing constant that expresses the probability of observing the data, averaged and weighted over all possible combinations of parameter values. The goal of Bayesian inference is to obtain the posterior distribution of the answer key values in order to draw inferences about it. Equation 7 is essentially the same as equation 24 in Batchelder and Romney (1988) accompanied with further derivation incorporating the competencies and guessing biases, and these results along with the constraint that $g = 0.5$ are implemented into the so-called matching method in ANTHROPAC and UCINET.

The same logic as in equation 7 can be applied to all parameters of the GCM, not only the answer key parameters. If we collect all model parameters in the vector ξ , the joint posterior distribution of all the parameters can be expressed as:

$$p(\xi|\mathbf{Y}) \propto p(\mathbf{Y}|\xi)p(\xi),$$

where, as mentioned earlier, \mathbf{Y} denotes the data. This formula is very similar to that in equation 7, except that we no longer explicitly mention the normalizing constant in the denominator and replace the equality sign with the proportionality sign, \propto . Computational statisticians have developed numerical sampling algorithms for estimating the posterior distribution, even when we can describe it only up to a constant of proportionality (Gilks et al. 1996; Robert and Casella 2004). BCCT relies on these algorithms to estimate the joint posterior distribution of all the GCM parameters.

In practice, when we have estimated the joint posterior distribution of all the parameters, we often want to explore the marginal probability density of each parameter. Turning back to our example, let us say, for example, that we want to see how likely it is that the first informant's competence parameter (D_1) is smaller than a certain value, like 0.7. For that, we explore the

marginal posterior density of that competence parameter, which is simply the full posterior distribution (the product of the prior and likelihood distributions) averaged over all parameters except the competence parameter in question. Sometimes if the statistical model is very simple, this multiplication leads to a known distributional form of the posterior density, in which case the probability that the parameter is less than some particular value is given immediately by the posterior cumulative density function (CDF) at that value. However, in most cases such as ours, the posterior distribution has a mathematical form for which this probability is difficult or impossible to compute directly. The standard strategy in this common situation is to use modern computational tools, called Markov Chain Monte Carlo (MCMC) sampling methods, to develop a sampler that can generate many random values from the posterior distribution, and use the sampled values to approximate the posterior distribution. An extensive literature exists (see, e.g., Robert and Casella 2004; Gamerman and Lopes 2006) regarding simulation procedures to draw samples from any kind of distribution. Also, there are freely available software packages that automate many of the steps required to run such a sampling procedure. However, these programs require the computational power of a modern computer to do the sampling in a timely manner.

Once we have explored the posterior density of the competence parameter in question (i.e., we have generated an ample amount of samples from its posterior density), several statistics can be calculated. For example, measures of central tendency, such as the mean and median can easily be attained. Posterior standard deviations can also be calculated to express uncertainty around these point estimates. Also, as mentioned earlier, Bayesian credible intervals can be constructed. Such BCIs designate a region where the true parameter lies with a certain probability, typically with probability 0.95. In addition, if we want to evaluate how likely it is that the competence parameter is smaller than a given value, say 0.7, computing this probability from the posterior density simply amounts to counting the proportion of sampled values for the parameter that fell below 0.7. The straightforward way of interpreting uncertainty in parameter estimation can be especially appealing for the social sciences, where parameters are often associated with personality traits, attitudes, skills, behavioral properties, and so on, where it makes particular sense to quantify uncertainty in inferences in an easily interpretable fashion.

In sum, the Bayesian statistical framework offers a rational and coherent way of drawing inferences. Despite its prevalence among computational statisticians, its use is not yet as widespread in the social and behavioral sciences as the classical inference framework. There are several reasons for

that. Most of all, as mentioned previously, in many cases, Bayesian parameter estimation using computational MCMC samplers requires some magnitude of computational power, which has become available only a couple of decades ago. Second, until recently, there have been only a few books and articles written especially for social scientists describing the logic behind using MCMC samplers to estimate the posterior distribution. The supplementary material to this article (user's guide) provides more technical details about sampling from the posterior.

Introducing the Bayesian Cultural Consensus Toolbox

We have written a user-friendly software application called the Bayesian Cultural Consensus Toolbox (BCCT) for conducting inference for various versions of the GCM for dichotomous response data discussed in earlier sections. BCCT is a freely distributed, stand-alone application that communicates with the user through a graphical user interface (similar to UCINET) and does not require any programming knowledge from the user. For carrying out the sampling from the posterior for the different GCMs, BCCT relies on a generic software package that deals with Bayesian statistical inference, namely JAGS (Plummer 2011). Moreover, a MATLAB engine is used to carry out the necessary computations as well as for supporting the user interface.² A straightforward manual entitled *User's Guide to Bayesian Cultural Consensus Toolbox* provides the researcher with all the information necessary to use the software.³

To demonstrate the merits of Bayesian CCT analysis, we make use of a data set published in Romney (1999). The data set was collected by Weller (1984) and analyzed by several authors such as Romney et al. (1986) and Batchelder and Anders (2012). The original study measured knowledge about contagiousness of certain diseases in a Guatemalan population. The sample consists of 24 informants who provided dichotomous responses to 27 items about a disease being contagious or not. As default in BCCT, vague priors are set for each parameter of interest, namely a uniform distribution for parameters with ranges between 0 and 1 and a Bernoulli distribution with parameter 0.5 for the answer key. Here we skip further technical details of how to carry out the analysis, as it is described in detail in the user's guide and we focus on the novel aspects of the results achieved through Bayesian CCT analysis.

We analyzed the previously mentioned data by a GCM in which persons were allowed to differ in their ability as well as in their guessing bias. Figure 1 shows a sample of the results on these model parameters. Each line

corresponds to the statistics of one of the model parameters named in the first column. The second and third columns show the posterior mean and median, which can be thought of as point estimates of the parameters. The fourth column displays the uncertainty in the parameter estimates in terms of posterior standard deviations. The fifth and sixth columns concern the spread of the posterior density by specifying its 2.5th and 97.5th percentiles, or in other words the 95% Bayesian credible interval. Then \hat{R} and AC in the next two columns show diagnostic values of the success of the MCMC sampler discussed more in details in the user's guide. The results displayed in Figure 1 can be written out into a standard text file by clicking the button below the table.

Looking at the first line for a competence parameter in Figure 1, we can see that a possible point estimate for the competence parameter of informant 001 is 0.93. Also, its 95% BCI lies above 0.7, so it is very likely that this informant's competence level is above 0.7. For deriving how probable that is exactly, we go back to the main window, select our parameter of interest, and press the Plot cumulative distribution button. It then displays its cumulative posterior density curve, as shown in Figure 2. Since this plot is interactive, by pointing at the particular parameter value on the x -axis, the plot shows us its place on the curve and displays numerically how likely it is that the parameter value is lower than that value. For $D_1 < 0.7$, it shows 0.0065, so it is rather unlikely that this informant has a competence value below this value.

The parameter selection in Figure 1 is only an example: The user can set all of the parameters or any subset of them. Moreover, at the end of the displayed table, there is always some summary on group results as we can see it in the lower part in Figure 1. They are based on the point estimates (posterior means) of the competence (or ability), guessing bias and item difficulty (if applicable) parameters, and their summary statistics such as mean, median, standard deviation, and 2.5th and 97.5th percentiles are also calculated. This way the user can have an idea about the average level of these parameters in the sample (mean, median), as well as about how much variability there is among informants and items (standard deviation, percentiles). These quantities would be the usual sorts of statistics about the competence level of a group of informants that one would publish in a research paper.

For example, we see that the mean competence level is 0.74 and the standard deviation is 0.11. In the case of the guessing bias, the mean is 0.55, which is close to the neutral guessing value of 0.50 assumed in earlier analysis of the contagion data by ANTHROPAC. However, the standard deviation of the guessing bias is 0.20, which suggests that there is a lot of heterogeneity in guessing bias in the group of informants. Certainly, one can look not only at the point estimates but the BCIs as well, and calculate,

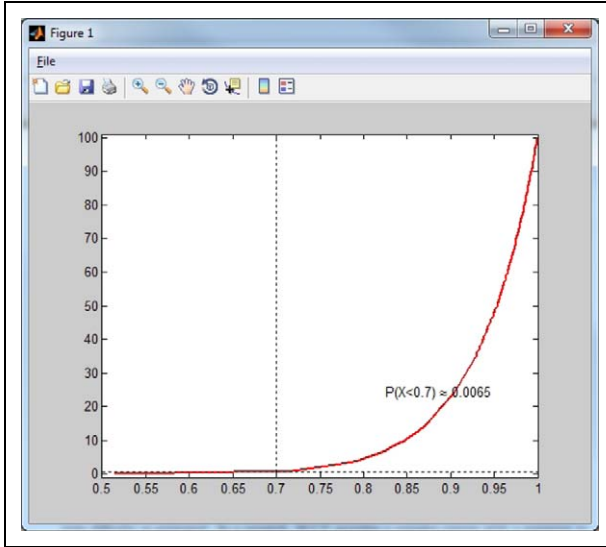


Figure 2. Screenshot of the cumulative distribution plot generated by BCCT.

for example, how many informants are most likely to have competency above 0.7, or how many of them will most likely always answer “False,” and so on. Moreover, although it is not displayed in the current example, ethnographers might also find it interesting to compare item difficulty levels for the different questions of the questionnaire, when the GCM with item difficulty is estimated. In a nutshell, BCCT provides a complex output with a potential to investigate various research questions.

Checking Model Fit

With BCCT, we can carry out several goodness-of-fit tests. Their detailed description as well as some guidelines for model selection can be found in the user’s guide, here we will only focus on the posterior predictive check of the one underlying culture assumption.

Posterior Predictive Check of the one Underlying Culture Assumption

The usual rule of thumb used by ethnographers (Weller 2007) for supporting the single culture assumption is that this ratio of the first and the second eigenvalues is above 3. This rule is actually problematic, as the ratio is very

dependent on the overall level of competence on the group of informants. The posterior predictive check option for the single culture assumption in the BCCT software improves on the classical eigenvalue test based on the recent work of Batchelder and Anders (2012). With BCCT, we can investigate the posterior predictive distribution of the ratio of the first to second eigenvalues of several simulated data sets based on the posterior samples of model parameters. Then we can check whether the data eigenvalue ratio is within the 95% posterior predictive distribution of the model based eigenvalue ratios. If not, BCCT returns an error message saying: “PPC failed”; otherwise, it returns “PPC passed.” By pushing the button under the statistic, we can see the smoothed posterior predictive distribution of these simulation based eigenvalues with the data’s eigenvalue ratio represented by a straight line. In this plot, the user can see the exact place of the data eigenvalue ratio (straight line) in the distribution of eigenvalues ratios based on simulated data.

Conclusions

CCT has been used extensively in field research in the last couple of decades. We believe that new developments in the statistical theory and computer science regarding the implementation of computational Bayesian inference will help us expand range of models that researchers can exploit for studying various topics in cultural consensus. Therefore, we provided an introduction into the Bayesian modeling framework that is gaining more and more traction in several fields. Although we consider the current encounter with Bayesian statistics sufficient for ethnographers to deal with Bayesian CCT inference, we encourage readers to further investigate the possibilities of using the Bayesian framework.

The Bayesian approach allows one to examine properties of the posterior distribution of a parameter. In addition, in cases with a small number of informants, the extra flexibility in allowing heterogeneity in the guessing process and in the item difficulty may contribute to better reconstruction of the answer key than the classical approach to estimation. Finally, the matching and covariance methods in ANTHROPAC and UCINET are specific to dichotomous true/false or multiple choice response data. The Bayesian approach to inference can be developed for any CCT model, although for more complex models more complex MCMC samplers may need to be developed (e.g., Batchelder *et al.* 2010).

We intend the current software package to be a starting point for the development of a larger variety of user-friendly BCCT tools for CCT

models for questionnaires that require other types of responses (e.g., Likert, ranking, matching, and continuous response scales). Based on feedback from users we intend to produce more complex BCCT packages that will enable researchers to investigate even more complex CCT models. More complex models could include, for example, population distributions for person-specific or item-specific parameters (i.e., random effects), or the possibility to introduce covariate information.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: We gratefully acknowledge support from research grants from the Air Force Office of Scientific Research (AFOSR) and the Army Research Office (ARO) to Batchelder, PI.

Notes

1. In principle, we could incorporate informative prior beliefs about model parameters, for example, distributions centered on a certain value (in our setting, beta distributions would be a suitable choice). In case of a longitudinal setting, for example, when the same informants are asked several times about the same topic, one might use information from one study to express prior information a subsequent study. Users interested in that option can implement this by making changes in the modeltxt.m file, although in this case BCCT should be run from MATLAB and not by BCCT.exe.
2. All MATLAB scripts that constitute the program are automatically downloaded as part of the BCCT package, and can be found in the folder called "Supplementary files." Users who do have a MATLAB license can use these files to run BCCT in its native environment.
3. As mentioned earlier, it is available on the *Field Methods*' website as well as on bayesian.zitaoravec.net.

Supplemental Material

The online data supplements are available at <http://fm.sagepub.com/supplemental>.

References

- Batchelder, W. H., and R. Anders. 2012. Cultural consensus theory: Comparing different concepts of cultural truth. *Journal of Mathematical Psychology* 56: 316–32.
- Batchelder, W. H., and A. K. Romney. 1988. Test theory without an answer key. *Psychometrika* 53:71–92.
- Batchelder, W. H., A. Strashny, and A. K. Romney. 2010. Cultural consensus theory: Aggregating continuous responses in a finite interval. In *Social computing, behavioral modeling and prediction*, eds. S. K. Chai, J. J. Salemo, and P. L. Mabry, 98–107. New York: Springer.
- Borgatti, S. P. 1996. *Anthropac 4.0*. Natick, MA: Analytic Technologies.
- Borgatti, S. P., M. G. Everett, and L. C. Freeman. 2002. *Ucinet for windows: Software for social network analysis*. Harvard, MA: Analytic Technologies.
- Gamerman, D., and H. F. Lopes. 2006. *Markov chain monte carlo: Stochastic simulation for Bayesian inference*. 2nd ed. Boca Raton, FL: Chapman & Hall/CRC.
- Gelman, A., J. Carlin, H. Stern, and D. Rubin. 2004. *Bayesian data analysis*. New York: Chapman & Hall.
- Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1996. *Markov chain monte carlo in practice*. London: Chapman & Hall.
- Gill, J. 2002. *Bayesian methods: A social and behavioral sciences approach*. Boca Raton, FL: Chapman & Hall/CRC.
- Hopkins, A. 2011. Use of network centrality measures to explain individual levels of herbal remedy cultural competence among the Yucatec Maya in Tabi, Mexico. *Field Methods* 23:307–28.
- Jackman, S. 2009. *Bayesian analysis for the social sciences*. New York: John Wiley.
- Karabatsos, G., and W. H. Batchelder. 2003. Markov chain estimation methods for test theory without an answer key. *Psychometrika* 68:373–89.
- Kruschke, J. K. 2011. *Doing Bayesian data analysis: A tutorial with R and BUGS*. New York: Academic Press.
- Lee, M. D., and E. -J. Wagenmakers. 2013. *Bayesian cognitive modeling: A practical course*. Cambridge: Cambridge University Press.
- Macmillan, N. A., and C. D. Creelman. 2005. *Detection theory: A users guide*. 2nd ed. Mahwah, NJ: Lawrence Erlbaum.
- Miller, E. 2011. Maternal health and knowledge and infant health outcomes in the Ariaal people of northern Kenya. *Social Science & Medicine* 73:1266–74.
- Plummer, M. 2011. Rjags: Bayesian graphic models using MCMC. *R package version 2.2.0-3*. <http://CRAN.R-project.org/package=rjags> (accessed February 27, 2012).
- Rasch, G. 1960. *Probabilistic models for some intelligent and attainment tests*. Copenhagen: Danish Institute for Educational Research.

- Robert, C. P., and G. Casella. 2004. *Monte Carlo statistical methods*. New York: Springer.
- Romney, A. K. 1999. Cultural consensus as a statistical model. *Current Anthropology* 40:103–15.
- Romney, A. K., and W. H. Batchelder. 1999. Cultural consensus theory. In *The MIT encyclopedia of the cognitive sciences*, eds. R. Wilson and F. Keil, 208–9. Cambridge, MA: The MIT Press.
- Romney, A. K., S. C. Weller, and W. H. Batchelder. 1986. Culture as consensus: A theory of culture and informant accuracy. *American Anthropologist* 88:313–38.
- Weller, S. C. 1984. Cross-cultural concept of illness: Variation and validation. *American Anthropologist* 86:341–51.
- Weller, S. C. 2007. Cultural consensus theory: Applications and frequently asked questions. *Field Methods* 19:339–68.